Evaluation of angiogram visualization methods for fast and reliable aneurysm diagnosis

Žiga Lesar, Ciril Bohak, Matija Marolt
Faculty of Computer and Information Science, University of Ljubljana, Slovenia

ABSTRACT

In this paper we present the results of an evaluation of different visualization methods for angiogram volumetric data - ray casting, marching cubes, and multi-level partition of unity implicits. There are several options available with ray-casting: isosurface extraction, maximum intensity projection and alpha compositing, each producing fundamentally different results. Different visualization methods are suitable for different needs, so this choice is crucial in diagnosis and decision making processes. We also evaluate visual effects such as ambient occlusion, screen space ambient occlusion, and depth of field. Some visualization methods include transparency, so we address the question of relevancy of this additional visual information. We employ transfer functions to map data values to color and transparency, allowing us to view or hide particular tissues. All the methods presented in this paper were developed using OpenCL, striving for real-time rendering and quality interaction. An evaluation has been conducted to assess the suitability of the visualization methods. Results show superiority of isosurface extraction with ambient occlusion effects. Visual effects may positively or negatively affect perception of depth, motion, and relative positions in space.

Keywords: volume rendering, angiogram visualization, evaluation, ray casting, image perception

1. INTRODUCTION

Angiography is a medical imaging technique used to visualize the inside of blood vessels. Computed Tomography (CT) and Magnetic resonance imaging (MRI) angiography is used to produce volumetric images, which are in a mathematical sense a sampling of a function $f : \mathbb{R}^3 \to \mathbb{R}$. Equipped with scalar data alone one has countless options for visualizing such images.

One of the oldest and most widely-used algorithms for visualizing volumetric data is ray casting, introduced in 1982 by Scott Roth.[1] The method relies on direct evaluation of the rendering integral. Its flexibility comes from the fact that the user can freely alter the transfer function (classification gradient), which maps volume data values to color and transparency, and effectively change how light reacts to different volume data values. The transfer function can be chosen in a way that it extracts isosurfaces or emphasizes values that correspond to specific tissues.

Another way of visualizing volumetric data is converting a raw volumetric image into a set of polygons - a polygonal mesh - which are more suitable for rendering on modern graphics cards. The method, called marching cubes,[2] is in result comparable to isosurface rendering with ray casting, though it has its pros and cons - most notably rendering speed and model complexity, respectively.

A different approach, known as multi-level partition of unity implicits,[3] is used to reconstruct surfaces from point cloud data. The method is capable of reproducing sharp features as well as smooth surface approximation, which may be desirable in the final rendering of vascular architecture. Its main advantage over marching cubes is more accurate feature reconstruction, though it takes more time to compute the model.

Different visual effects may be added during the rendering process or as a post-processing filter. Screen-space effects are generally faster and thus preferable over object-space effects, which require additional sampling of volumetric data to produce results. Use of effects should though be limited to image and feature enhancement. Although often useful and visually appealing, the magnitude of the effects should be chosen with care as to not make the final rendering worse in terms of harder or unreliable diagnosis.

Further author information:
Žiga Lesar: ziga.lesar@lgm.fri.uni-lj.si, Ciril Bohak: ciril.bohak@fri.uni-lj.si, Matija Marolt: matija.marolt@fri.uni-lj.si

2. DESCRIPTION OF PURPOSE

We are actively working on visualization of 3D angiography data, focusing on facilitating aneurysm detection. As aneurysms are in general quite difficult to observe, a straightforward quality visualization system is needed to speed up and improve the process of detection, identification and diagnosis. On the other hand, better visualization means faster identification of healthy samples. Speed and reliability of the process are of great importance here, so we intend to identify the properties of visualization systems that significantly improve on this process. Hopefully our findings will yield better visualization systems in the future.

3. METHODS

This section describes various visualization methods and effects that we have applied to angiography. Pros and cons are presented for each method.

3.1 Ray casting

Ray casting is an image-based rendering method, which directly evaluates the rendering integral based on how light reacts to the volumetric data at a certain point on a ray. Scalar field reconstruction can be done using an approximation of a 3D analogue of a sinc function. The sinc function is in theory the perfect choice for signal reconstruction, but its extent is infinite and its use is thus intractable in real-time applications. In our implementation we use box and tent filters - nearest neighbor and trilinear interpolation, respectively. For performance reasons the filter is chosen dynamically during rendering by the ray casting algorithm, while mostly preserving the quality of the final rendering.

Image synthesis is done using different compositing schemes. We have implemented first-hit (isosurface extraction), maximum intensity projection,[4] and alpha compositing. Each compositing scheme has unique properties, which can be taken advantage of to speed up the rendering process and to provide the user with specific relevant information. These compositing schemes are presented in the following paragraphs.

Isosurface extraction is suitable for basic analysis of volumetric data. Aiming for speed, we have implemented an octree-based traversal algorithm kd-restart[5] and screen-space sparse sampling.[6] These two methods combined speed up our application by about 20%. Once the isosurface is found we use the Phong shading model[7] to simulate light interaction as it's fast and provides the user with a lot of information about the model's topology and surface orientation. The normal vector used in the algorithm is the gradient of the scalar field, which is estimated with the central difference formula. For better visual quality trilinear interpolation of the normals may be chosen adaptively. There were several other options for the shading model, but none of those were simple and fast enough for a real-time implementation. The application is in turn faster, encouraging the viewer to take advantage of camera animation.

Maximum intensity projection (MIP) is a compositing scheme where the maximum value of the scalar field along a ray is projected on the screen with a color corresponding to the maximum value. Although this may be the perfect method for identifying areas of the volume with large value differences (e.g. vessel walls), it does not provide many clues for better depth perception. In fact, if the projection is orthographic, it is impossible to distinguish between left or right and front or back, but this can be partially overcome by projecting local maxima only. Our implementation makes use of perspective projection along with user-controlled camera animation, which should make a huge difference from the popular static rendering approach. Most volumetric images in the medical field also contain a certain amount of noise that can, using MIP, obstruct important features. Assuming its magnitude is relatively low one can pre-filter the image with a low-pass filter. Furthermore, as MIP is mostly used to extract the parts of the volumetric image with high numerical values, simple thresholding can emphasize this distinction.

Alpha compositing, also known as the *emission-absorption model*, is a way to display different structures contained in the volume on the same image. It achieves that by introducing an optical model that simulates transparency. This is done using a so-called *transfer function*, which maps scalar values to color and opacity. Users may adjust the transfer function interactively and in effect choose the tissues to be displayed. The method is very successful when analyzing still pictures, but is is rather slow for interactive applications, though real-time

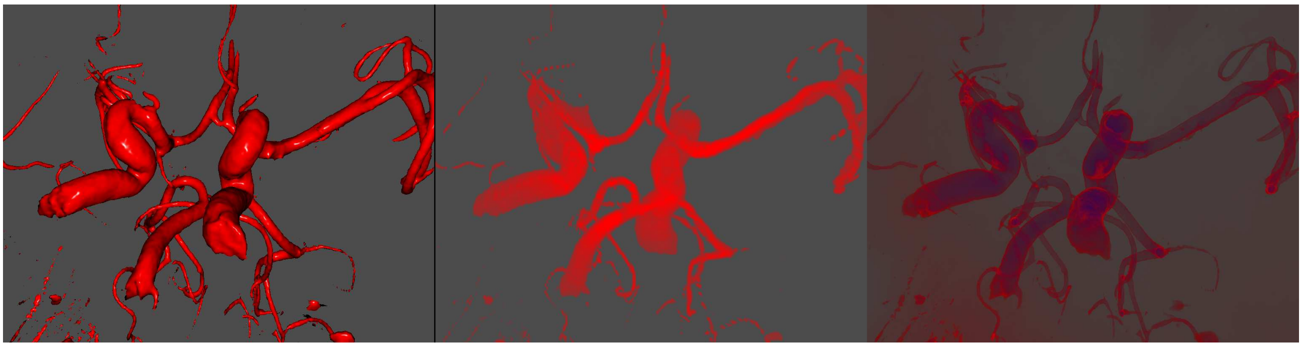

Figure 1: From left to right: isosurface rendering, maximum intensity projection (MIP), and alpha compositing. Phong shading in the leftmost image is a big improvement over flat-colored MIP, though the latter could be additionally equipped with shading, essentially defeating the purpose of the method. Alpha compositing combines both transparency and color mapping, producing quality renders at interactive rates. It could further be extended with shading and shadows, but that may be intractable on commodity hardware.

rendering with this method is already possible. Transparent features may also affect depth perception, but this may be overcome with animation and proper choice of the transfer function.

Comparison of the three methods presented can be seen in Fig. 1.

3.2 Marching cubes

A well-known algorithm for converting volumetric data into triangular meshes is marching cubes.[2] It takes eight neighbor volume values at a time and classifies them as *inside* or *outside* (0 or 1), if their values are higher or lower than the isovalue, respectively. This 8-bit result of classification is then used as an index in a lookup table to construct a few triangles for the output mesh. The downside of this approach is obvious: to view a different isosurface of the same volumetric data set one must construct an entirely new mesh, which is of course time-consuming. Implementing the algorithm in OpenCL is a lot faster, but still slower than ray casting. For a large volume this method also potentially yields lots of triangles.

The algorithm has some advantages over ray casting though. Well-established algorithms for triangle mesh manipulation may be used to enhance the result and make it more visually appealing, possibly facilitating identification of important volumetric features. The mesh can also be easily textured, colored, and/or shaded, a common practice in computer graphics to enhance certain features and improve depth perception. As modern graphics hardware is built specifically for rendering triangular meshes, this approach is also a lot faster than ray casting - after the mesh is already built. However, the resulting model often lacks fine features and thus is not always preferable in medical applications. A number of improvements have been developed over the years, one of them being multi-level partition of unity implicits, which can preserve sharp features and details contained in the original volumetric data and in the same time take advantage of modern graphics hardware.

3.3 Multi-level partition of unity implicits

This is an algorithm for reconstructing surfaces from point cloud data. Piecewise quadratic functions, which nicely approximate local surface curvature, are blended together with weighting functions (the partitions of unity). Having a 3D volume as an input, we must first convert that data into a point cloud. Although this gives us much control over what is passed forward to the MPU algorithm, it is a relatively slow process. A MPU implicit is in mathematical sense a function $f : \mathbb{R}^3 \to \mathbb{R}$, which maps points in space to distances from the surface - a variant of a distance transform. This simplifies the process of visualization. The authors demonstrate this using Bloomenthal's polygonizer[8] and Hart's sphere tracing.[9] Also, the user can enhance the output with offsetting, and may even blend two MPU implicits together, as described in the original paper.

In our case the algorithm produces results comparable to those obtained with marching cubes. The main difference is MPU implicits' ability to accurately represent sharp features of the volume, such as blood vessel

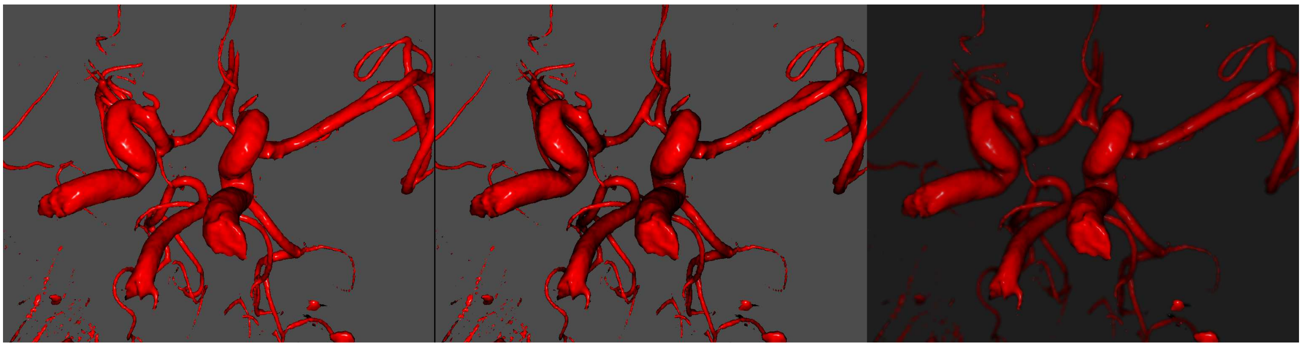

Figure 2: From left to right: no effects, screen-space ambient occlusion, depth of field. Identification of overlapping blood vessel segments should be easier with occlusion effects and depth illusion hints as shown in the central and rightmost rendering. Focused segments in the rightmost image seem to stand out and should attract the user's attention by being rendered sharp and in a brighter tone. However, parts of the image may become too dark, so it is possible for important features to go unnoticed.

walls at vessel branchings. When the visualization method of choice is polygonization, the benefits described in the marching cubes sections also apply here.

The main downside of this approach is computational time of the MPU algorithm. For large data sets it performs rather poorly, though hardware acceleration is possible. Once the transform is computed, the benefits described above can be taken advantage of.

3.4 Ambient occlusion

Ambient occlusion is an efficient method for approximating effects of global illumination. This is achieved by detecting proximity of surfaces. If a point on the surface is surrounded with other solid structures, those structures obstruct light that would otherwise hit the point on the surface. If the light is coming from a light source this creates a shadow, and if the light is indirect (e.g. reflected from a nearby object) the effect is called ambient occlusion. Computing realistic occlusion includes exhaustive sampling of the volumetric data, thus being computationally expensive. It may possibly prolong the rendering time beyond the scope of real-time when using transparency. Only rough approximations are possible in real-time, but fortunately this still represents a major improvement for the user's perception of proximity.

Since the user is mainly interested in the structures that are actually visible in the final rendering, the effect can be simplified enough to give satisfactory results with minimal computational time. Whereas the approximation is carried out in the final rendering, the effect is called screen-space ambient occlusion. We have implemented both an object-space version and a screen-space version. The screen-space ambient occlusion algorithm has been developed with speed in mind. Its main advantage is independence of the complexity of the scene. This post-processing algorithm provides a lot of clues for better depth perception, in exchange for dark patches, which may hide relevant information from the user.

Both algorithms were implemented as an estimation of the ambient occlusion integral.[10] The computational time of the screen-space variant is negligible compared to the benefits of the effect, most notably better and faster perception of how volume segments overlap in the final rendering.

Comparison of ambient occlusion with plain shading can be seen in Fig. 2.

3.5 Depth of field

Depth perception can be further enhanced by simulating the lens of a camera. A lens can precisely focus only objects at focal distance from the camera, so other parts of the scene appear blurred in the final rendering. This effect is called depth of field. Out-of-focus scenery should redirect the user's attention to the focused parts of the image. Focusing may be additionally emphasized by darkening defocused objects.

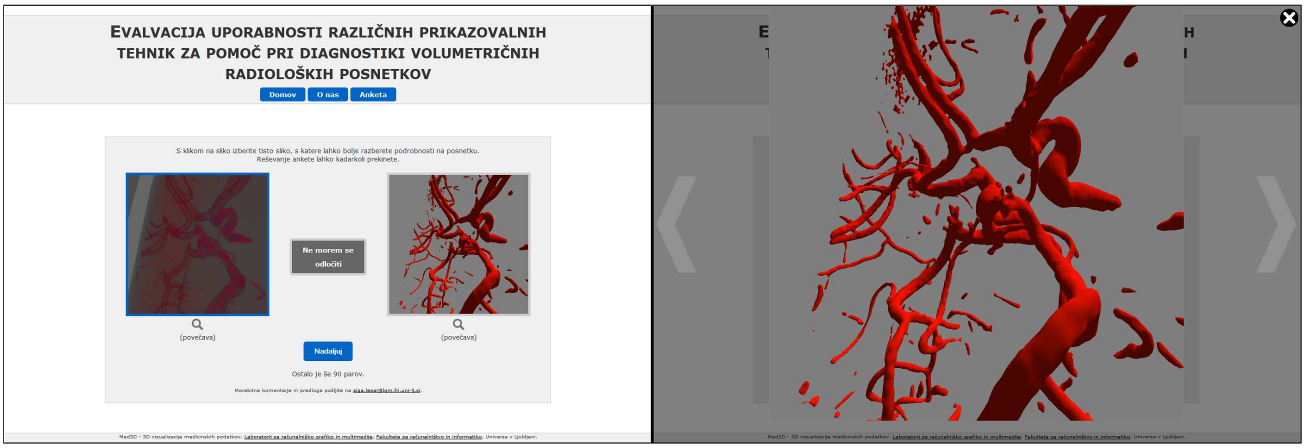

EVALVACIJA UPORABNOSTI RAZLIČNIH PRIKAZOVALNIH TEHNIK ZA POMOČ PRI DIAGNOSTIKI VOLUMETRIČNIH RADIOLOŠKIH POSNETKOV

S klikom na sliko izberite tisto sliko, s katere lahko bolje razberete podrobnosti na posnetku.
Reševanje ankete lahko kadarkoli prekinete.

Ostalo je še 90 parov.

Morebitne komentarje in predloge pošljite na ziga.lesar@lgm.fri.uni-lj.si.

Med3D - 3D vizualizacija medicinskih podatkov. Laboratorij za računalniško grafiko in multimedije, Fakulteta za računalništvo in informatiko, Univerza v Ljubljani.

Figure 3: Two screenshots of the survey website. The users were asked to choose the rendering in which important structures, shapes and surfaces could be recognised faster and easier (left). One could enlarge a rendering at will to take a closer look (right).

Depth of field can be approximated in screen space. Our implementation uses linear interpolation of sharp and blurred versions of a rendering, with the interpolation factor being the distance from the focal point, which we obtain from the depth image. The computational time for this effect is also relatively small compared to the rendering method itself, which is also a welcomed feature.

Since a human eye is in a way similar to a camera lens, depth of field simulation seems an obvious way to enhance a rendering by giving it a more natural look. However, the blurring removes a fair amount of details and fine structures from the image for which we have strived for by using appropriate visualization methods.

An example rendering with depth of field is shown in Fig. 2.

4. EVALUATION

To determine the desirable qualities of a visualization system we have conducted an evaluation consisting of an on-line survey and a hands-on test. The participants were mostly radiology specialists - the potential users of advanced medical imaging software - from whom we obtained opinions on different visualization methods and visual effects. The main purpose of the evaluation was to give an in-depth comparison of those methods and assess their suitability for the purpose of aneurysm diagnosis. The results of the evaluation will in effect help to make diagnoses easier, faster, and more reliable.

From each participant we collected some basic demographic information, such as age, occupation and specialization. Then the participants were confronted with a series of pairwise comparison tests between several renderings of angiograms. The users were prompted (but not compelled) to choose the rendering in which important structures, shapes and surfaces could be recognised faster and easier. They had the option to enlarge a rendering to take a closer look. The survey is shown in Fig. 3. The set of renderings was composed of 3 different volumes with 2 viewing positions each. Comparison of renderings of different volumes or the same volume from different viewing positions was not included. 6 different visualization methods, which are in widespread use in practice, were compared:

- marching cubes (MC),
- multi-level partition of unity implicits with polygonization of cubes (CUBE),
- multi-level partition of unity implicits with polygonization of tetrahedra (TET),
- isosurface extraction with object-space and screen-space ambient occlusion (ISO),
- isosurface extraction with both variants of ambient occlusion and depth of field (DOF), and

- alpha compositing (ALPHA)

This sums up to $\binom{6}{2} \cdot 3 \cdot 2 = 90$ pairs of renderings.

5. RESULTS

We have received opinions of 24 participants of which 6 have evaluated all 90 comparison tests. A total of 659 pairs of renderings have been compared. The participants were mostly radiology specialists (including all 6 of those whose responses included all 90 comparison test) with an average age of 38.67.

For the analysis we have chosen balanced rank estimation[11] (BRE) and voting (VOTE). In the former case (BRE), when a user chose one method over another, that method received a point and the other one lost a point. In the latter case (VOTE), only the chosen method received a point. A higher sum of points means a better method. The results of BRE and VOTE ranking is shown in table 1.

	BRE	VOTE	BRE*	VOTE*
ISO	67	103	86	128
MC	11	71	15	87
CUBE	10	66	7	79
TET	-7	57	-12	66
DOF	-22	61	-15	80
ALPHA	-59	41	-81	65

Table 1: Results of BRE and VOTE analysis of the evaluation. BRE* and VOTE* denote the analysis with all responses taken into consideration, while BRE and VOTE denote just the responses with a complete set of answers.

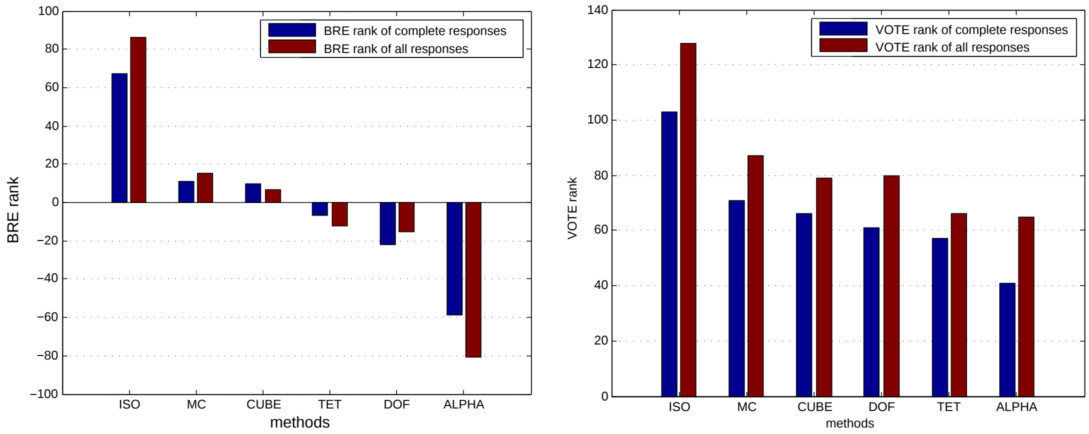

Figure 4: BRE (left) and VOTE (right) analysis of evaluation responses.

We can from Fig. 4 that isosurface extraction with object-space and screen-space ambient occlusion (ISO) has been considered most suitable for the task. Although other visualization methods received comparable scores, alpha compositing (ALPHA) stands out in the final ranking, most probably because of a poor choice of the classification gradient, which made for a rather low-contrast appearance. Low contrast and brightness seemed to also negatively affect the ranking of the depth-of-field effect (DOF).

Pairwise comparison of the methods is shown in Fig. 5, where we show how one method compares to any other in terms of BRE. Both matrices are antisymmetric. The figures undoubtedly support the outcomes seen in Fig. 4 and give more insight into the outcome of MC, CUBE and TET comparison. CUBE seems to have scored better results than MC and TET, whereas these two methods are completely comparable. Taking all survey responses into consideration gives us similar results with only minor discrepancies.

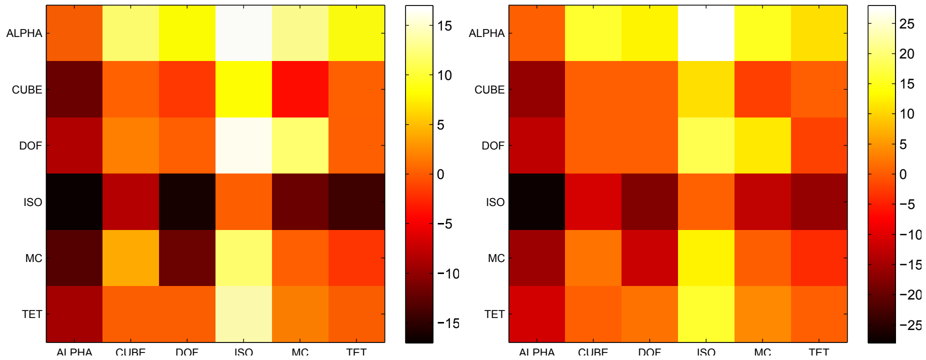

Figure 5: Pairwise analysis of complete responses (left) and all responses (right).

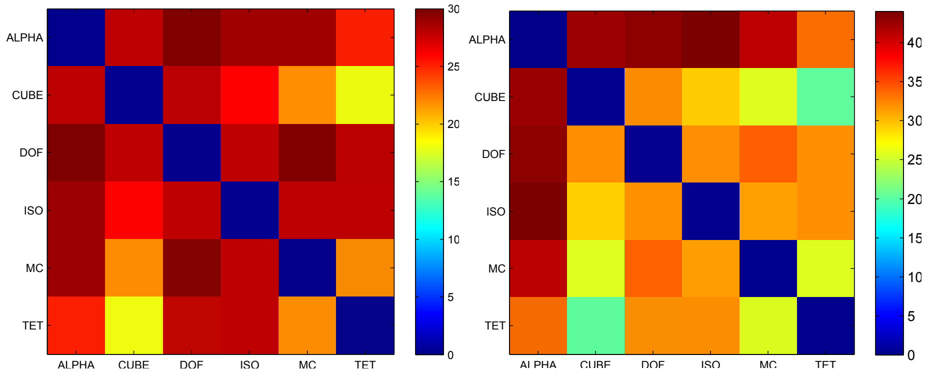

Figure 6: Pairwise analysis of undecidability of complete responses (left) and all responses (right).

We have also analysed undecidability of survey results. Fig. 6 shows the number of comparisons where a rendering method was chosen over another. A lower number thus indicates higher undecidability, which is clearly visible on the diagonals, since we did not allow comparison of a method with itself. The two matrices are symmetrical. As expected, the figures show higher undecidability of TET and CUBE, as these methods share the same algorithmic basis but employ different polygonization techniques. The same goes for MC in comparison with TET and CUBE, as these three methods share the same shader. We can see in the figures that the users had little trouble deciding between DOF and ALPHA or DOF and MC. We find the preference of DOF over ALPHA quite surprising, but that may be a consequence of the fact that the classification gradient used with ALPHA was somewhat unsuitable. Other pairs' scores are fairly straightforward, with low undecidability of the pairs including ALPHA or ISO.

In general, the results of the evaluation have shown better and faster identification of model features (shapes, details, size), as well as better depth perception when using additional effects. Use of visual effects that imitate natural lighting or effects in photography shows to be beneficial for the illusion of depth. It seems that transparency and coloring of different tissues conveys a lot of additional information that may be hidden from the user when viewing isosurfaces only, although a careful choice of the classification gradient is a must. Being able to control the strength of visual effects we conjecture that neither extreme is beneficial to the viewer in the process of feature identification.

We have also performed a live hands-on test with a radiology specialist, confirming our conjectures. While animation seems to enhance depth perception, we also estimate that the user is less focused on details when the features are not stationary in screen-space. Lowering the rendering quality while the parameters of the camera and/or the scene are changing we can effectively raise the rendering speed, adding to the perception of motion and facilitating tracking of specific features of the model. Also, details seem less important in out-of-focus or out-of-scope features. Careful thought should though be aimed at lighting and ambient occlusion effects, which seem to affect general volume perception the most.

6. CONCLUSIONS

To our knowledge, no evaluation of visualization systems for faster aneurysm diagnosis has yet been conducted. We believe that our research is fundamentally important for the quality of decision support. Although most rendering methods rely on reproducing photorealistic effects, photorealism may not be the best solution in this field. Since we have implemented a variety of visualization methods, we have a wide palette of options for comparison. So far it seems that isosurface rendering with advanced lighting effects such as ambient occlusion is the most suitable for aneurysm diagnosis. Although most modern volume visualization methods employ different physically inspired methods for simulating light interaction, our research shows that simple isosurface extraction may be more suitable in this case.

In the other hand, visual effects may positively affect perception of depth, motion and relative positions in space. Since this is of fundamental significance for having a good mental image of the situation, we believe that it may help the users in diagnostic procedures. Moderate use of visual effect should be used only to enhance certain features of the model and to give the users a more explicit understanding of the rendering.

This evaluation is a move in the right direction, but with a large enough set of evaluation samples we can make stronger conclusions. We plan on extending the evaluation further with more experts in this particular field. Hopefully this will improve on the decision making process by making it faster and more reliable, as well as easier and straightforward.

ACKNOWLEDGMENTS

The angiography data sets are courtesy of Özlem Gürvit, Institute for Neuroradiology, Frankfurt, Germany, and are available for download at `http://volvis.org/`. We also gratefully acknowledge the support of NVIDIA Corporation with the donation of the Quadro K5000 graphics card for this research.

REFERENCES

[1] Roth, S. D., "Ray casting for modeling solids," *Computer Graphics and Image Processing* **18**, 109–144 (Feb. 1982).

[2] Lorensen, W. E. and Cline, H. E., "Marching cubes: A high resolution 3D surface construction algorithm," *ACM SIGGRAPH Computer Graphics* **21**(4), 163–169 (1987).

[3] Ohtake, Y., Belyaev, A., Alexa, M., Turk, G., and Seidel, H.-P., "Multi-level partition of unity implicits," *ACM Transactions on Graphics* **22**, 463–470 (July 2003).

[4] Mroz, L., König, A., and Gröller, E., "Maximum intensity projection at warp speed," *Computers & Graphics* **24**, 343–352 (June 2000).

[5] Foley, T. and Sugerman, J., "KD-tree acceleration structures for a GPU raytracer," in [*Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware*], **45**, 15–22 (2005).

[6] Kratz, A., Reininghaus, J., Hadwiger, M., and Hotz, I., "Adaptive screen-space sampling for volume raycasting," Tech. Rep. February, Konrad-Zuse-Zentrum für Informationstechnik Berlin (2011).

[7] Phong, B. T., "Illumination for computer generated pictures," *Communications of the ACM* **18**, 311–317 (June 1975).

[8] Bloomenthal, J., [*An implicit surface polygonizer*], vol. 1, Academic Press Professional, Inc. (1994).

[9] Hart, J. C., "Sphere tracing: a geometric method for the antialiased ray tracing of implicit surfaces," *The Visual Computer* **12**(10), 527–545 (1996).

[10] Miller, G., "Efficient algorithms for local and global accessibility shading," in [*Proceedings of the 21st annual conference on Computer graphics and interactive techniques*], 319–326 (1994).

[11] Wauthier, F. L., Jordan, M. I., and Jojic, N., "Efficient ranking from pairwise comparisons," *Journal of Machine Learning Research* **28**(3), 109–117 (2013).