

The VesselGlyph: Focus & Context Visualization in CT-Angiography

Matúš Straka[¶]
Commission for Scientific
Visualization, Austrian Academy
of Sciences

Arnold Köchl[£]
Medical University Of Vienna

Michal Červeňanský[‡]
Comenius University, Bratislava

Miloš Šrámek^{*}
Commission for Scientific
Visualization, Austrian Academy
of Sciences

Dominik Fleischmann[§]
Stanford University Medical
Center

Alexandra La Cruz[□]
Institute of Computer Graphics and
Algorithms, Vienna University of
Technology

Eduard Gröller[¥]
Institute of Computer Graphics and
Algorithms, Vienna University of
Technology

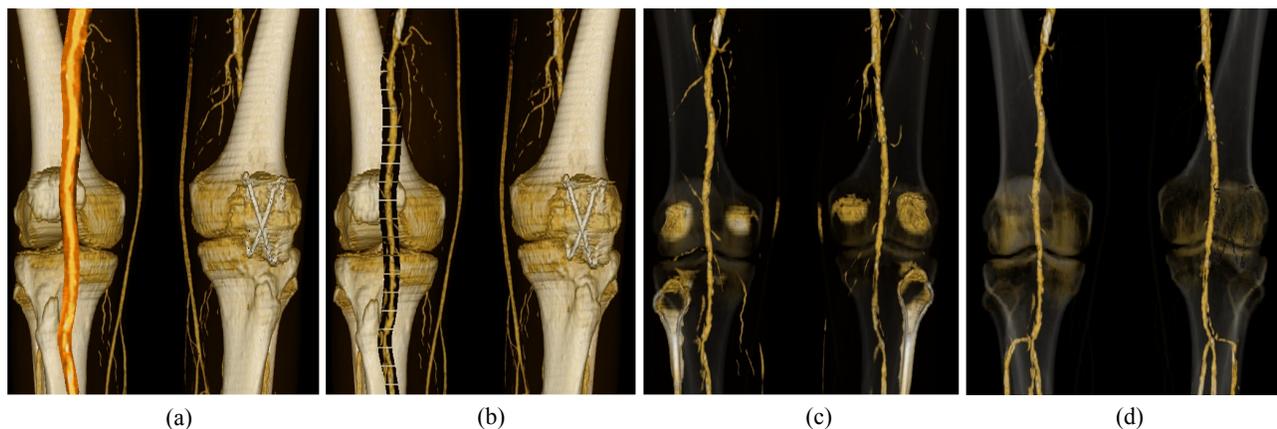


Fig. 1. Visualization using the Vessel Glyph: (a) CPR + DVR, (b) foreground-cleft in DVR with occlusion lines, (c) ThickSlab rendering (DVR), (d) tubular rendering (DVR)

Abstract

Reliable and complete blood-vessel segmentation is still a challenging problem. This is especially true in the presence of morphologic changes resulting from atherosclerotic diseases. In this paper we take advantage of partially segmented data with approximately identified vessel centerlines to comprehensively visualize the diseased peripheral arterial tree.

[¶] matus.straka@assoc.oeaw.ac.at

[‡] mcervenansky@hotmail.com

[□] alacruz@cg.tuwien.ac.at

[£] arnold.keochl@univie.ac.at

^{*} milos.sramek@oeaw.ac.at

[¥] groeller@cg.tuwien.ac.at

[§] d.fleischmann@stanford.edu

We introduce the VesselGlyph as an abstract notation for novel focus & context visualization techniques of tubular structures such as contrast-medium enhanced arteries in CT-Angiography (CT-A). The proposed techniques combine direct volume rendering (DVR) and curved planar reformation (CPR) within a single image. The VesselGlyph consists of several regions where different rendering methods are used. Region type, the used visualization method and region parameters depend on the distance from the vessel centerline and on viewing parameters as well. By selecting proper rendering techniques for different regions, vessels are depicted in a naturally looking and undistorted anatomic context.

In this paper we furthermore present a way how to implement the proposed techniques in software and by means of modern 3D graphics accelerators.

CR Categories: J.3.2 [Life And Medical Sciences]: Medical information systems, I.3.3 [Computing Methodologies]: Computer Graphics—Picture/Image Generation

Keywords: focus & context technique, direct volume rendering, curved planar reformation, vessel visualization

1 Motivation

1.1 Introduction

Peripheral arterial occlusive disease (PAOD) is a disease that influences lower extremities of elderly people. It is characterized by formation of atherosclerotic plaque and vessel wall calcifications, which lead to luminal stenoses (narrowing) or complete occlusion of the involved arteries [Baum and Pentecost 1997]. Diminished blood flow to the legs causes restricted mobility, pain and even necrosis, eventually leading to death. Therefore, early diagnostics and proper treatment is needed.

Today's method of choice for vascular imaging is the Digital Subtraction Angiography (DSA), which produces in many cases high-resolution 2D images of good quality. A contrast agent used for DSA is active only for a very short period of time (a second or less). Therefore the agent needs to be injected closely to the point where the vessel tree is investigated. In the case of lower extremities this means an injection directly to the aorta, which might be dangerous for the patients and requires a relatively long stay in the hospital.

On the other side, Computed Tomography Angiography (CT-A) uses a contrast agent that is active for a much longer time (typically 20–40 seconds). This allows injections in points that are further away from the investigated area, e.g. in forearm vessels. The examination is less dangerous, requires a shorter stay in the hospital and finally is also cheaper. Furthermore, CT-A produces 3D data, which allows more possibilities for the visualization of vessels. Nowadays, CT-A of the peripheral arteries is performed with helical CT scanners. They produce more than 1200 slices of 512x512 pixels in grayscale. Individual volume-element is smaller than 1 mm³. This produces very high quality images, but also enormous amounts of information. The large data sets are difficult to process for the radiologist and to store (600 – 1000 MB per patient). The whole scanning procedure takes about 40 seconds. The post-processing and evaluation, however, take much longer. It is in the order of hours even for a well-trained operator using dedicated processing and visualization software. Such an evaluation is not feasible if several patients shall be investigated on a daily basis.

Previous research from our environment led to promising techniques for segmentation and visualization of peripheral CT-A datasets [Kanitsar et al. 2001; Kanitsar et al. 2002]. This allowed partial identification of main vessels in the 3D dataset in the form of vessel centerlines and also enabled the clinical application of dedicated visualization techniques and their further development.

1.2 Research goals

The objective of imaging the peripheral arterial tree in patients with PAOD is to detect, localize, and gauge the degree and length of vascular changes which cause diminished blood flow to the legs, in order to initiate the appropriate therapeutic measures. The culprit lesion in atherosclerosis is the so-called atherosclerotic plaque, which develops on the inner surface of an affected artery, protrudes into the vessel lumen and thus causes luminal narrowing (stenosis). A greater than 50% diameter reduction is generally referred to as a significant stenosis and – if associated with clinical symptoms – considered an indication for treatment. Significant flow obstruction of the main conducting arteries is often accompanied by the development of a network of collateral vessels, which maintain a basic blood supply to the tissue distal to the lesion. The presence of collateral vessels is thus an indirect

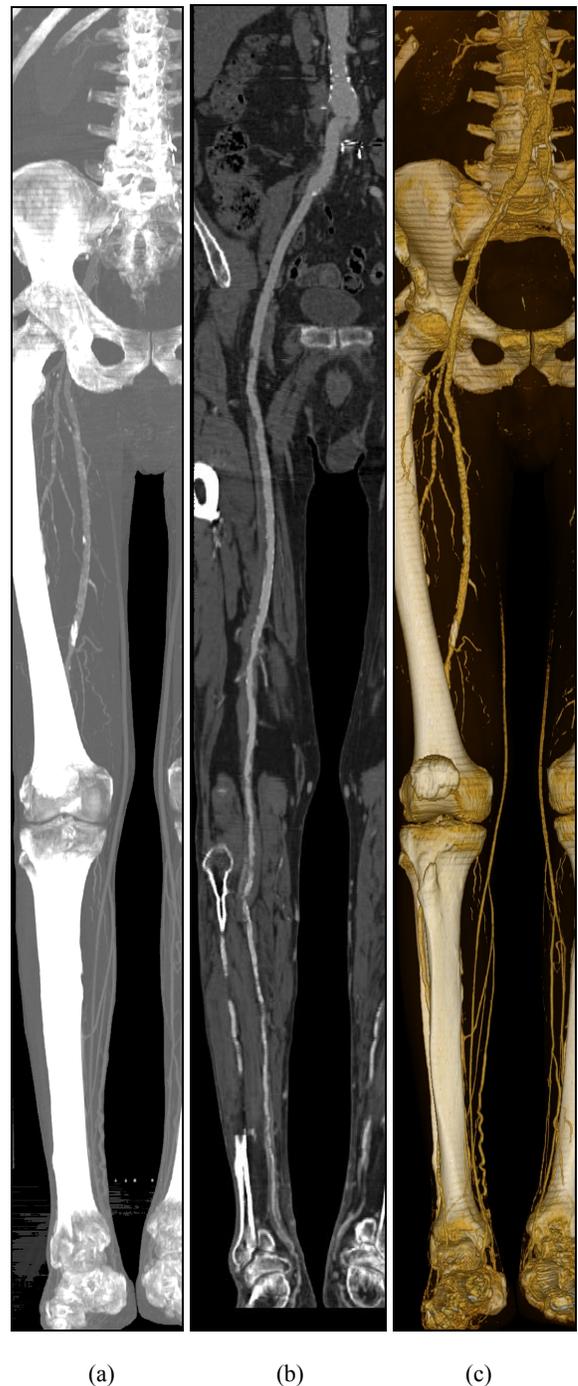


Fig. 2. Visualization of a CT-A dataset: (a) MIP image, (b) single-path CPR image through the aorta, right pelvic, femoral, popliteal, and posterior tibial arteries, (c) DVR image

sign of the presence of a hemodynamically significant reduction of blood flow. Atherosclerotic plaque may not only narrow, but also completely fill and occlude the vessel lumen. Long segments of complete occlusion may require a surgical bypass-graft to reestablish flow, whereas short segment luminal stenoses can often be treated by catheter techniques (balloon angioplasty). Initially, atherosclerotic plaque is of soft-tissue density (i.e. similar to muscle), and thus hypodense to the contrast-medium opacified intravascular blood stream. Over time, however,

atherosclerotic plaque calcifies and its CT density is then greater than contrast-medium opacified blood. This poses a significant problem for imaging diseased arteries with CT, because the hyperdense, calcified plaque obscures the vessel lumen when maximum intensity projection (MIP, [Rossnick et al. 1986]) or direct volume rendering (DVR, [Levoy 1988; Drebin et al. 1988]) is used. Because of the different therapeutic consequences it is clinically important though, to distinguish between calcified plaque, versus significant narrowing, versus complete occlusion. Another problem of MIP and DVR techniques for vessel visualization is that vessels are frequently obscured by bones (unless the bones are segmented out in a separate procedure). Both MIP (Fig. 2a) and DVR (Fig. 2c), however, provide excellent spatial perception for the observer, and provide bony landmarks for vessel segment identification, and also nicely display collateral vessels.

A reliable and automatic segmentation of diseased vessels in their entirety is still elusive. There are, however, robust algorithms that can approximately calculate the center paths of the major conducting arteries [Kanitsar et al. 2001]. The peripheral arteries of the lower limbs constitute a tree like structure and their centerlines represent a three-dimensional tree of paths. In the following these paths are referred to as the vessel tree.

The vessel tree can be used as an input function for the creation of curved planar reformations (CPR, [Kanitsar et al. 2002]) These longitudinal cross-sections through the arteries provide an excellent view on the vessel lumen, unobstructed by calcified plaque or overlying bony structures. A CPR is generated by resampling the 3D data set along a free-form cut surface as determined by the vessel centerline of the artery of interest. All other objects of the dataset (which may provide clinically useful information such as anatomic context, bony landmarks, or collateral vessel) are either distorted or not visualized at all (Fig. 2b). This limits the spatial perception of CPR images and thus their ability to communicate the findings to the treating physicians.

From the previous considerations it follows that the vessels as focus objects should be treated differently than the surrounding context data. The focus & context approach has been already widely used in information visualization [Shneiderman et al. 1999]. Visualization of vessels is a good example where this concept can be applied to the advantage of medical imaging. In this setting, the available information from the pre-segmented vessel tree can be used by our algorithms for defining the focus area in the dataset, which is treated differently from the rest of the data (context).

In this paper we introduce the VesselGlyph for focus & context rendering of the blood vessels. The VesselGlyph is an abstract representation aimed at incorporation of different rendering styles for the creation of an image. An example would be a CPR rendering of a vessel and its immediate neighborhood, combined with a DVR rendering of the surrounding data. The VesselGlyph serves several purposes. On the one hand it gives an abstract notation on which rendering styles are combined and shows how the transition between the styles is handled. The VesselGlyph also facilitates a systematic exploration of various combinations of rendering techniques which otherwise would have eluded the investigator. Furthermore, the VesselGlyph may serve as an interface element, which allows the user to change styles or parameters. Examples in this respect could be the type of CPR (CPR or thick-CPR), width of the transitional region between the focus and context areas and the context area rendering algorithms (DVR or MIP).

The VesselGlyph allows visualization and investigation of objects in areas where they would otherwise be obscured if volume visualization techniques alone would be used. Spatial configuration of the VesselGlyph that suppresses objects in front of the focus area may solve this problem.

This paper is organized as follows: in the Section 2 we give an overview of related work; in Section 3 we introduce the VesselGlyph concept and present applications thereof. In Section 4, we discuss implementation details of the concept using current graphics hardware. Finally, we summarize the results of our work and sketch possible future research and applications.

2 Related Work

Several techniques for vessel visualization are currently available. They range from very simple ones like MIP, through more complicated (CPR or multi-path CPR) to highly sophisticated techniques as triangulation of vessel walls and modeling of the vessels.

Kanitsar et al. [2002] developed various approaches based on the idea of curved planar reformation. This approach allows an unobscured investigation of the vessel lumen by creating a longitudinal section through an approximate vessel centerline. Eccentric plaques and calcifications are investigated by rotating the CPR around the central axis of the vessel. The limitation of CPR is that it is sensitive to incorrect vessel-centerline position, which may lead to so-called “pseudo-stenosis”. It may also fail to visualize very thin vessels due to the same problem. Another inherent disadvantage of CPR, however, is its ambiguous display of the anatomic context. The viewer cannot reliably identify the anatomic segment that is displayed in the CPR image, which is essential for clinical decision-making. Kanitsar et al. also developed extensions of the basic CPR concept, called thick-CPR and multi-path CPR. Thick-CPR performs MIP rendering in a close vicinity of the curved plane through the vessel centerline. Multi-path CPR incorporates CPR renderings of an entire vessel tree within one image, which improves spatial perception.

Hauser et al. [2001] developed two-level rendering of medical images. In their work, they fuse DVR and MIP techniques in one image. A density based pre-classification step determined which technique to use for which object.

Zhou et al. [2002] extended this approach and developed a system where a region in focus is defined through a “magnifying glass”. Within the magnifying glass, the data is rendered photorealistically, e.g. using DVR. Outside this region, data is rendered using non-photorealistic styles, e.g. only contours, etc.

3 The VesselGlyph

In this section we introduce the VesselGlyph for focus and context rendering of vessels. The VesselGlyph is an abstract notation describing the combination of different rendering styles within one image. Let us consider a single axial slice of a CT-A data set showing a round cross section of an artery oriented perpendicularly to the CT section. The focus area is close to this circular object. The context area is the area surrounding the vessel. The vessel path intersects the axial slice close to the center of the vessel. The VesselGlyph (Fig. 4) describes the spatial arrangement of focus and context regions, where different techniques are used for rendering with a possible smooth transition between them. Sweeping the VesselGlyph along the

curved vessel centerline would cover those regions of the data that constitute the focus and context subset in the original volume data. E.g., the VesselGlyph shown in Fig. 4 combines CPR rendering of the focus with a DVR rendering of the context, including a smooth transition between the two rendering styles. The CPR rendering is depicted as a horizontal line as CPR resamples the vessel along an infinitesimally thin free-form surface. The DVR style of the context is depicted as two symmetric grayish areas. The larger extent of these areas reflects the fact that DVR takes samples from a larger region of the data

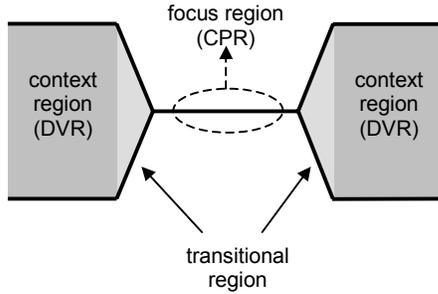


Fig. 4. VesselGlyph for Focus & Context rendering.

set. If ray casting is used for rendering then all samples along each ray are taken into account. The VesselGlyph further describes a smooth transition from CPR to DVR. Close to the focus region only a thin slab of data samples is considered in the rendering stage. Close to the context region more data samples along an individual ray are included. Fig. 3 gives examples of several possible configurations of the VesselGlyph. This figure also shows that previous vessel visualization techniques (DVR, CPR) can be considered as special cases of VesselGlyphs.

The VesselGlyph can therefore be regarded as an extension of known approaches. Fig. 3a represents the VesselGlyph that describes standard DVR. It provides excellent context visualization, but as no focus region is defined, the vessels of interest may be obscured by overlying bony structures. Fig. 3b describes the CPR rendering, which provides a detailed view on the vessel and its flow channel. Data far away from the vessel center are intersected by the thin curved surface, which does not provide a good context overview. Fig. 3c and Fig. 3d show a combination of CPR for the vessel and DVR for the context. The abrupt transition in Fig. 3c produces discontinuities in the resulting image, which makes it easy to distinguish between focus and context. In Fig. 3d there is a transitional region that ensures a smooth change from focus to context. In Fig. 3e a combination of a thick-CPR in the focus and DVR in the context is illustrated. A thick-CPR in this case may not only implement MIP, but can also aggregate the data samples within the slab using shaded or unshaded DVR. Thick CPRs alleviate the problem of imprecise vessel centerlines, which is especially important for small vessels. The VesselGlyph in Fig. 3f describes a CPR rendering where the information in the context area is fading out. Data far away from the vessel center are thus not rendered. This ensures that the important focus information immediately stands out in the result images. Fig. 3g describes a Foreground-Cleft VesselGlyph, where DVR is used both for focus and context. To allow a clear view of object in focus, the VesselGlyph ensures that objects in front of the focus region are suppressed. The Foreground-Cleft VesselGlyph can be completed with “occlusion lines” indicating front-to-back spatial ordering of objects (Fig. 3h).

In the following we investigate the properties of images generated by various configurations of the VesselGlyph, with

special attention to their ability to display vascular detail and anatomic context simultaneously, using a peripheral CT-A dataset obtained from a patient with peripheral arterial occlusive disease.

3.1 CPR + DVR VesselGlyph

CPR images are very important for vessel investigation in the clinical environment. The radiologists prefer them mainly for their ability to show the vessel interior - to see whether the calcification occludes the vessel partially or fully. On the other side, CPR images lack context information. Therefore we propose to use CPR only in the focus region and apply DVR in rendering the context area (Fig. 3c). There is an abrupt change between the two regions (Fig. 5).

Visualization with the CPR+DVR VesselGlyph can be performed in color or in grayscale. Color images create higher contrast between different tissues, while the grayscale images possess 12-bit grayscale information. This is important in the clinical environment, because the radiologists are used to “window” the investigated images. This window shows only relevant range of densities. This helps to emphasize the information that would be suppressed in case of resampling of 4096 possible gray levels in the CT data to 256 levels of gray in the PC workstations. Ideal is the combination of color DVR context and grayscale rendering of focus area with the “windowing” possibility.

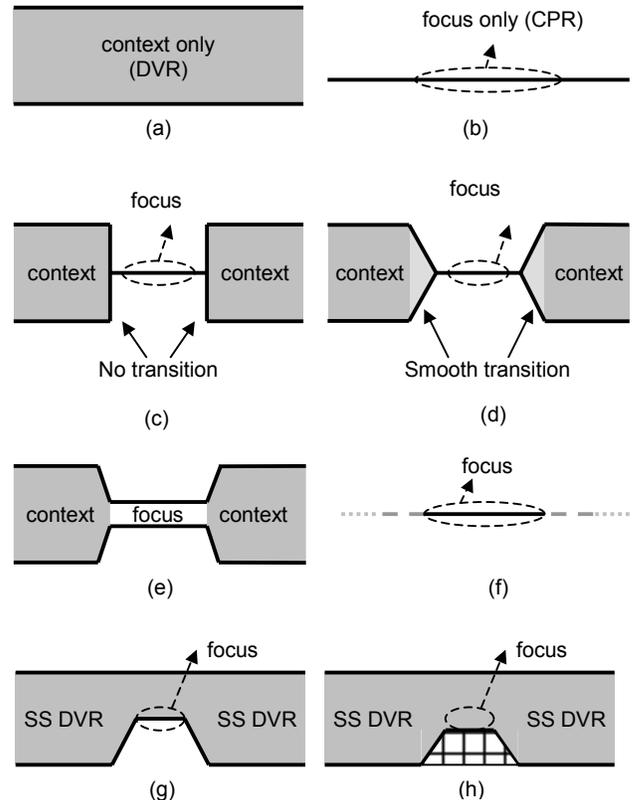


Fig. 3. VesselGlyph examples: (a) DVR of the entire dataset, (b) CPR through vessel axis, (c) combination of CPR and DVR without transition, (d) combination of CPR and DVR with transition region, (e) ThickSlab, (f) CPR with fading in context area, (g) Foreground-Cleft, (h) Foreground-Cleft with occlusion lines



Fig. 5. Application of VesselGlyph with CPR for focus and DVR for context (DVR+CPR VesselGlyph).

3.2 Blended CPR + DVR VesselGlyph

We propose a transition area between the focus and the context area in the VesselGlyph (Fig. 3d). This eliminates the sometimes disturbing abrupt transition in images produced with the CPR+DVR VesselGlyph. The crossover region ensures a smooth transition from the technique used for the focus area to the technique used in the context area. This can be implemented in various ways, e.g. by linear interpolation of both techniques or by

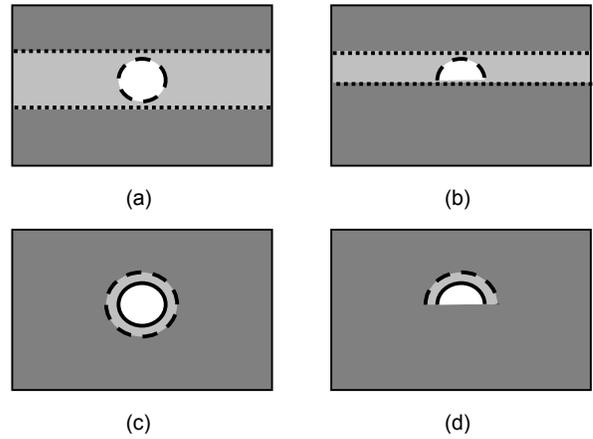


Fig. 6. (a) Full-vessel ThickSlab VesselGlyph, (b) half-vessel ThickSlab VesselGlyph, (c) full-vessel Tubular VesselGlyph, (d) half-vessel Tubular VesselGlyph

changing the width of the context within the transition area. Fig. 7a shows an example of such blended VesselGlyph.

3.3 Foreground-Cleft + DVR VesselGlyph

CPR images provide a good view of the vessel interior; but do not give sufficient information on the overall vessel shape and its front-to-back spatial ordering. Moreover, they fail to visualize thin vessels, when the centerline is incorrectly identified. In these cases, the radiologists prefer DVR images. As mentioned before, object occlusion is a problem then. Therefore, we propose a VesselGlyph, which ensures that the objects in front of the focus region will be suppressed, resembling a cleft in the foreground

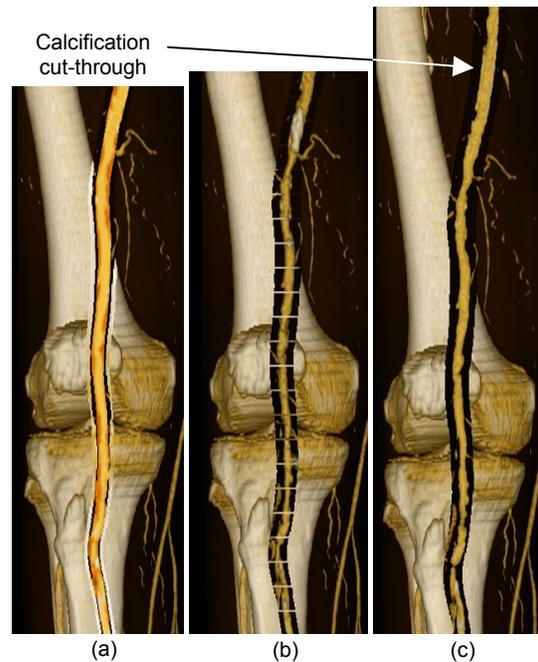


Fig. 7. Close-ups of a knee are for: (a) CPR blended in DVR, (b) Foreground-Cleft with occlusion lines (full-vessel), (c) Foreground-Cleft (half-vessel)

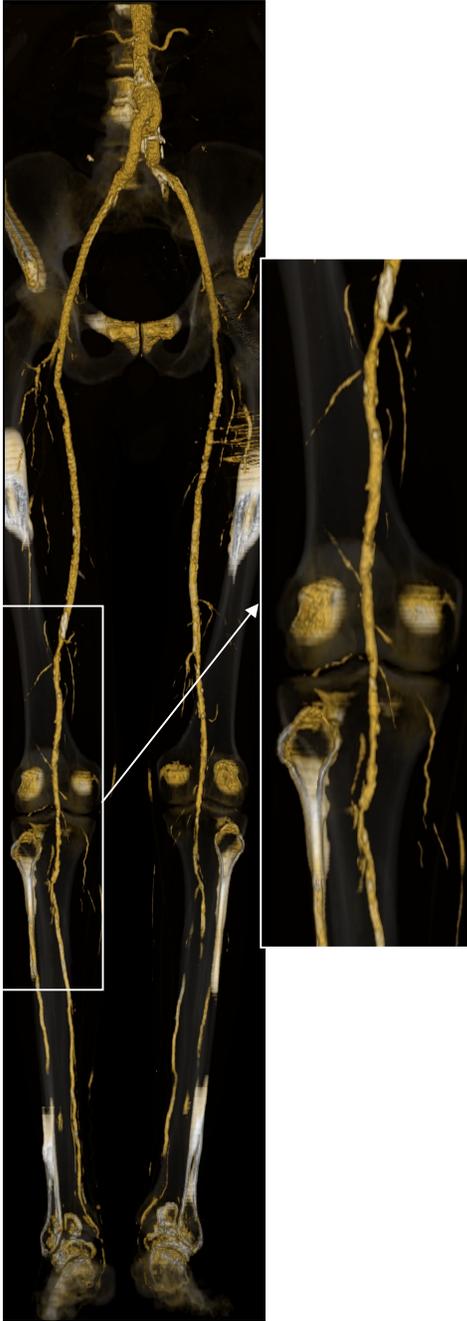


Fig. 8. ThickSlab VesselGlyph rendering. Context is rendered very transparent and slab in focus is rendered normally.

(Fig. 3h).

For DVR visualization of vessels, one has two possibilities – either to show the entire vessel or to show vessel interior. The first case we call “*full-vessel configuration*” (Fig. 6a). The second case we “*half-vessel configuration*” (Fig. 6b).

The Foreground-Cleft VesselGlyph prevents the loss of spatial arrangement information. To emphasize this information in images, we suggest to render occlusion lines (Fig. 7b), which indicate whether the area of focus is behind the actual object. Using this concept, the investigation of visualized objects is possible and still the front-to-back information is available.

3.4 ThickSlab VesselGlyph

The CPR+DVR or the Foreground-Cleft VesselGlyph show clearly and correctly the focus regions in their context. However, the radiologists are also sometimes interested in other structures surrounding the main vessels, e.g. collateral vessels. These structures lie in the context area of the previously mentioned VesselGlyphs where they might be occluded by dense objects, e.g. bones. As mentioned above, this problem is partially solved by CPR or thick-CPR, which are actually special cases of the proposed ThickSlab VesselGlyph.

The ThickSlab VesselGlyph has a focus region in the shape of a slab curved along the vessel path (Fig. 3e). We have several possibilities how to render this VesselGlyph. The context region can be set fully transparent and MIP can be used in the focus area (which is actually the thick-CPR technique of Kanitsar et al, [2002]). As an alternative, DVR can be used within the focus with the context area set either fully or partially transparent (Fig. 8). Fig. 10b shows half-vessel rendering of ThickSlab VesselGlyph.

3.5 Tubular VesselGlyph

In case the radiologists want to see vessels clearly depicted within the suppressed context, we propose the VesselGlyph illustrated in Fig. 6c and Fig. 6d. DVR is used for both focus and context, but the context is rendered very transparent, giving only a supplementary information, whereas objects in focus are rendered photorealistic and in high quality. Fig. 10a depicts half-vessel rendering and Fig. 9 shows a full-vessel application of the tube VesselGlyph. Note, that bony structures in close vicinity to the vessels of interest are also fully rendered. This may cause “pseudo-calcifications” of the vessel wall, or small areas of occlusion. Changing the viewing direction may alleviate this problem. This problem has a reason in the fact that we do not have precise information about vessel diameter in the given point. We only estimate this value, but the estimation is not very precise and therefore the focus region might include also objects that lie in close vicinity of the vessels, e.g. adjacent bones.

4 Implementation Details

The VesselGlyph can be implemented as an extension of a DVR algorithm, where a transparency modifier is assigned to each volume element through the VesselGlyph. The value of the modifier is based on the orientation and the distance of the actual volume element to the vessel centerline. This modifier later

Computation time for:	Software (Athlon 1400 CPU, 2GB RAM)	Hardware (nVidia 5700, 128MB RAM)
VesselGlyph transparency modifiers	1-5 min	-
DVR with transparency modifiers	10 min	1 min

Tab. 1 Computation time for: CT-A dataset preprocessing with the VesselGlyph and image production based on the VesselGlyph (final image: 724x1298 pixels, 2x oversampling in viewing direction, DVR with Phong shading model, gradient/shading evaluated only for voxels where the final transparency is lower than 0.9)

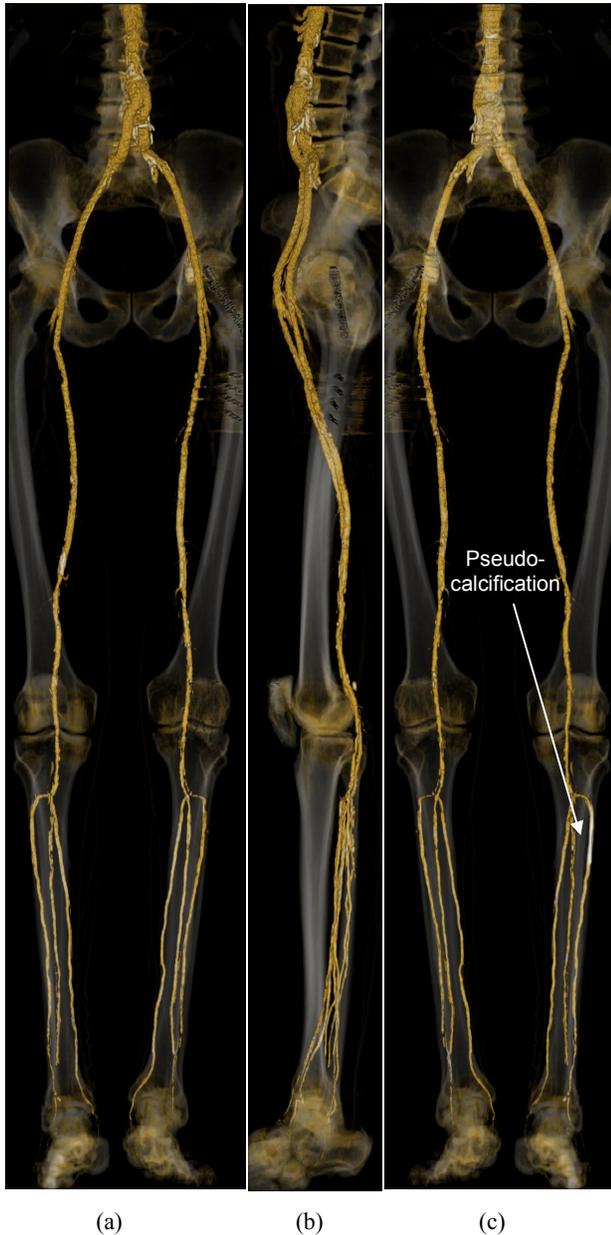


Fig. 9. Full-vessel Tubular VesselGlyph: (a) front view, (b) side view, (c) rear view. These images show multi-path possibility of the VesselGlyph.

controls the transparency function in the DVR stage. The actual software implementation is described in Appendix A.

As described in Tab. 1, image generation using the VesselGlyph is not possible in real-time. The off-line generation of images takes a significant amount of time, which only allows the user to see pre-computed images for a limited number of view angles. Therefore we investigated the possible acceleration by modern graphics hardware in the image production stage.

Modern graphics hardware, driven by the gaming industry, evolves rapidly. Today, many graphic accelerators support 3D textures and are usable for 3D volume visualization [Cabral et al. 1994; Rezk-Salama et. al 2000].

Visualization of 3D textures is done by defining of a set of cut polygons with specified 3D texture coordinates. These 2D

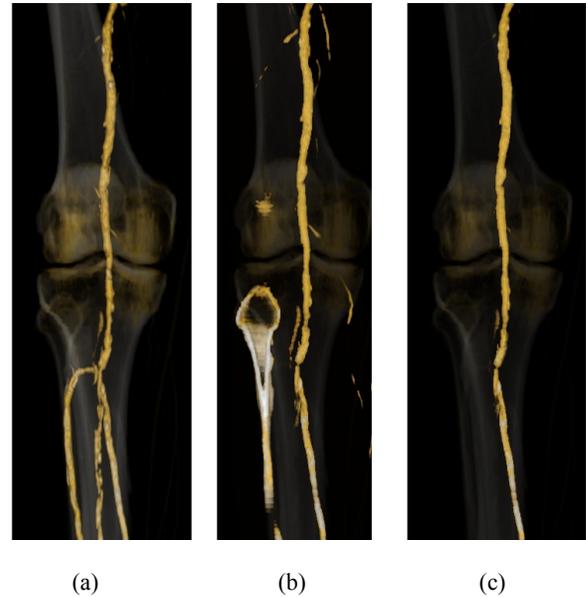


Fig. 10. (a) Full-vessel rendering with Tubular VesselGlyph (multipath version), (b) Half-vessel rendering with Tubular VesselGlyph, (c) Half-vessel rendering with ThickSlab VesselGlyph

polygons are then rendered taking advantage of hardware interpolation of the 3D data to 2D planes. The texture unit of the GPU maps the densities to colors and transparency values which are used for direct volume rendering. This part of the GPU is user-programmable (fragment program) and can merge different 3D textures in the rendering stage. In our case, this feature can be used then for multiplication of the transparency of the original data with the transparency modifier at a given location.

Using this approach we were able to render datasets of 512x512x128 voxels with 256 cutting polygons in 9 fps on an nVidia 5700 GPU with 128 MB RAM. The 128 slices is the maximum size of the dataset due to the RAM size on the graphics card. We expect that in the future hardware-supported rendering of full-sized datasets will be possible. To overcome this problem today, we developed a system, where 128 slices-thick slabs are rendered sequentially and the final image is composed part by part.

5 Conclusion and Future Work

In this paper we have proposed a set of novel focus & context visualization techniques for vessel investigation, based on the generalized notation referred to as the VesselGlyph. The results obtained by various modifications of the basic concept have been presented and confirm the feasibility of the VesselGlyph to comprehensively display vascular pathology within a diagnostically meaningful anatomic context.

Further evaluation and comparison of the presented techniques is necessary to identify the most useful set of parameters for clinical application in the setting of peripheral arterial occlusive disease. In addition, we believe that the concept of the VesselGlyph is favorably suited for the application in other vascular territories, notably the coronary arteries. The general problem of visualizing a vascular tree with morphologic changes due to atherosclerosis is comparable to the peripheral arterial tree.

In the future, the VesselGlyph concept should allow incorporation of more techniques, e.g. average intensity projection or allow rendering of transparent vessels in opaque context using the half-vessel approach. In this way, the VesselGlyph concept should provide more alternatives for the vessel visualizations. For this, further development and implementation of the VesselGlyph in hardware is required.

Acknowledgements

This work has been supported by the Austrian Science Fund (FWF) grant no. P15217.

Appendix A

Mathematically, we can describe the basic algorithm for VesselGlyph techniques as follows: let \bar{x} be the location vector and let $D = D(\bar{x})$ be the density in a given location. Then, for DVR we have mapping functions dependent on densities (f_{red} , f_{green} , f_{blue}) that return coloring information (c_r , c_g , c_b) in the respective location and a mapping function for transparency (f_{trans}) that return transparency t based on the density.

$$\left. \begin{aligned} c_r &= f_{red}(D) \\ c_g &= f_{green}(D) \\ c_b &= f_{blue}(D) \end{aligned} \right\} \text{ for color}$$

$$t = f_{trans}(D) \text{ for transparency.}$$

Let us also we assume have the distance transform function (DXF) on which the VesselGlyph is based:

$$t_m = f_{DXF}(\bar{x}), \text{ which gives us the transparency modifier } t_m.$$

The transparency modifier modifies the original transparency values or can have special values (e.g. one special value results in setting the original transparency to 0.0 – this is a way how CPR can be achieved).

For the iterative DVR algorithm, we get:

```
sum_alpha = 1.0
iterate{
  t_final(x) = f(t(x), t_m(x))
  ilum = f(∇(x))
  sum_red_i(x) = sum_red_{i-1} + sum_alpha_{i-1} * (1 - t_final(x)) * c_r(x) * ilum
  sum_green_i(x) = sum_green_{i-1} + sum_alpha_{i-1} * (1 - t_final(x)) * c_g(x) * ilum
  sum_blue_i(x) = sum_blue_{i-1} + sum_alpha_{i-1} * (1 - t_final(x)) * c_b(x) * ilum
  sum_alpha_i = sum_alpha_{i-1} * t_final(x)
}
```

References

- BAUM, S., AND PENTECOST, M. J. 1997. *Abrams' angiography: vascular and interventional radiology*. Little Brown, Boston.
- DIEPSTARTEN, J., WEISKOPF, D., AND ERTL., T. 2003. Interactive cut-away illustrations. In *Proceedings of EUROGRAPHICS '03*, 523-532.
- GUTWIN, C., DYCK, J. AND FEDAK, C. 2003. The effects of dynamic transparency on targeting performance. In *Proceedings of Graphics Interface '03*, 105-112
- HAHN, H., PREIM, B., SELLE, D., PEITGEN, H.-O. 2001. Visualization and interaction techniques for the exploration of vascular structures. *IEEE Visualization 2001*, 395-402.
- KANITSAR, A., WEGENTKITTL, R., FELKEL, P., FLEISCHMANN, D., SANDNER, D., AND GRÖLLER, E. 2001. Computed Tomography angiography: a case study of peripheral vessel investigation. In *Proceedings of IEEE Visualization 2001*, 477-480
- ZHOU, J., DÖRING, A., AND TÖNNIES, K. D. 2004. Distance Transfer Function Based rendering. Tech. Rep. TR-ISGBV-01-04. Institute for Simulation and Graphics, University of Magdeburg, Germany.
- KANITSAR, A., FLEISCHMANN, D., WEGENTKITTL, R., FELKEL, P., AND GRÖLLER, E. 2002. CPR: Curved planar reformation. In *Proceedings of IEEE Visualization 2002*, 37-44
- ZHOU, J., HINZ, M., AND TÖNNIES, K. D. 2002. Hybrid-Focal Region Based Rendering Of Medical Data. In *Bildverarbeitung für die Medizin 2002*, Leipzig. 113-116
- HAUSER, H., MROZ, L., BISCHI, G.-I., AND GRÖLLER, E. 2001. Two-level rendering. Published in *Journal IEEE Transactions on Computer Graphics and Visualization (IEEE TVCG)*, 242-252
- BÜHLER, K., FELKEL, P., LA CRUZ, A. 2003. *Geometric Methods for Vessel Visualization and Quantification - A Survey*. Geometric Modelling for Scientific Visualization. G. Brunnet, B. Hamann and H. Müller (eds.). Kluwer Academic Publishers
- LEVOY M. 1998. Display of surfaces from volume data. *IEEE Computer Graphics & Applications*, 8(5):29-37
- CARD, S., MACKINLAY, J., AND SHNEIDERMAN, B. 1999. *Readings in Information Visualization: Using Vision to Think*, Morgan Kaufmann Publishers.
- REZK-SALAMA C., ENGEL, K., BAUER, M., GREINER, G., AND ERTL, T. 2000. Interactive Volume Rendering on Standard PC Graphics Hardware Using Multi-Textures and Multi-Stage Rasterization. In *Proc. Eurographics/SIGGRAPH Workshop on Graphics Hardware 2000 (HWWS00)*
- CABRAL, B., CAM, N., AND FORAN, J. 1994 Accelerated volume rendering and tomographic re-construction using texture mapping hardware. In *Proc. of IEEE Symposium on Volume Visualization*, 91-98
- ROSSNICK, S., LAUB, G., BRAECKLE, R. 1986. Three dimensional display of blood vessels in MRI. In *Proceedings of the IEEE Computers in Cardiology Conference*, 193-196.
- DREBIN, R.A., CARPENTER, L. AND HANRAHAN, P. 1988. Volume Rendering. In *Proceedings of SIGGRAPH '88*, Computer Graphics 22, 65-74.