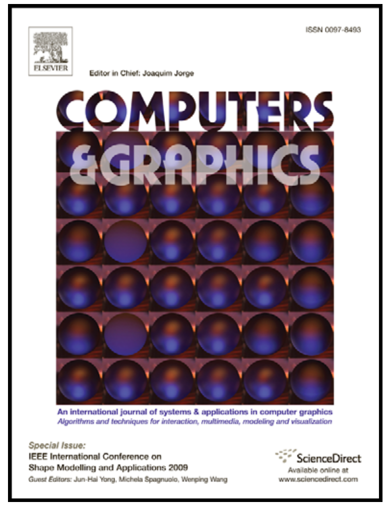


www.elsevier.com/locate/cag

# Perceptual enhancement of two-level volume rendering

Andrew Corcoran, Niall Redmond, John Dingliana

PII: S0097-8493(10)00053-1
DOI: doi:10.1016/j.cag.2010.03.014
Reference: CAG 1999

To appear in: *Computers & Graphics*

Received date: 15 October 2009
Revised date: 22 March 2010
Accepted date: 24 March 2010

Cite this article as: Andrew Corcoran, Niall Redmond and John Dingliana, Perceptual enhancement of two-level volume rendering, *Computers & Graphics*, doi:10.1016/j.cag.2010.03.014

Elsevier Editorial System(tm) for Computers & Graphics
Manuscript Draft

Manuscript Number: CAG-D-09-00136R2

Title: Perceptual Enhancement of Two-level Volume Rendering

Article Type: Special Section: IllustraVis

Section/Category: Technical Section

Keywords: Volume rendering
Non-photorealistic rendering
Visualisation

Corresponding Author: Mr. Andrew Corcoran,

Corresponding Author's Institution: Trinity College Dublin

First Author: Andrew Corcoran

Order of Authors: Andrew Corcoran; Niall Redmond; John Dingliana

Manuscript Region of Origin: IRELAND

Abstract: We present a system for interactive visualisation of 3D volumetric medical datasets combined with perceptual evaluation of how such visualizations can affect a user's interpretation of scenes and attention. Enhancements to traditional volume renderings are provided through a two-level volume rendering strategy which employs fast GPU-based Direct Volume Rendering (DVR) coupled with an additional layer of perceptual cues derived from various techniques from the non-photorealistic rendering (NPR) literature. The two-level approach allows us to successfully separate the most relevant data from peripheral extraneous detail enabling the user to more effectively understand the visual information. Peripheral details are abstracted but sufficiently retained in order to provide spatial reference. We perform a number of perceptual user experiments which test how this approach affects a user's attention and ability to determine the shape of an object. Results indicate that our approach can provide a significant improvement in user perception of shape in complex visualizations, especially when a user has little or no prior knowledge of the data. Our approach would prove extremely useful in technical, medical or scientific visualizations to improve understanding of detailed volumetric datasets.

# Perceptual Enhancement of Two-level Volume Rendering


**Abstract**

We present a system for interactive visualisation of 3D volumetric medical datasets combined with perceptual evaluation of how such visualizations can affect a user's interpretation of scenes and attention. Enhancements to traditional volume renderings are provided through a two-level volume rendering strategy which employs fast GPU-based Direct Volume Rendering (DVR) coupled with an additional layer of perceptual cues derived from various techniques from the non-photorealistic rendering (NPR) literature. The two-level approach allows us to successfully separate the most relevant data from peripheral extraneous detail enabling the user to more effectively understand the visual information. Peripheral details are abstracted but sufficiently retained in order to provide spatial reference. We perform a number of perceptual user experiments which test how this approach affects a user's attention and ability to determine the shape of an object. Results indicate that our approach can provide a significant improvement in user perception of shape in complex visualizations, especially when a user has little or no prior knowledge of the data. Our approach would prove extremely useful in technical, medical or scientific visualizations to improve understanding of detailed volumetric datasets.



*Keywords:*
volume rendering, non-photorealistic rendering, visualisation




## 1. Introduction

3D volume visualisation is an important technique with widespread use in a number of scientific and engineering fields. Popular application areas include medical imaging, scientific visualisation and flow visualisation. In recent years, more advanced graphics hardware and new algorithms have

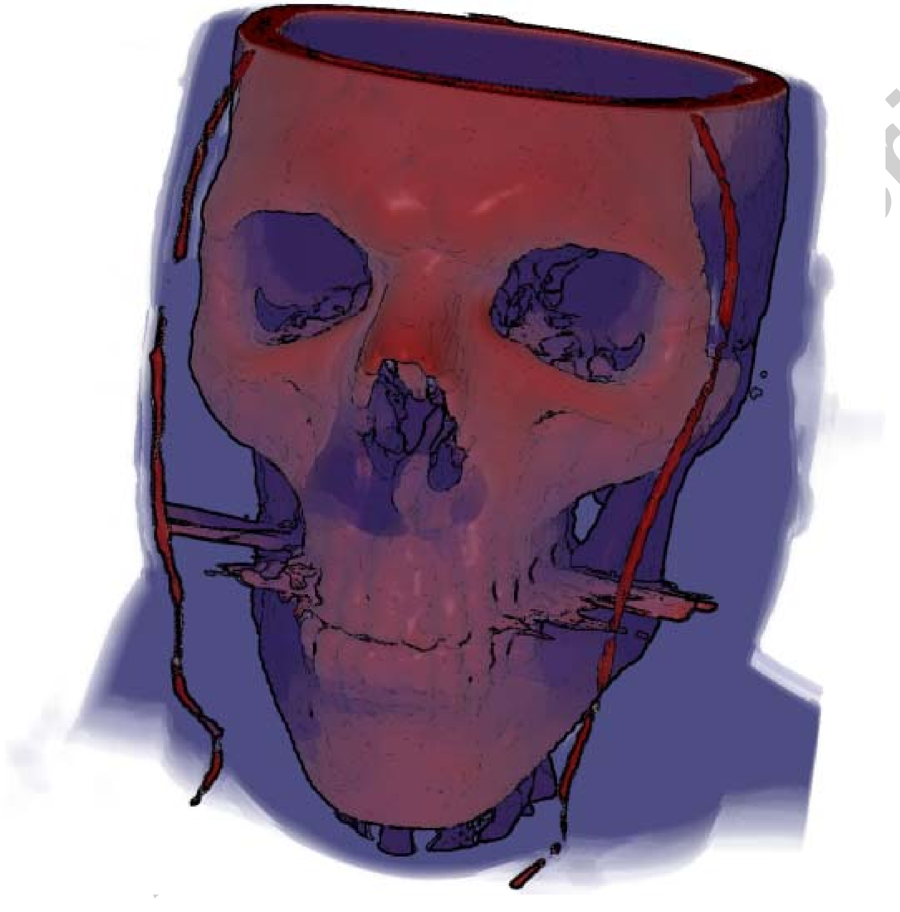

Figure 1: Example from our system; Sobel edge detection combined with suggestive contours and the Kuwahara filter applied to a volumetric model of a skull.

facilitated highly optimised pipelines for fast volume rendering. However, congruent advances in the capabilities of acquisition, storage and simulation technologies have also led to a demand for visualising increasingly complex data.

In addition to the increased performance demands, successful interactive rendering of complex datasets is problematic due to the visual complexity of the data. Rendered images of complex volume data tend to have excessive detail, which is often multi-layered and overlapping. This causes significant demands on the human vision system to process what is often a perceptually challenging type of visual output. In interactive visualisation this problem is further exacerbated as the view of the data is continuously changing and there is even less time per frame for a user to assimilate the information. Thus, there is a clear need for strategies to address perceptual issues in interactive visualisations.

Whilst volume visualisations have been widely used for analysis of complex data, another long-standing approach taken in texts and static imagery is the use of hand-drawn illustrations. These approaches have the benefit of human intuitive input in deciding what are the most significant visual elements of a dataset and how best to communicate this in a visual manner. Interestingly there has been significant work in the past decade in the field of Non-Photorealistic Rendering, which deals with simulating various intuitive mechanisms that human illustrators employ in creating perceptually effective renderings. Such techniques could be exploited in interactive visualisations to alleviate the demands of assimilating information and improving user perception and understanding of 3D volumetric data.

In this paper we present an interactive system that successfully combines a new technique for two-level direct volume rendering [1] with a stylistic NPR overlay that enhances the saliency and fidelity of the most pertinent data. The system simultaneously exploits enhanced DVR and real-time NPR techniques to provide enhanced detail without complete loss of peripheral information, which is retained for spatial context. The paradigm of presenting both a detailed region and a surrounding context in this manner for intuitive imagery has been widely used by artists for hundreds of years and previous research [2] has demonstrated that there are strong potential benefits to adopting such a paradigm in contemporary applications of volume visualisation. In our system, we employ a combination of image-space and object-space line drawing as well as textural abstraction of peripheral data using image-space techniques. The 3D data is first segmented and rendered into two layers of high and low importance and the layers are then post-processed with stylised enhancements and then merged. We provide a mechanism by which the salience of the two layers can be modulated and integrated to provide optimal balance of focus and context.

As the success of the proposed two-level volume rendering approach aims to affect user perception of the scene, a number of perceptual experiments were conducted to evaluate how the abstraction techniques implemented can influence both a user's attention within a scene, and the user's understanding of object shape. The first evaluation looked at how the approach can influence saliency within the scene and how user attention can be affected. A second evaluation was run which explored how each abstraction style affected user shape perception. The experiments show that our system has proven advantages for the exploration of volume data and that saliency of relevant parts of the image can be affected by the techniques used. The results also provide an indication of an optimal combination of stylisations to create the most effective two-level renderings.

## 2. Related Work

Numerous approaches exist for volume visualisation which range from the very accurate to the very interactive [3, 4]. In recent times, the gap between these two extents has narrowed with technologies now enabling the interactive visualisation of very complex volume data. While some approaches first obtain a surface representation from the volume before applying standard 3D renderings, DVR techniques operate directly on the volume data. This can be achieved through the use of 3D textures, generated from view aligned slices of volume data and optimised using GPU hardware for real-time performance [5, 6]. Other approaches employ techniques such as volume ray-casting [7], facilitated through adaptive optimisations for real-time visualisation of implicit surfaces defined by a volumetric grid of sample data.

Apart from the speed-accuracy tradeoffs, another prevalent concern in the literature is the perceptual optimisation of the rendered output. For this, a range of approaches are employed from post-process image operations to direct processing of the 3D data. Engel et al. describe a process for pre-integrating volume data, which eliminates some artifacts related to low volume sampling resolution [8]. Viola et al. discuss methods for highlighting and enhancing pre-segmented volume data through the use of cut-away and ghosted views of volumetric data [9]. Hauser et al. propose the use of two-level volume rendering, which merges DVR and MIP techniques in an interactive tool [1]. They argue that such an approach, based on a *focus-and-context* strategy, provides intuitive benefits to users, allowing them to peer inside inner structures, while keeping surrounding objects integrated for spatial reference. Focus and emphasis is also discussed by Preim et al [10], who survey a number of existing areas where various enhancement and NPR techniques can be used to improve tasks in medical visualisation. Kruger et al. present a focus and context visualisation framework called Clearview, assigning varying levels of importance across an object and rendering the object adaptively using texture-based volume raycasting [2]. Kim and Varshney use a visual saliency operator to compute an emphasis field which is used to elicit viewer attention to relevant parts of a volume rendering [11].

Although some existing visualisation approaches already exploit a degree of multi-level rendering using NPR or other visual enhancements, there are numerous more compelling perceptual techniques to be exploited in the field of NPR, which is closely related to perception and intuitive human techniques for creating imagery. Halper et al. discuss the potential of NPR to affect a user's response to a rendered scene, and show that a user's judgement or choices can be influenced by certain rendering and modelling choices [12]. Related works to this in the NPR literature include the work of Santella and DeCarlo, who describe how variable levels of abstraction can be generated in NPR to achieve modulated saliency and to visually prioritise certain regions of an image over others [13, 14]. Cole et al. discuss specifically directing user gaze in 3D with stylised focus [15]. This is achieved by variable levels of NPR abstraction and stylisations, creating higher saliency in the most relevant parts of a scene. In similar work, Redmond et al. utilise a combination of object and image-based strategies accelerated on the GPU to provide adaptive abstractions, resulting in user attention being drawn to selected parts of a rendered 3D scene [16, 17, 18]. Although these approaches deal with 2D images and/or 3D polygonal data, they serve as evidence of a strong potential for exploiting similar strategies to aid or influence a user's interaction and understanding of volume visualisation.

Of particular relevance, there is a large body of work in NPR that has explored the perceptual and mathematical basis of line drawings, which are particularly useful in a multi-level approach as they provide a non-occluding self-contained dataset that can be integrated relatively easily to augment other rendering styles. Well known classes of line-drawing which can be automatically extracted include: silhouettes, occluding contours, geometric ridges and valleys, suggestive contours [19], apparent ridges [20], and suggestive and principal highlights [21]. Cole et al. perform experiments to determine correspondences between automated line drawing techniques and hand-drawn line illustrations [22]. The results provide an indication of lines that artists believe would most intuitively indicate relevant shape and structure and it is observed that these tend to largely correspond to image-space lines and occluding contours. A more recent study by Cole et al. suggests that users are able to interpret certain shapes almost as well from line drawings as from shaded images and that current computer generated line drawing techniques can match the effectiveness of artists' drawings in depicting shape [23].

Existing works that have already explored NPR techniques in volume rendering include that of Burns et al. [24], who extract a number of classes of lines directly from volume data, although they do not address how this could be combined with other rendering modes for a more complete volume visualisation. Treavett and Chen employ a stylistic rendering approach resembling pen-and-ink illustrations for volume visualisation [25]. They note that such renderings can be overlaid with photo-realistic rendering to create selective abstraction of detail and to improve, for instance, user understanding of translucent parts of the volume. Lum and Ma employ hardware-accelerated parallel non-photorealistic styles in their volume renderings [26], utilising object-space GPU techniques for enhancement of the image. They use tone shading, silhouettes and depth cues to provide perceptual improvements to the rendered volume data. Contour enhancement on the fragment shader and tone-shading are employed in combination with MIP compositing and DVR approaches for focus and context two-level volume rendering by Hadwiger et al. [27]. In separate work, Hadwiger et al. also integrate ridges and valleys and tone shading into their system for real-time volume ray casting [7]. Svakhine et al. [28] present image-based outlining techniques that can generate pure line drawings or helpful feature enhancements for illustrative volume visualisation.



## 3. Pipeline Overview

Our current implementation combines object-space line extraction, two pass volume rendering, image-space edge extraction and texture abstraction, as shown in Figure 2, in order to generate a final perceptually enhanced image. In this section we provide a general overview of our implementation before explaining each technique in detail.

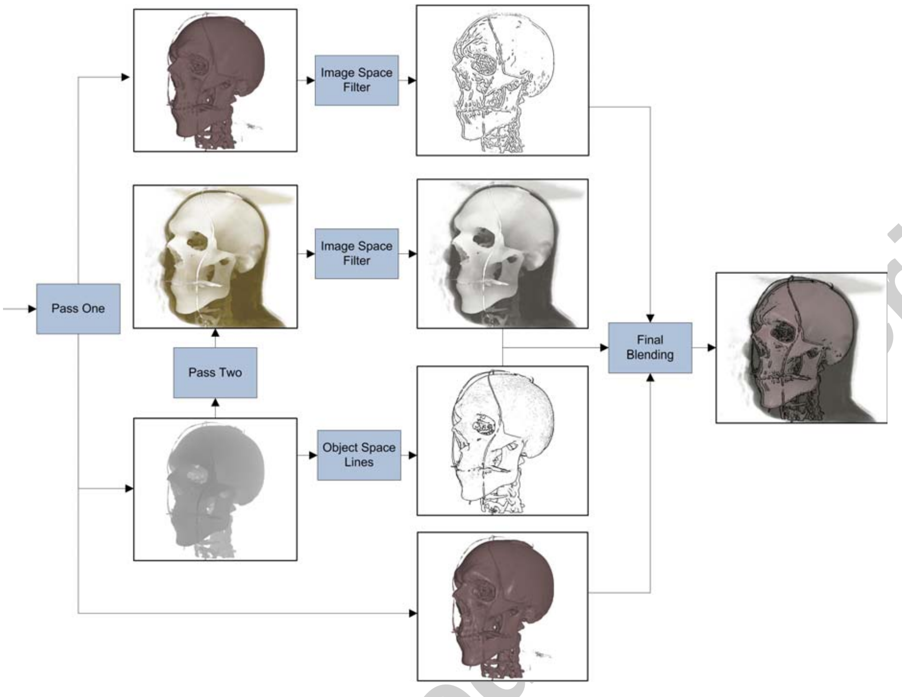


Figure 2: Pipeline Overview

The first step of our implementation is to perform a pre-processing step on the volume data using a 3D Sobel filter in order to calculate the normal vector for each voxel in the dataset. Since we are currently focusing our efforts on static volume data this normal information does not change from frame to frame and we are able to dramatically reduce our system's overall rendering time by performing this calculation once and saving it for future use in subsequent frames.

The next step of our implementation performs two pass volume rendering. Using the normal information calculated in the previous step and by generating a custom transfer function based on the user's input, we are able to generate two separate images from this process: one of the user selected isosurface and one of the remaining volume. These two images are passed on to our GPU image shaders for further post-processing.

In the third stage of our pipeline, we extract object-space lines directly from the volume data using an optimised and parallelised version of an algorithm proposed by Burns et al [24]. Our optimisations result in approximately a 50% increase in performance in comparison to the original implementation by Burns et al. As this stage is performed primarily on the CPU, it is a major bottleneck for performance and unfortunately even our optimised algorithm is unable to render both silhouettes and suggestive contours in real-time. One of the largest computational costs for the object-space line drawing is in the line visibility calculations. We are able to completely eliminate this step by utilising the depth image generated by our volume rendering stage in order to depth cull the generated lines.

We dynamically change the complexity of the lines based on user interaction with the program. This is done by utilising an algorithm by Burns et al. which exchanges line drawing reliability for reduced computational costs. This algorithm is used by our system when the user is interacting with the volume dataset. When the scene is static our optimised line-drawing technique, as described in Section 5.2, is enabled which displays all possible lines, in contrast, the Burns et al. algorithm does not guarantee that all lines will be drawn.

The penultimate step of our implementation is the application of post processing filters to the volume images, implemented through the use of GPU shaders. The types of filters we use fall into two general categories: edge extraction calculations applied to add detail to the target isosurface; and texture smoothing filters applied to abstract extraneous information from the surrounding volume.

The final step in our pipeline is to blend all the previously generated images together and display them to the screen. This is accomplished by rendering multiple view aligned quads positioned so that they exactly fill the viewport. By texturing these quads with the generated images we can utilise the GPU's fixed function pipeline to perform blending and compositing between the images. Our two level approach can be applied to multiple volume rendering techniques including high dimensional transfer functions as shown in Figure 10(b).

### *3.1. Volume Rendering*

In order to apply different abstractions to the two levels of the volume we modify a standard 3D texture slice volume renderer to output two images: one of the areas of the volume that the user has chosen to focus on and one of the remaining sections of the volume. We modified the basic algorithm so that the different segments of the volume are rendered in separate passes, enabling us to use custom rendering parameters for each segment in the volume. We further modify the output images from each pass by applying various image-space filters as a post-processing step. In a final stage the two images are blended together and rendered to the screen resulting in a complete volume view. If no modifications are performed on the intermediary images, the resulting volume rendering is visually identical to a traditional one-pass render.

#### *3.1.1. First Rendering Pass*

Our first pass is used to generate a solid rendering of the user selected isosurface. We flag voxels which are part of the isosurface and render these voxels as normal. The pixel shader utilises multiple render targets in order to render multiple images in one pass: we generate an image using phong lighting rendering parameters; an image using phong lighting without any specular terms; and also a depth image.

It was found that the majority of image-space edge filters used to perform perceptual enhancement gave optimum results when using a volume rendering with only ambient and diffuse lighting, however the addition of specular highlights produce a more aesthetically pleasing final image. Thus we render both images and use the rendering with specular highlights for the final blending step while using the diffusely lit image for image-space line extraction. The depth image is used to perform depth culling for both the object-space lines and the second volume rendering pass.

#### *3.1.2. Second Rendering Pass*

The second pass is used to render the remaining section of the volume. The naive approach would be to render all parts of the volume that were not displayed previously. This does follow the same approach as used by a normal one-pass renderer however it neglects to deal with occlusions that would normally be calculated by the GPUs fixed function pipeline. The isosurface rendered in the first pass would normally occlude various parts of the volume, however since we are not rendering the isosurface in the second pass, no occlusion calculations occur and areas of the volume are rendered when they shouldn't be.

It is possible to avoid this problem by using a custom pixel shader which uses the depth buffer generated in the first pass to perform depth testing.

This discards any voxels which would normally be occluded and results in the correct final image.

### 3.2. Non-Photorealistic Rendering Techniques

Edge stylisation can play an important part in NPR. Emphasizing edges in a scene enhances distinction between regions and can increase the saliency of a particular object or area. Five edge detection techniques were implemented for testing with the system. Three of these were image-space edge detectors; Sobel, Difference of Gaussian and Canny, while the other two were object-space edge detection techniques; silhouettes and suggestive contours. The three image-space detectors were of varying complexity, as were the two object-space edge detectors. This allowed direct comparison between techniques and therefore we were able to determine whether more expensive techniques outperformed the simpler, faster methods in tests.

#### 3.2.1. Object-space Edge Detectors

Our current implementation for extracting object-space lines uses a parallelised version of the line extraction algorithm as described by Burns et al [24]. We were able to optimise this algorithm in a number of different ways in order to provide real-time navigation of the volume dataset with object-space line enhancement.

We perform a pre-processing step which calculates the normal value for each voxel to a high accuracy using a 3D mask based on the Sobel filter. This information is stored for use by the line drawing calculation and converted into a 3D texture and uploaded to the GPU to be used in shading calculations for the volume rendering. By performing this step as a pre-process and avoiding repeated GPU calculation every frame we were able to significantly increase the performance of our line extraction.

A problem that frequently occurs in highly parallel algorithms is that the overhead of creating and destroying threads can be significant in comparison to the actual work that each thread performs. We employ a thread pool structure that allows threads to continually pull work jobs from a centralised queue using hardware supported x86 atomic operations, this structure eliminates the overhead of creating new threads and our use of atomic operations reduces contention for the shared queue.

One advantage of the serial algorithm of Burns et al. is that it exploits locality of reference since the serial nature of their algorithm increases the likelihood that a large amount of the data required for the current calculation will have been loaded in to the CPU cache from previous steps. We have modified our parallel approach in order to exploit this paradigm by splitting the volume into groups of neighbouring voxels, each of which represents a discrete work item thus allowing us to reduce contention on the work queue and also exploit cache coherency as the voxels are processed in spatially localised batches.

*3.2.2. Peripheral Abstraction Techniques*

Previous work by Redmond et al. [17] [18] has shown that a user's perception of a scene can be altered by not only adding detail to the target object, but also by abstracting extraneous information. In our system, this is achieved by implementing two varied abstraction styles chosen due to their success in past work: Saturation variation and the Kuwahara filter. Both techniques were implemented using GPU shaders and achieve real-time frame rates.

The Kuwahara filter [29] produces painterly like results through textural abstraction. The filter has the effect of smoothing internal regions while keeping edges sharp and has been shown to produce perceptual improvements in experiments from previous papers [17]. Saturation abstraction alters the saturation component of each pixel to reduce the colour of objects that are not of importance. These two filters abstract unimportant parts of the scene leaving important areas in full detail which therefore increases their saliency.

## 4. Evaluation

*4.1. Saliency Tests*

An automatic saliency metric [30] was used to evaluate whether adding stylisation can increase the saliency of important elements while decreasing the saliency of extraneous detail within the scene. The metric uses a multi-scale approach for finding the location of salient regions based on local image structure. Regions of high contrast are found for colour, intensity and orientation of an image on a number of scales. The results are then normalised and summed to create the final saliency map for an image.

A number of images were taken from the system using the styles described in Section 3 and each was compared with unstylised renderings. Each set of images was taken from identical viewpoints. The images were then tested using the saliency metric and while results were somewhat mixed, it was seen that by adding stylisation to the images, the saliency of the isosurface within

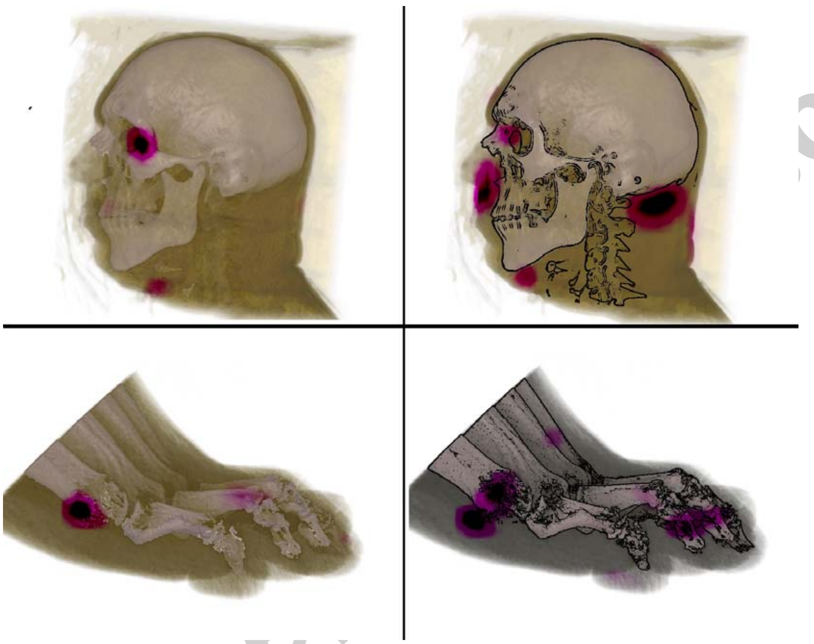

Figure 3: Examples of how adding stylisation can affect colour contrast within scenes. The left images are normally rendered while the right images use the Sobel filter (top) and a combination of suggestive contours and saturation variation (bottom). Colour contrast is indicated in purple and measured using the automatic saliency metric presented by Itti et al. [30]. These images show how the saliency of the target isosurface can be increased and manipulated using various styles.

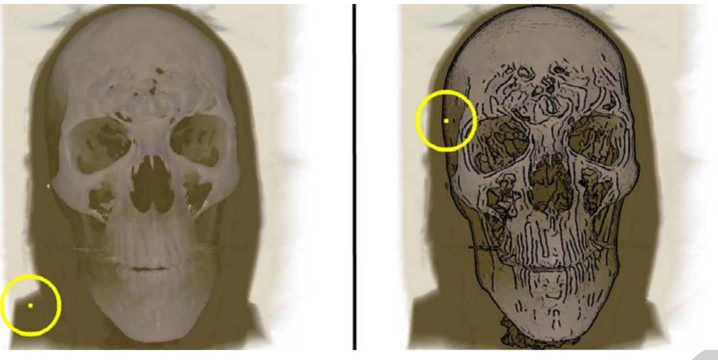

Figure 4: Example of how using stylisations can change the most salient point in the image, marked with a yellow circle in each image; Normal rendering (left) and Canny (right). As can be seen, adding stylisation can change the most salient point to the target isosurface.

the image can be increased while the impact of the surrounding volume can be reduced. This is especially true for increasing colour contrast within images using edge detectors, as can be seen in Figure 3. There was no large difference in how the edge detectors performed against each other although suggestive contours increased the saliency within the centre of the isosurface more than others. Figure 4 shows that most salient point of each image (marked by a yellow circle) can be changed using these stylisations.

These results from the metric show that by adding stylisation, visual saliency within scenes, and therefore user eye-gaze behaviour, can be manipulated. Previous authors [15] [31] have shown the benefits of gaze direction in emphasizing details computer imagery. While the automatic metric is based on a low-level model of visual interest, the results show that the stylisations used could make scenes faster to comprehend as eye-gaze is drawn to the most important parts of the scene.

*4.2. User Study*

An experiment was performed to investigate how each style affected a user's ability to determine the shape of the target isosurface. The primary aim of the visualizations is to use all volumetric data to give a clear impression of a certain part of each dataset while showing the rest of the object data for reference. For this reason it was necessary to test how well each style represented each dataset which meant a shape evaluation experiment was

Figure 5: An example image shown to participants in the user study. The controllable gauge can be seen in red.

the most suitable test and it was chosen over task-based or eye-tracking experiments. We hypothesised that by adding edges and abstractions to the dataset, the background volume would be de-emphasized and the target isosurface level would be clearer to the user, therefore increasing a users understanding of the dataset.

To test a user's perception of shape, an experiment was run which involved placement of gauges on static images. This protocol was described in previous experiments such as [32] and [23]. Participants were shown a series of images from the system and asked to rotate gauges which overlaid the images to match the surface normal of the isosurface at that point. Participants had no control over gauge position, only orientation. The orientation of the gauge was controlled by the mouse and the space bar was used to indicate that the participant was happy with the gauge position and ready to move on to the next trial. Each gauge is drawn as an ellipse and a single line as seen in Figure 5.

Gauge placement was pre-determined and each gauge was placed in an area of interest for each dataset. Three diverse datasets were used in the experiment. These were scans of a male head, a brain and a foot. The scale of each dataset in each image was constant although the rotation of the dataset was changed slightly to avoid any learning curve with regards to the normals. Each image was shown 8 times with different gauge placements in each trial to fully test a user's shape perception with each style. This resulted in a total of 360 trials. 14 naive participants took part in the experiment (10M-4F), each with a knowledge of graphics and therefore surface normals.

There were 15 styles tested on each dataset, which resulted in 45 images, a sample of these images can be seen in Figure 11. These styles can be seen listed below. Each of the 2 object-space line drawing, 3 image-space edge detection and 2 abstraction techniques were tested individually on each dataset. It was also tested how combining the image-space and object-space edges would affect user estimations, which resulted in 6 additional styles. Also tested was a style containing suggestive contours, Difference of Gaussian edges and Kuwahara painterly effect. This style was added to determine how using 3 different stylisation types together could affect user perception of an image. Finally a set of images using the two-level volume rendering system was run with no stylisation at all as a baseline comparison. The two-level rendering was chosen as a baseline instead of a single shaded isosurface as our focus is on investigating the effectiveness of different NPR techniques in order to emphasise and abstract data. A single isosurface only has one area of interest and does not allow for both emphasis and abstraction to be applied at the same time.

Style 1: Sobel edge detection

Style 2: Difference of Gaussian (DoG) edge detection

Style 3: Canny edge detection

Style 4: Silhouettes

Style 5: Suggestive contours

Style 6: Sobel / Silhouettes

Style 7: DoG / Silhouettes

Style 8: Canny / Silhouettes

Style 9: Sobel / Suggestive contours

Style 10: DoG / Suggestive contours

Style 11: Canny / Suggestive contours

Style 12: DoG / Suggestive contours / Kuwahara filter

Style 13: Kuwahara filter

Style 14: Saturation variation

No Style (Normal rendering)

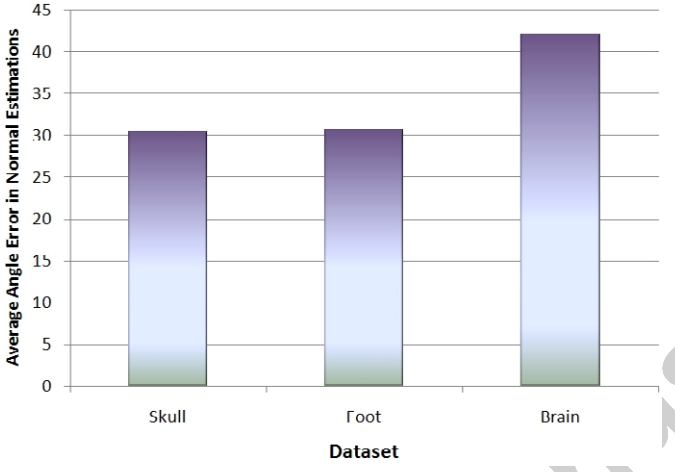


Figure 6: Average normal estimation errors for each dataset.

### *4.3. Results and Analysis*

For each trial the angle between the correct surface normal and the estimated surface normal was calculated. The average angle was then calculated for each style across each dataset and an ANalysis Of VAriance (ANOVA) was then performed on the results where the conditions were *Dataset*(3) and *Style*(15). It was found that there was a main effect of dataset where the type of dataset shown affected user accuracy ($F(2, 26) = 15.253, p < 0.0005$). Post-hoc analysis was then performed using a standard Newman-Keuls test for pairwise comparison among means. As can be seen from Figure 6, error was significantly higher for the brain dataset than the skull or foot datasets ($p < 0.0003$ in both cases). Several participants noted after the experiment that they had prior knowledge of what shape a skull or foot skeleton should be, whereas they had no knowledge of the localised shape of a brain dataset. Therefore, we hypothesise that this previous knowledge affected the experiment results.

It was also seen that there was a main effect in style ($F(14, 182) = 7.6138, p < 0.00001$), where certain styles performed significantly better than others. Post-hoc analysis was performed on the data using a Newman-Keuls test and it was seen that the angle estimations from images with Difference of Gaussian edges, Canny edges ($p < 0.015$ in both cases) and suggestive

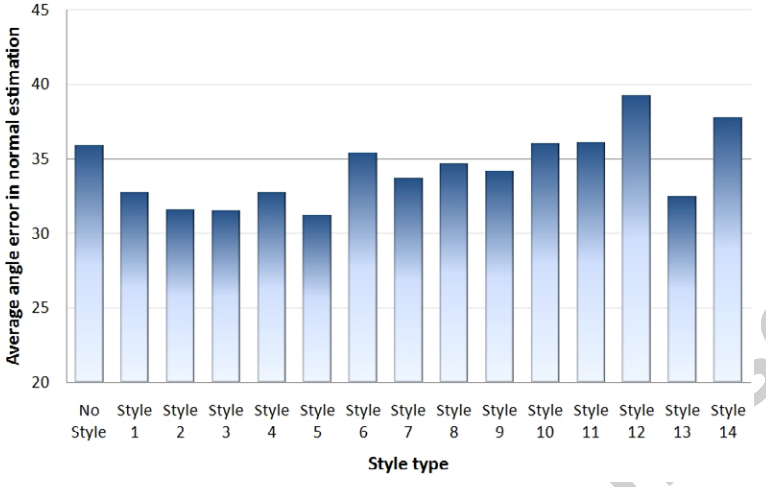


Figure 7: Average normal estimation errors for each style type over all datasets.

contours ($p < 0.007$) were significantly better than those from renderings with no stylisation. These differences can be seen in Figure 7. Results also show that estimations from the images which used either basic silhouettes or Sobel edges did not have any significant difference than those from the unstylised renderings. This was somewhat expected as Sobel edge detection is a far simpler and faster image-space technique than either Difference of Gaussian or Canny edge detection. Similarly, object-space silhouette detection is a much faster method than suggestive contours. These results imply that speed cannot be traded off if reliable and effective results are to be obtained.

Results also showed that there was no significant difference between results from images with no stylisation and those both with the Kuwahara filter and saturation variation. This shows that abstracting the background volume data does not have a large effect on user perception of the main volume data over all datasets. It can also be seen from our results that combining multiple styles had no significant impact on user accuracy. This result is quite surprising, especially in the case of combining suggestive contours with Difference of Gaussian edges or Canny Edges, as they performed well individually. However, the results may be explained by the fact that adding too many edges of different types may clutter the object and obscure vital shape cues.

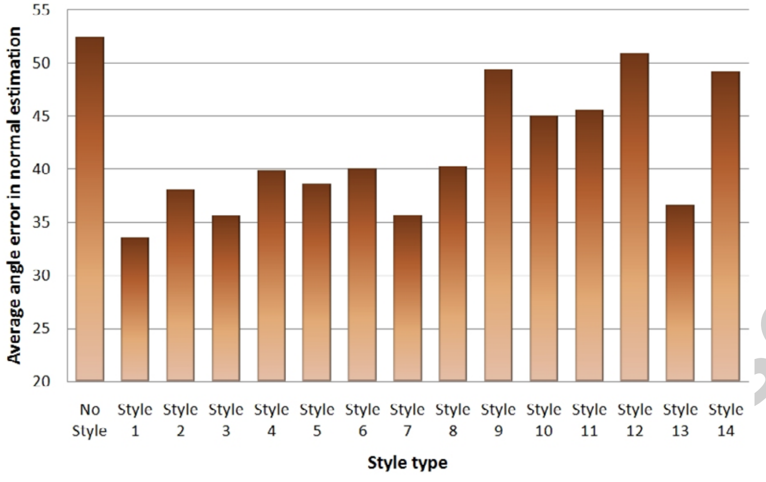


Figure 8: Average normal estimation errors for each style type for the brain dataset.



It was also found there was an interaction between dataset and style ($F(28, 364) = 13.576, p < 0.00001$) which indicates that certain styles performed better within certain datasets. To investigate this effect, the results from each style were isolated and compared. While it was found that there was no change in the styles which performed well for the foot and skull datasets, a number of styles performed significantly better than normal rendering in the brain dataset. As Figure 8 illustrates, all styles apart from the mixed style, saturation variation and the combination of Sobel edges and suggestive contours performed significantly better than normal rendering in this set ($p < 0.02$ in all cases). Within the brain dataset there was much occlusion of the isosurface from the background volume data, this meant there were less cues for shape than in the other datasets and the styles were therefore more effective in conveying shape. Also, as mentioned earlier, participants were less familiar with the shape of the brain dataset so the cues added made more of a difference than in the other sets.

It was noted that there was an effect of gauge orientation in the experiments. It was found that if the correct gauge position was less than 15 degrees from the camera view vector then user accuracy was significantly increased. If the angle between the correct gauge position and the camera view vector was above 60 degrees then the gauge estimation was significantly

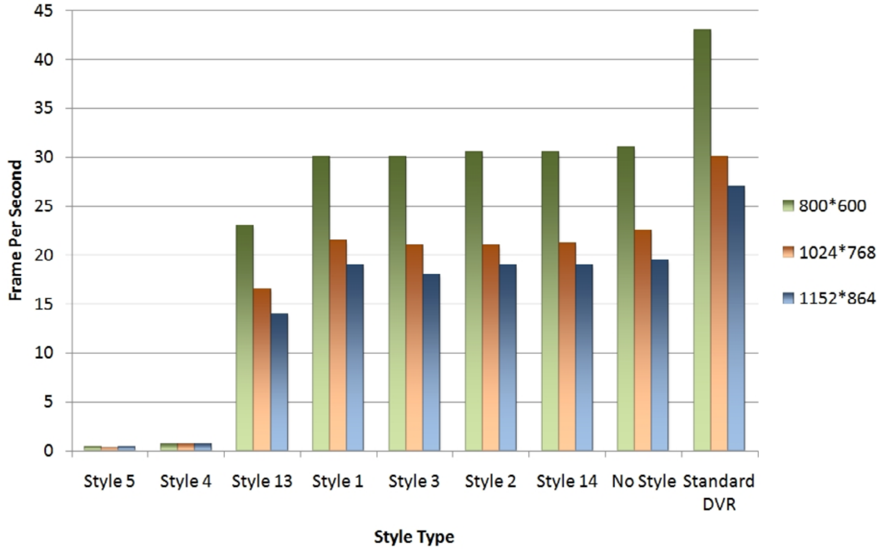


Figure 9: Average frames per second for the brain dataset at varying resolutions and styles.

harder. There was no significant difference in accuracy if the angle was between 15 and 60 degrees, which represented approximately 80 percent of the gauges tested in the experiment. Within the experiments the gauge angles tested were distributed randomly which means that more experiments would be necessary to determine what effect, if any, gauge position has on how each style performed. This is an interesting result, although it is unexpected as such an effect was not noted in previous work which used the same system [23].

Figure 9 shows the average frames per second for the brain dataset at varying resolutions and styles. Our two pass modification incurs a computation penalty of between 16-40% in comparison to an unmodified volume render but allows custom rendering parameters and post-processes to be applied separately to each volume segment. Currently our implementation is a proof of concept with no optimisations and several areas where duplicate calculations are performed and as such the frame times in Figure 9 are not reflective of the performance of our technique. In the future we are planning to implement our algorithm using a ray tracing based volume renderer which we believe will dramatically reduce the cost of segmenting the dataset. Once the volume has been segmented, applying the majority of the image-space NPR filters is negligible, however both the object-space lines and Kuwahara filtering incur a relatively large performance penalty.

## 5. Conclusions

We have presented a system which integrates DVR and NPR components into a two-level volume rendering system for various types of datasets. The DVR system is implemented on the GPU and employs parallelism to provide improved rendering speed for interactive visualisation. The NPR components include object-space silhouettes and suggestive contours, implemented on the CPU, as well as image-space edge detection and textural abstraction, implemented as shaders on the GPU. A two-pass mechanism in the DVR system generates two separate images that are variably abstracted, post-processed and then merged to create a combined rendering featuring the various stylistic enhancements.

We also carried out a number of experiments that show that our technique is capable of increasing object saliency and improving user understanding of shape whilst preserving context with the surrounding peripheral data. Results from the experiments showed that certain edge detectors can have a significant effect on a user's perception of object shape. We have shown that combining multiple edge detectors does not result in significantly better results than when the edge detectors are applied individually. It was also shown that the techniques described here are especially effective for datasets which contain a large amount of peripheral volume data, and also for datasets which users are not familiar with.

Future work will involve us integrating more detailed DVR techniques such as ray-casting into our current framework which we believe will greatly increase the performance of our system and also enable potentially higher-fidelity in the rendered image. We would also like to investigate further speed optimisations to achieve high interactivity so that the system might be used to render temporally variant volume data sets such as 4D MRI and 3D ultrasound. Additionally we would like to test the system with a wider range of datasets, as the approach is suitable for a wider variety of domains. An example of a technical dataset using our approach can be seen in Figure 10, where an engine block is rendered with internal components highlighted.

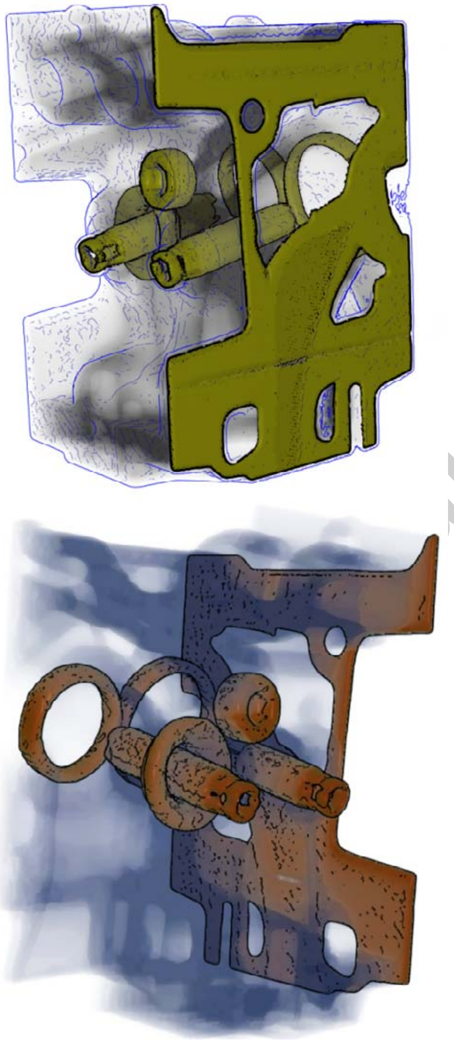

Figure 10: Two avenues of research we are investigating, line enhancements for the surrounding volume and multi-dimensional transfer functions. The first image shows an engine block with the target isosurface highlighted using Difference of Gaussian edges and the surrounding volume enhanced using suggestive contours. The second image shows the target isosurface highlighted using suggestive contours and the Kuwahara filter while also using a multi-dimensional transfer function.

This image also shows future avenues of research, line enhancements for the surrounding volume and multi-dimensional transfer functions.

We have performed user experiments that show an increase in user understanding of surface shape at different parts of the visualisation but further experiments could be performed to compare different textural and illumination stylisations. It would also be of interest to conduct studies which determine what levels of abstraction are most effective as well as determining how the two levels should be best combined to ensure the target isosurface level is clear while ensuring it does not seem separate from the object. Furthermore we intend to do studies on incorporating additional NPR stylisations into our system including highlights, halos, more complex painterly effects and line based shading such as curvature dependent hatching and stippling. Finally, we have employed two levels of abstraction for the in-focus and contextual background data however we would like to investigate if more than two levels of abstraction across the image can be exploited to provide fine-grained cues for more control of the saliency of different parts of the volume.

**Acknowledgements**

This research is funded by the Science Foundation Ireland Research Frontiers grant SFI08/RFP/CMS1076.

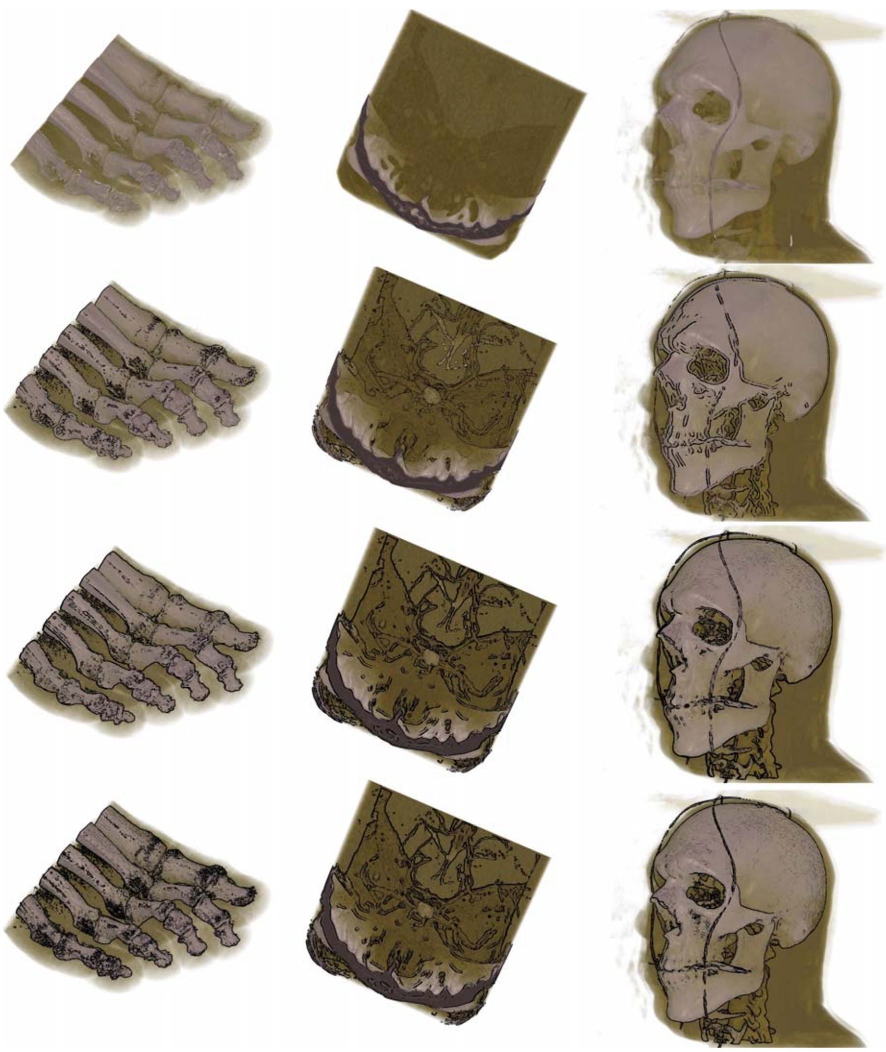

Figure 11: Some sample image from our system: First Row - Normal rendering; Second Row - silhouettes, suggestive contours, Canny; Third Row - Sobel, Canny combined with silhouettes, Sobel combined with suggestive contours; Fourth Row - Difference of Gaussian combined with suggestive contours, Canny combined with suggestive contours, Difference of Gaussian combined with suggestive contours and Kuwahara filtering.

## References


[1] H. Hauser, L. Mroz, G.-I. Bischi, E. Gröller, Two-Level Volume Rendering-Fusing MIP and DVR, in: VISUALIZATION '00: Proceedings of the 11th IEEE Visualization 2000 Conference (VIS 2000), IEEE Computer Society, Washington, DC, USA, ISBN 0-7803-6478-3, 2000.

[2] J. Kruger, J. Schneider, R. Westermann, ClearView: An Interactive Context Preserving Hotspot Visualization Technique, IEEE Transactions on Visualization and Computer Graphics 12 (5) (2006) 941–948, ISSN 1077-2626, doi:http://dx.doi.org/10.1109/TVCG.2006.124.

[3] M. Chen, A. Kaufman, R. Yagel (Eds.), Volume Graphics, Springer-Verlag New York, Inc., Secaucus, NJ, USA, ISBN 1852331925, 2001.

[4] B. Preim, D. Bartz, Visualization in Medicine: Theory, Algorithms, and Applications (The Morgan Kaufmann Series in Computer Graphics), Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, ISBN 0123705967, 2007.

[5] T. J. Cullip, U. Neumann, Accelerating Volume Reconstruction With 3D Texture Hardware, Tech. Rep., Chapel Hill, NC, USA, 1994.

[6] B. Cabral, N. Cam, J. Foran, Accelerated volume rendering and tomographic reconstruction using texture mapping hardware, in: VVS '94: Proceedings of the 1994 symposium on Volume visualization, ACM, New York, NY, USA, ISBN 0-89791-741-3, 91–98, doi: http://doi.acm.org/10.1145/197938.197972, 1994.

[7] M. Hadwiger, C. Sigg, H. Scharsach, K. Bhler, M. Gross, Real-Time Ray-Casting and Advanced Shading of Discrete Isosurfaces, Computer Graphics Forum 24 (3) (2005) 303–312.

[8] K. Engel, M. Kraus, T. Ertl, High-quality pre-integrated volume rendering using hardware-accelerated pixel shading, in: HWWS '01: Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware, ACM, New York, NY, USA, ISBN 1-58113-407-X, 9–16, doi:http://doi.acm.org/10.1145/383507.383515, 2001.

[9] I. Viola, A. Kanitsar, M. E. Gröller, Importance-Driven Feature Enhancement in Volume Visualization, IEEE Transactions on Visualization and Computer Graphics 11 (4) (2005) 408–418.

[10] C. T. Bernhard Preim, C. Dorge, NPR, Focussing and Emphasis in Medical Visualizations, Simulation and Visualisierung (2005) 139–152.

[11] Y. Kim, A. Varshney, Saliency-guided Enhancement for Volume Visualization, IEEE Transactions on Visualization and Computer Graphics 12 (5) (2006) 925–932.

[12] N. Halper, M. Mellin, C. S. Herrmann, V. Linneweber, T. Strothotte, Psychology and Non-Photorealistic Rendering: The Beginning of a Beautiful Relationship (2003) 277–286.

[13] A. Santella, D. DeCarlo, Stylisation and abstraction of photographs, in: SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques, ACM Press, 769–776, 2002.

[14] A. Santella, D. DeCarlo, Visual interest and NPR: an evaluation and manifesto, in: NPAR '04:, ACM Press, 71–150, 2004.

[15] F. Cole, D. DeCarlo, A. Finkelstein, K. Kin, K. Morley, A. Santella, Directing Gaze in 3D Models with Stylized Focus, in: Eurographics Symposium on Rendering, 377–387, 2006.

[16] N. Redmond, J. Dingliana, A Hybrid Technique for Creating Meaningful Abstractions of Dynamic Scenes in Real-time, in: Proceedings of WSCG'08: The International Conference in Eastern Europe on Computer Graphics, Vision and Visualisation, 2008.

[17] N. Redmond, J. Dinglana, Influencing User Attention Using Real-Time Stylised Rendering, in: TPCG 2009: Proceedings of Theory and Practice of Computer Graphics Conference, Eurographics UK Chapter, 2009.

[18] N. Redmond, J. Dingliana, Investigating the effect of real-time stylisation techniques on user task performance, in: APGV '09: Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization, ACM, New York, NY, USA, ISBN 978-1-60558-743-1, 121–124, doi:http://doi.acm.org/10.1145/1620993.1621017, 2009.

[19] D. DeCarlo, A. Finkelstein, S. Rusinkiewicz, A. Santella, Suggestive contours for conveying shape, ACM Trans. Graph. 22 (3) (2003) 848–855, ISSN 0730-0301, doi:http://doi.acm.org/10.1145/882262.882354.

[20] T. Judd, F. Durand, E. Adelson, Apparent ridges for line drawing, ACM Trans. Graph. 26 (3) (2007) 19, ISSN 0730-0301, doi: http://doi.acm.org/10.1145/1276377.1276401.

[21] D. DeCarlo, S. Rusinkiewicz, Highlight lines for conveying shape, in: NPAR '07: Proceedings of the 5th international symposium on Non-photorealistic animation and rendering, ACM, New York, NY, USA, ISBN 978-1-59593-624-0, 63–70, doi: http://doi.acm.org/10.1145/1274871.1274881, 2007.

[22] F. Cole, A. Golovinskiy, A. Limpaecher, H. S. Barros, A. Finkelstein, T. Funkhouser, S. Rusinkiewicz, Where Do People Draw Lines?, ACM Transactions on Graphics (Proc. SIGGRAPH) 27.

[23] F. Cole, K. Sanik, D. DeCarlo, A. Finkelstein, T. Funkhouser, S. Rusinkiewicz, M. Singh, How Well Do Line Drawings Depict Shape?, in: ACM Transactions on Graphics (Proc. SIGGRAPH), vol. 28, 2009.

[24] M. Burns, J. Klawe, S. Rusinkiewicz, A. Finkelstein, D. DeCarlo, Line drawings from volume data, ACM Trans. Graph. 24 (3) (2005) 512–518, ISSN 0730-0301, doi:http://doi.acm.org/10.1145/1073204.1073222.

[25] S. M. F. Treavett, M. Chen, Pen-and-Ink rendering in volume visualisation, in: VIS '00, IEEE Computer Society Press, 203–210, 2000.

[26] E. B. Lum, K.-L. Ma, Hardware-accelerated parallel non-photorealistic volume rendering, in: NPAR '02: Proceedings of the 2nd international symposium on Non-photorealistic animation and rendering, ACM, New York, NY, USA, ISBN 1-58113-494-0, 67–ff, doi: http://doi.acm.org/10.1145/508530.508542, 2002.

[27] M. Hadwiger, C. Berger, H. Hauser, High-Quality Two-Level Volume Rendering of Segmented Data Sets on Consumer Graphics Hardware, in: VIS '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03), IEEE Computer Society, Washington, DC, USA, ISBN 0-7695-2030-8, 40, doi:http://dx.doi.org/10.1109/VISUAL.2003.1250386, 2003.

[28] N. A. Svakhine, D. S. Ebert, W. M. Andrews, Illustration-Inspired Depth Enhanced Volumetric Medical Visualization, IEEE Transactions on Visualization and Computer Graphics 15 (1) (2009) 77–86, ISSN 1077-2626, doi:http://dx.doi.org/10.1109/TVCG.2008.56.

[29] S. E. M. Kuwahara, K. Hachimura, M. Kinoshita, Digital processing of biomedical images, Plenum Press, 1976.

[30] L. Itti, C. Koch, E. Niebur, A Model of Saliency-Based Visual Attention for Rapid Scene Analysis, IEEE Transactions on Pattern Analysis and Machine Intelligence 20 (11) (1998) 1254–1259.

[31] A. McNamara, R. Bailey, C. Grimm, Improving search task performance using subtle gaze direction, in: APGV '08: Proceedings of the 5th symposium on Applied perception in graphics and visualization, ACM, New York, NY, USA, ISBN 978-1-59593-981-4, 51–56, doi: http://doi.acm.org/10.1145/1394281.1394289, 2008.

[32] J. Koenderink, A. V. Doorn, A. Kappers, Surface Perception in Pictures, Perception and Psychophysics 52 (5) (1992) 487.