Noname manuscript No.
(will be inserted by the editor)

An experimental study on the effects of shading in 3D perception of volumetric models

Jose Díaz · Timo Ropinski · Isabel Navazo · Enrico Gobbetti · Pere-Pau Vázquez

Received: date / Accepted: date

Abstract Throughout the years, many shading techniques have been developed to improve the conveying of information in Volume Visualization. Some of these methods, usually referred to as realistic, are supposed to provide better cues for the understanding of volume data sets. While shading approaches are heavily exploited in traditional monoscopic setups, no previous study has analyzed the effect of these techniques in Virtual Reality. To further explore the influence of shading on the understanding of volume data in such environments, we carried out a user study in a desktop-based stereoscopic setup. The goals of the study were to investigate the impact of well-known shading approaches and the influence of real illumination on depth perception. Participants had to perform three different perceptual tasks when exposed to static visual stimuli. 45 participants took part in the study, giving us 1152 trials for each task. Results show that advanced shading techniques improve depth perception in stereoscopic volume visualization. As well, external lighting does not affect depth perception when these shading methods are applied. As a result, we derive some guidelines that may help the researchers when selecting illumination models for stereoscopic rendering.

Keywords Volume Illumination · Volume Visualization · Desktop-based Stereoscopy · Depth Perception

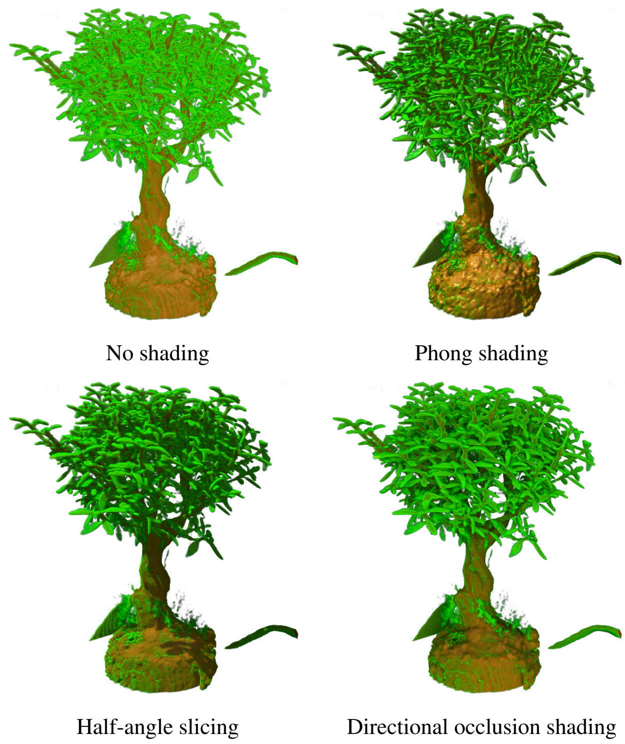

Fig. 1 The presented study analyzes the influence of four volume shading models on depth perception in stereoscopic images. The compared techniques are: no shading, Phong shading, half-angle slicing [14] and directional occlusion shading [30] .

1 Introduction

Volume rendering is a widely used technique that is nowadays applied in many fields, ranging from life sciences, through medicine to oil and gas exploration. It enables the interactive exploration of volumetric data sets, which are often acquired through advanced measurement techniques. As these data sets capture real world phenomena which are often of high geometric complexity, an effective exploration of these data sets is of uttermost importance. While most advanced volume rendering algorithms have been developed for desktop-based

J. Díaz and E. Gobbetti, ViC Group, CRS4 (Italy)
E-mail: jdiaz@crs4.it, gobbetti@crs4.it
T. Ropinski, Ulm University (Germany)
E-mail: timo.ropinski@uni-ulm.de
I. Navazo and P. Vázquez, VIRViG Group, UPC (Spain)
E-mail: isabel@cs.upc.edu, ppau@cs.upc.edu

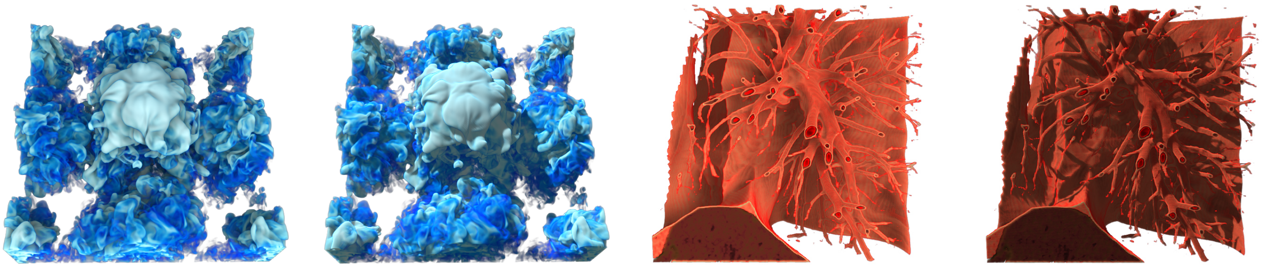

(a) Directional occlusion shading (b) Half-angle slicing (c) Directional occlusion shading (d) Half-angle slicing

Fig. 2 Two of the volume data sets used in the conducted study. (a) and (c) are rendered using the undirected directional occlusion approach, while (b) and (d) show an application of the light direction dependent half-angle slicing technique, whereby the light is positioned in the top left corner.

environments, recent hardware advancements also enable interactive volume rendering in virtual reality (VR) setups, where large screen resolutions and stereoscopic images make the rendering process more complex. Due to this application and the projected role of VR-based volume rendering in the future, it needs to be investigated in how far the benefits reported in desktop environments are also transferable to VR-based setups [17].

In the past years, one focus in the area of volume rendering was to improve the perceptual qualities of volume rendered images by incorporating advanced volumetric illumination models [13]. While the images resulting from advanced volumetric illumination algorithms are not only of higher fidelity (see Figure 1), it could also be shown that they improve scene perception as compared to standard volume rendering techniques [39, 21]. Unfortunately, as the reported findings are made in the context of monoscopic desktop-based environments, it is unclear whether the made conclusions can be transferred to stereoscopic setups. As recent studies on the perception of distances in virtual reality environments (VREs) for instance indicate that there is no correlation between visual quality and egocentric absolute distance perception [37], it is important to investigate if the made findings for volumetric illumination models are transferable to VREs. Furthermore, as many of the presented volume illumination techniques exploit approximations, it needs to be determined if a combination with stereoscopic projection reveals the resulting shortcomings. Therefore, in this paper we analyze the impact of advanced volumetric illumination techniques on the perception of volume rendered images in stereoscopic desktop-based environments. Within a user study, we investigate which would be the preferred illumination approach when considering scene perception. Additionally, as there is usually some real-world illumination present in VREs, we determine which role real-world lighting plays in such setups. As it is widely accepted that the illumination direction has an impact on scene perception [34], it is essential to find out how this extends to VREs, in order to support effective volume exploration. Thus, we have also analyzed the effect of real-world illumination apparent in a VRE on the perceptual qualities. To investigate these questions, we have conducted a user study with three independent tasks, where participants had to perform depth judgments based on volume rendered images. The images have been generated with four different illumination methods (see Figure 1), i.e., no shading, Phong shading, directional occlusion shading, and half-angle slicing, which can be considered as a representative subset of existing volume illumination techniques. Opposed to using no shading, Phong shading introduces local shading effects, while directional occlusion shading and half-angle slicing introduce more global effects. The fundamental difference between these two techniques is the fact that directional occlusion shading is independent of the light source direction, while half-angle slicing incorporates the incoming light direction (see Figure 2). Thus, the set of tested techniques does not only enable us to compare the techniques' individual perceptual benefits, but also to investigate the influence of incoming real-world lighting. To our knowledge, this is the first study investigating volume illumination within stereoscopic desktop-based environments in combination with real world illumination. We chose this concrete setup as it allows us to isolate the influence of the shading model with respect to the effects of head tracking, that, by itself, has a great impact in depth perception.

This work is a significantly extended version of our CGI2015 contribution [6]. Besides supplying a more thorough exposition, we provide here significant new material. In particular, we provide more details on the experimental setup to improve reproducibility of our experiments and verification of our results, we extend our analysis with an in-depth study of the main potentially confounding variables, we analyze time-related user performance as a function of shading models, we compare stereoscopic and monoscopic findings, and discuss the reasons that may lead some shading methods to influence perception more than others. In the remainder of this paper, we will first discuss related studies as well as relevant volumetric illumination techniques. Section 3 details our experimental setup and discusses the underlying research questions in greater detail. The achieved results and their implication for volume illumination in stereoscopic setups are discussed in Sections 4. In Section 5, we deeply analyze important aspects of the experiments, including marker positioning, difference between left and right eyes, luminance difference of the marked points, and relationships between depth map and shading. Section 6 discusses our findings. Finally, the paper concludes in Section 7, whereby we summarize our findings, we provide the derived guidelines, and state implications for current VREs.

2 Related work

Previous studies have shown that advanced illumination models are beneficial when perceiving a 3D scene [41, 18, 22]. However, to the authors' knowledge, no previous analysis has been carried out to stereoscopic systems. In this section we review the previous work related to our problem. First we deal with the evaluation of perception in Virtual Reality setups, and then we focus on different volumetric illumination techniques.

VR Evaluation. Many studies have investigated the usability of VREs, in order to understand the impact on the human user and to identify usability problems [5]. Used methodologies range from heuristic evaluation [35], over inspections and user testing [2], to metric-based approaches [9]. Despite this range of methods, usually the measurement of human behavior with respect to perception, action and task-performance are involved [23]. Many of the presented studies deal with visual perception of scenes presented in VREs [36, 8, 37, 3]. Thompson et al. have for instance investigated whether the quality of the displayed graphics in a virtual world affects the estimation of distances [37]. Therefore they have analyzed egocentric distance perception, whereby they focus on absolute distance measurements. For this scenario the authors report that accommodation, binocular convergence, linear perspective and familiar sizes are the primary cues for measuring distances. Based on their experiments they conclude that there is no significant correlation between image quality and distance perception. McMahan et al. present a study in which they have analyzed the impact of display and interaction fidelity in VR games [25]. They conclude that display and interaction fidelity significantly affect strategy and performance, as well as subjective judgments of presence, engagement, and usability. When analyzing search task performance in relation to visual realism, Lee et al. found only a significant difference for four of the sixteen search tasks they have analyzed [19]. It is not clear what differentiates these four tasks from the others, so further investigations are required to determine under which conditions realistic visual representations matter.

While most of the presented studies in VREs deal with representations of polygonal surface-based scenes, only until recently the importance of understanding the implications of volumetric representations has been acknowledged [15, 17, 16]. Laha et al. have investigated the effect of VR immersion on the visual analysis of volumetric data. Their results indicate that head-tracking and stereoscopy is beneficial when performing selected volume analysis tasks in a CAVE-like environment. Later, Laha et al. analyzed the impact on volume analysis performance, when using head-mounted displays [16]. Based on their study, they conclude that VR systems with a high field of regard combined with head tracking are helpful for visual search tasks involving volume data. Furthermore, they state that VR systems with fewer encumbrances might produce more significant benefits of higher immersion for visual task analysis with volume data sets. Based on the later finding, we have decided to use a 3D TV in the test setup described in this paper, while we avoid head tracking in order to exclude parallax effects, which have been shown to have a huge impact on the perception of volumetric data sets [4].

Besides considering VREs as an isolated space, several authors focus on the impact of the real world environment with respect to scene perception and immersion. In augmented reality applications for instance, depth is one of the major concerns when fusing real and virtual objects [31].

While other authors report a correlation between presence and depth perception in 3D TV setups [12], our study more focuses on perceptual benefits, than immersion or presence. The goal is to bridge the gap between the findings made with respect to realism in VREs, and the impact of volume illumination models on task performance as reported in the visualization literature [39, 40, 21].

Volume Illumination Techniques. While standard volume rendering is already widely used and appreciated as an enabling technology, more advanced volume rendering algorithms promise even more effective volume exploration [21]. Initially, volumetric shading was computed on a local level, whereby the gradient was derived and used to enable directional shading effects [20]. Thus, the resulting images resemble the illumination effects expected from standard Phong shading [26]. Through a formalization of the underlying light transport theory [24], it became possible to develop more sophisticated volumetric illumination algorithms. However, in order to support interactive exploration, these techniques usually incorporate approximations to the underlying light transport equations. The class of ambient occlusion based techniques for instance, discards directional light components to enable a visibility based shading. While the first such approach was presented for isosurface rendering

Technique	Local/global	Illumination frequency	Light direction dependent	Approximation
No shading	-	-	-	emission/absorption only
Phong shading	local	low/high	yes	surface orientation only
Half-angle slicing	global	high	yes	approximated light direction
Directional occlusion sh.	global	low	no	light visibility only

Table 1 The four tested techniques vary in the way illumination is incorporated in the rendering process. This table shows the main features of the tested models according to different categories. The illumination frequency refers to the interaction between the model and the scene (high frequency models produce sharp effects for even small features of the volume (i.e., hard shadows) while low frequency models produce softer results that may seem more natural (i.e., soft shadows)).

only [33], more modern techniques enable real-time ambient occlusion for semi-transparent structures [11,28,29,7]. At the same time when the first ambient occlusion techniques were proposed for volume data, also directional illumination techniques were developed. The earliest and still most widely used technique is the half-angle slicing approach [14]. With this approach it becomes possible to compute directional shadows with hard shadow borders at interactive frame rates. To achieve this interactivity, the incoming light direction is approximated, by slicing the volumetric data set along the half-angle in between viewing and light direction. More recently, other slice-based illumination techniques followed [30]. As discussing all relevant techniques would be beyond the scope of this paper, we refer to a recent state-of-the-art report on this topic [13].

3 Experiment

This section describes the designed user study. Since we want to evaluate the effects of shading in depth perception of stereo visualizations, we did not consider other immersive tools such as head tracking, which have a huge impact on user's perception due to parallax effects [4], and may disrupt the results of the experiment.

3.1 Goals and hypotheses

As stated before, the main goal of the presented study is to understand the effects of shading in communicating depth cues when rendering volume models in Stereoscopic Desktop-based Environments (SDEs). Concretely, we want to evaluate the previous findings in monoscopic environments, which assess that advanced shading techniques enhance the perception of shape and depth. To perform the analysis we chose two realistic and two non-realistic shading techniques. The second goal of the study is to determine whether the non-complete darkness that often is present in VREs (due to half-open doors, windows, or computer screens) plays any role on the perception of depth.

These goals lead us to the following research questions:

1. *Does advanced volume illumination perform better at communicating depth cues for volume models in SDEs?*
2. *Does the relation between real and virtual lighting in such stereoscopic setups influence depth perception?*

which lead us to formulate the following hypotheses:

1. *[H1] Advanced volume illumination techniques have an influence on the perception of depth cues for the understanding of volumetric data sets in SDEs.* Since the effects and acceptance of hard shadows vs soft shadows may be quite different [1], we will test both separately and analyze if any of these techniques performs better than more classical methods.
2. *[H2] Real illumination affects depth perception in SDEs.*

In addition, we want to verify if the lack of coherence between real lighting and virtual lighting has an influence on depth perception in stereoscopic desktop-based setups. Although some shading techniques may not define a virtual light source, others do. We will evaluate this scenario for such techniques if [H2] is accepted.

The experiment is based on asking the users to perform depth judgments on pairs of points. Since these are placed at different depths, and they might correspond to different parts of the volume data set, we will also analyze other parameters (such as their relative position and depth) to see if they play a role that supersedes the shading model effect. However, our tasks are designed to evaluate the influence of shading.

3.2 Datasets and shading models

A mixture of CT models and synthetic data sets are used in the study. In total, users were exposed to six different models, half of them were CTs and the other half were synthetic. The generated stereo images are representative of volume visualization, since they present cluttered regions with complex shapes and combine semi-transparent with opaque layers. Note that we used specific data sets and added arbitrary rotations to them in order to reduce the possibility that a previous knowledge of the data or too evident shapes facilitate the recognition of depths aside from the proper perception. For this reason we did not consider well-known anatomic models such as parts of the body, which might drive the users to have conflicting inputs (from their proper visual perception and

Fig. 3 The users performed the experiment using a 3D TV and made the selections using a customized keyboard. We placed a lamp in the room at top left, to be aligned with the lighting direction of methods that use a virtual light position.

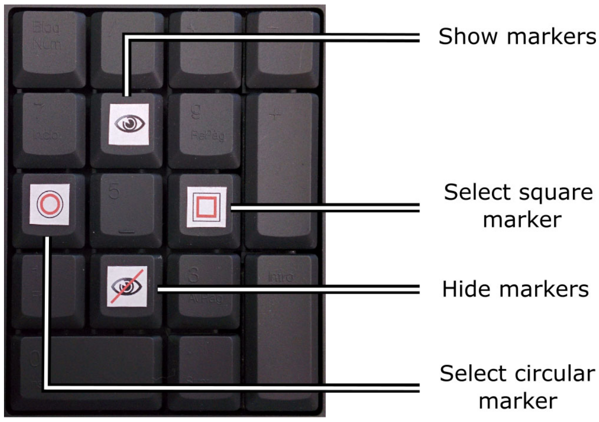

Fig. 4 The modified keyboard to facilitate the selection. By painting the keys in black we eliminate distractions, and the white labels facilitate the identification of the proper keys during the experiment, since the light is highly dimmed.

their previous knowledge on the shape of the structures) and thus yield unclear results.

The shading methods evaluated in the study have been selected among a wide range of traditional techniques. The reasons that guided our selection were two: popularity of the method and unique visual look. Since there is previous evidence in monoscopic visualization environments that physically-based rendering seems to yield better perception than local illumination models, we used two non-realistic approaches: emission-absorption model (NO shading) and Phong illumination (PH), and two advanced illumination techniques: half-angle slicing (HA) and directional occlusion shading (DOS).

- *NO shading:* one of the most basic forms of rendering a volumetric model is by assigning a color and opacity to the density of each voxel, and process rays only by taking into account the emission and absorption in the ray, without further processing. We used this method as the basis of our comparisons.
- *Phong shading:* it is probably the most common shading model used in Volume Visualization. Originally developed by Levoy [20], it implements the classical Phong model [26] by using volume gradients as surface normals. It requires placing a virtual light source.
- *Half-angle slicing:* shadows in general improve the perception of spatial relationships in scenes. From the available methods to simulate hard shadows, we chose half-angle slicing [14], since it is among the most popular shadowing techniques. It is based on slices and performs two rendering passes for each slice, one from the point of view of the observer and another one from the light source. The result of this method is an image with high frequency shadows such as in shadow mapping techniques. It also relies on placing a virtual light source.
- *Directional occlusion shading:* soft shadows are commonly more accepted than hard shadows [1] in VR and AR environments and they are perceived as more realistic. From the many existing techniques for volume shading, the method by Schott et al. [30] seems to be the technique that performs better in depth communication for monocular systems [21]. This method is also based on slices, and simulates shadows as would be produced by the light from a conical area originating at the observer position. The result is a scene with soft shadows that resemble ambient occlusion techniques.

In this way, we confront four different shading approaches that produce visually different images (see Figure 1). Two of them are more physically based (HA and DOS) and the other two are more empirical (NO and PH). Furthermore, two of them rely on the use of a directional light source (PH and HA) while the others have fixed light or no light position (NO and DOS). Table 1 summarizes the main features of the compared shading techniques.

3.3 Apparatus

For the experiment, we used a passive stereo system consisting on a 46” JVC 3D TV (GD-463D10 model) with polarized glasses. Moreover, we added a lamp placed top-left and some objects around the TV (two balls, a plastic glass and a couple of books) so that the lamp generates shadows of different sizes inside the participants’ Field of Regard (at least the peripheral vision was aware of such shadows, see Figure 3). The lamp position was chosen to be coherent with the virtual light source that is used in half-angle slicing and Phong shading (the position of the virtual light source is changed in task 3 to make it non-coherent with the lamp position). In order to facilitate the task, we customized a keyboard by painting

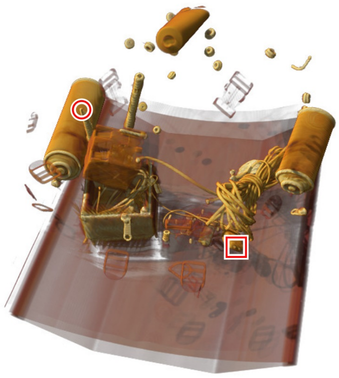

Fig. 5 Users were shown an image with two markers that identify two points in the model that must be classified. The markers disappear after popping up for three times and the users must determine which of the points is closer.

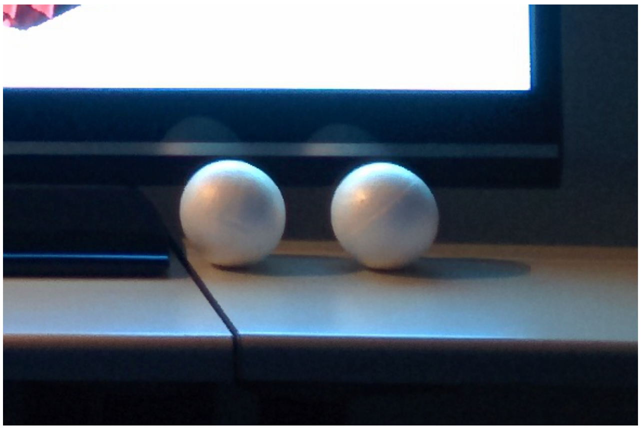

Fig. 6 A close-up of the objects placed in front of the TV, that cast obvious shadows over the table.

all the keys in black and putting some stickers to mark the necessary keys (see Figure 4). Users sat during the study (2 m from the screen) to avoid movements and thus, limiting the impact of the perspective distortion.

The application shows static stereoscopic images to avoid depth inferring via other elements such as the model motion. We created a small application that shows two markers in each image, and lets the users select the one which is placed closer to the observer. The markers are designed as small windows with two shapes: circular and square. They indicate the points of the model to be classified by the users. The markers are placed at the same distance from the observer, and users were instructed to classify the point that the markers show, not the markers themselves. These markers pop up three times, and then disappear. If necessary, the users may request the application to show the markers again. In Figure 5 we show one of the used images.

3.4 Design and procedure

The study comprises three different tasks. The first one evaluates the perception of depth using different shading techniques. Then, the second task evaluates the same methods in the presence of a controlled external light source (we use a lamp that casts obvious shadows inside the Field of Regard of the user). This task is tailored to determine if real illumination has any influence on the responses. Apart from the presence of the lamp, the experiment and the setup in both tasks are the same. Finally, the third task evaluates the influence on depth perception of having coherent real and virtual light directions. This task was meant to be done for those shading models that allow to change the virtual light position, only if the analysis of the previous tasks showed a real influence of external illumination on depth perception. The setup used in the three tasks is similar to others present in depth perception experiments [10] and in all cases comparisons were carried out mostly between users (only three users repeated on task 1 and 2).

Task 1. This task had the objective of evaluating the influence of different shading techniques in the perception of depth in volumetric models. For the development of the task, we used the four shading techniques illustrated in Figure 1.

Participants were shown 72 images covering 3 different points of view, four different techniques and six different models. Since the variable to analyze was the shading model, the images were systematically sorted based on the shading technique using Latin Squares, and the 18 images of each technique (variations of models and views) were then presented randomly. The users were introduced the task with a first example that had to be solved before the experiment effectively started. Everybody showed a proper understanding of the procedure and all the users completed their tasks to their satisfaction (as proved by post-hoc questionnaires). After each choice, the users had a neutral screen to let them recover from fatigue if necessary. Participants were instructed to determine the closer point from a pair indicated by two markers, and they were asked to take as much time as necessary. Markers popped up three times before the actual selection might start. In case of necessity, the users could make the markers visible again. The test was previously checked with two extra users whose results were not included in the data analysis. Throughout the experiment, we measured the correctness of the answers and the time spent in each choice. In Figure 5 we already showed the task as seen by the user.

Task 2. In order to determine the effect of external illumination, we performed the same experiment of the first task but with the lamp switched on. This adds obvious shadows that are in the field of regard of the users (see Figure 6) As in the previous case, we measured the correctness of the answers

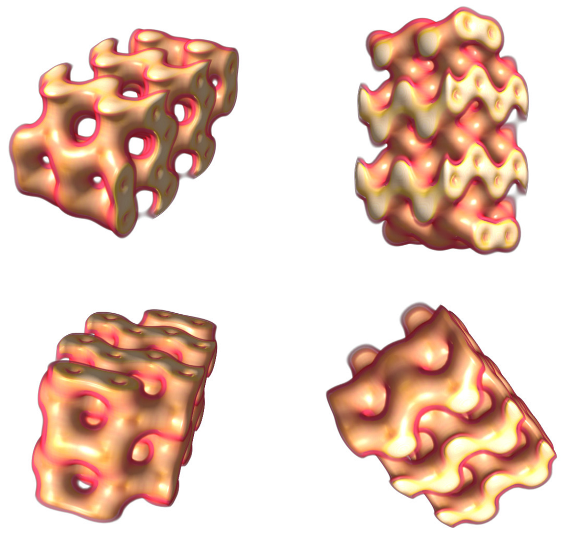

Fig. 7 Example of images shown in task 3. The displayed data sets were illuminated from different angles: the two on top are lit from top-left and top-right and the two at the bottom from bottom-left and bottom-right. Images were generated using Phong shading.

and time spent by the users, aside of the classical previous and post-hoc questionnaires.

Task 3. The goal of this task was to study whether conflicting directions between physical light source and virtual light sources leads to a worse perception of depth. The setup was the same as the previous task, but the images were generated using one out of four different light positions: top-left, top-right, bottom-left, and bottom-right. The users were presented images generated using Phong shading, which was the unique method affected by real illumination (see Section 4) that generate different images depending on the virtual light source. Figure 7 shows some examples of placing the virtual light source in different positions (top-left, top-right, bottom-left, and bottom-right, respectively). We used six different volume data sets, and generated three views for each of the four virtual light source positions. We systematically changed the order of the images for the users by using the light position as variable, so we applied a 4×4 Latin Squares scheme. The images in the same group were randomly sorted. A total of 72 images were shown to each user. Since the setup was the same as for the previous case, we only tested the study with an extra single user whose data was not used for the data analysis.

3.5 Participants

The study was carried out by 45 users. They had to fill in two different questionnaires: one before the task with personal information (age, gender, quality of eyesight, experience with VREs, etc.) and another after the task, asking about a subjective evaluation of the performed activity. After analyzing the post-hoc questionnaires, we found that all of them understood their task properly and found it easy to achieve. Therefore, there was no need of discarding any user.

Task 1. 16 volunteers (12 males, 4 females) took part in this task, none of them was color blind. They had ages between 24 and 56 (6 people did not reveal their age). Most of them were recruited among faculty and students, with some participants being from outside the university. All of them declared to perceive 3D correctly in VREs. Participants had from undergraduate to PhD degrees. All except one (who was not able to declare his ability to perceive 3D in VREs) declared good 3D perception when using stereoscopic displays. Most of them use the computer for above 4-6 hours a day, while two declared to spend around two hours a day. 72 images were shown in this task, which gave us 1152 trials.

Task 2. We recruited 16 volunteers (15 males, 1 female, none of them was color blind) among faculty and students of the university for this task, with ages between 23 and 56 (one person declined to declare the age). All of them presented a technical background and had already experience with some VR setups (cinema, 3D TV, CAVE, Powerwall, etc.). Except one participant (who was not able to assess his level of 3D perception), all declared good perception of 3D when exposed to stereoscopic displays. All the users except one commonly spend more than 6 hours on a computer (the other one works around 4 hours on a computer). All of them have an undergraduate, graduate or PhD degree. The users were shown 72 images, which led to 1152 trials.

Task 3. 16 users participated in the third task (13 males, 3 females, none of them was color blind), with ages ranging between 19 and 60 (three users did not reveal their age). As well as the previous tasks, time of response was measured throughout the experiment, although users were instructed to devote as much time as necessary to properly complete the experiment (as in the previous tasks). Like in the second task, people were selected among faculty and students, so their academic degree ranged from undergraduate to PhD. All of them spend 4 hours or more on a computer. Users were shown 72 images, which gave us 1152 trials.

4 Results

In this section we present the main results of our study. The mean correctness of the answers for each task was analyzed by using a one-way analysis of variance (ANOVA) with a significance level of $\alpha = 0.05$. When significant differences between the means were found, we used a post-hoc

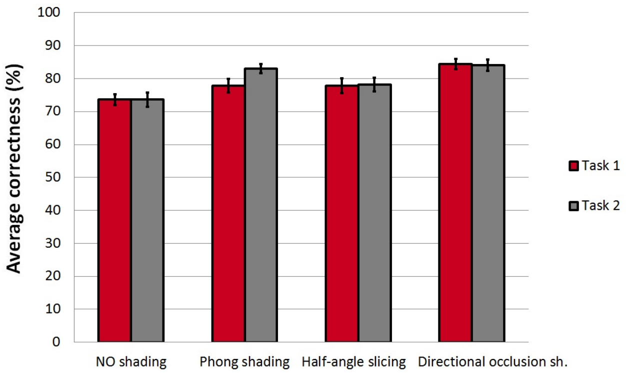

Fig. 8 Task 1 and task 2 results: average correctness of the user's answers for each individual shading technique with standard error.

Variables	χ^2	p-value	Correct answers for each depth category
T1: Relative depth vs. Users' answers	5.991	< 0.0001	$< 0.05 : 66\%$ $0.05 - 0.1 : 88\%$ $> 0.1 : 86\%$
T2: Relative depth vs. Users' answers	5.991	< 0.0001	$< 0.05 : 63\%$ $0.05 - 0.1 : 86\%$ $> 0.1 : 87\%$

Table 2 Results of the Chi square test of association for the categorical variables relative depth and users' answers from tasks 1 and 2. Results show that relative depth and user's answers are associated. Correct answers for each category of relative depth are also provided.

Bonferroni's pairwise test with the same significance level ($\alpha = 0.05$). The same analysis was performed with the means of the time spent by the users to complete each task. In order to test linear correlation between the measured variables, we used the Pearson's r statistic and assessed the linear model testing the regression coefficient β_1 with $\alpha = 0.05$. Finally, the Chi square test of association with a significance level of $\alpha = 0.05$ was used to analyze categorical variables.

4.1 Evaluation of Task 1 and Task 2

As stated before, we compared four different shading techniques in the first and the second tasks (T1 and T2): no shading (NO), Phong shading (PH), half-angle slicing (HA) and directional occlusion shading (DOS). The main difference between both tasks was the presence of real illumination in the second one. Results were analyzed individually for each task.

Analysis of shading models. The average correctness of the answers obtained from each task shows that users performed better with shaded images (PH, HA and DOS) than without shading (NO), independently on the presence of real illumination. Furthermore, after rejecting the null hypothesis that all correctness means were equal between techniques (ANOVA T1: $p = 0.002$, T2: $p = 0.0004$), Bonferroni's test revealed that DOS behaved significantly better (avg. correctness T1, T2: 84%) than NO (T1, T2: 73%) in both tasks and PH (T1: 77%, T2: 82%) performed significantly better than NO when real illumination was present. On the contrary, though the average correctness of HA (T1: 77%, T2: 78%) was higher than the average correctness of NO in both scenarios, no significant difference was found between them. The percentage of correct answers for each shading technique with its standard error is shown in Figure 8.

Influence of real illumination. The analysis of the correctness means of the answers for each individual shading technique with and without the external light shows no significant difference in the behavior of NO (ANOVA: $p = 1$), HA (ANOVA: $p = 0.91$) and DOS (ANOVA: $p = 0.87$). Instead PH performs significantly better (ANOVA: $p = 0.046$) in the presence of real illumination. This means that just PH seems to be affected by the external light, whereas the other methods perform in a similar way (see Figure 8). For this reason, and because of PH takes into account the light direction to shade the model (see Table 1), the third task only evaluates PH.

Avg. time and relative depth. The time spent by the users to perform each task was about 30 and 40 minutes, and no significant differences were found between shading techniques with respect to the elapsed time (ANOVA T1: $p = 0.20$, T2: $p = 0.34$). Instead, a moderate linear correlation was found between the average time and the average correctness of the answers in T1 (Pearson's $r = -0.62$, $p < 0.0001$), which means that users spent more time on images whith lower average correctness. When real illumination was present no linear correlation was found ($r = -0.08$, $p = 0.50$). Regarding the difference of depths between the markers and the average correctness of the answers, a weak linear correlation was found in both cases (T1: $r = 0.34$, $p = 0.004$, T2: $r = 0.29$, $p = 0.001$). Despite the weakness of this correlation, we observed that the number of correct answers depended on the relative depth between the markers. The bigger the difference of depth was, the higher the number of correct answers. The results of a Chi squared test showed this association between users' answers and relative depth for both tasks (see Table 2).

4.2 Evaluation of Task 3

In order to know whether conflicting directions between real and virtual light may influence on depth perception when applying Phong shading (PH), we confronted visualizations generated with the same light direction as the real one (EQ), placed at top-left position, and with different directions (DIFF), placed at top-right, bottom-left and bottom-right positions.

Shading technique and variables	χ^2	p-value	Correct answers for each depth category
Relative depth vs. Users' answers	5.991	< 0.0001	$< 0.05 : 72\%$ $0.05 - 0.1 : 92\%$ $> 0.1 : 99\%$

Table 3 Results of the Chi square test of association for the categorical variables relative depth and users' answers measured from task 3. As well as happened in the previous tasks, the test shows that the relative depth and user's answers are associated.

Analysis of lighting coherence. The ANOVA test for the mean correctness of user's answers showed no significant differences ($p = 0.98$) between having aligned real and virtual light directions (EQ) or disaligned ones (DIFF). The avg. correctness of the answers was 85% in both of cases.

Avg. time and relative depth. Users also spent between 30 and 40 minutes to complete this task and no significant differences were found (ANOVA: $p = 0.76$) with respect to the average elapsed time in both different scenarios (EQ and DIFF). Furthermore, as happened in the first task, the Pearson's r statistic revealed a moderated linear correlation between avg. time and avg. correctness ($r = -0.55$, $p < 0.0001$). As well, a weak linear correlation was found between depth differences and avg. correctness when using Phong shading ($r = 0.46$, $p < 0.0001$). The results of a Chi square test show that user's answers and the relative depth of the markers are associated (see Table 3).

5 Data Analysis

In this section we provide additional results extracted from the data collected during the study. First, we analyze the possible influence of other candidates to confounding variables, such as the distance of the markers, or the differences between the left and right eye. Then, we proceed to further analyze the information of the shading method, and see whether any of the shading models provides more information about the shape of the object.

5.1 Marker's positions

The position on the screen of the marked points was an important aspect to consider. We were curious to know if there would be any difference regarding depth perception when users can see at a glance the points to compare (they are close in screen coordinates) or they have to look to different screen positions. It is easier to perceive the relative depth between two points in the first case? To answer this question, we considered three different scenarios: *i)* checking for linear correlation between the difference in pixels of their x coordinates on the screen and avg. correctness (T1: $r = 0.03$, $p = 0.74$, T2: $r = 0.02$, $p = 0.82$), *ii)* the same with their y coordinates (T1: $r = 0.03$, $p = 0.74$, T2:$r = 0.06$, $p = 0.57$), *iii)* and the Euclidean distance between the markers (T1: $r = 0.03$, $p = 0.78$, T2: $r = 0.05$, $p = 0.66$). The results show that no linear correlation can be determined in any case. This suggests that the position on the screen of the markers did not influence depth perception.

5.2 Differences between left and right eye.

Since we generated two different images (one for each eye) to produce a single stereo visualization, one may wonder if the quantity of information provided by these images is related with user's depth perception. In order to check it, we computed the quantity of information for the images of left and right eyes by using the multiscale entropy metric [32]. Multiscale Entropy is a measure of quantity of information that uses a multiresolution pyramid to evaluate information contents in signals. It has demonstrated its utility measuring the illumination information in color images [38]. Thus, we used it to evaluate the amount of perceptible information contents in left and right eyes. As a consequence, we looked for a linear correlation between the obtained values and the average correctness of the answers obtained in tasks 1 and 2. Pearson's r statistic showed that there was no linear correlation in any task, neither with the left eye images (T1: $r = 0.07$, $p = 0.54$, T2: $r = 0.11$, $p = 0.33$) nor the right eye ones (T1: $r = 0.07$, $p = 0.53$, T2: $r = 0.11$, $p = 0.33$). As well, we performed the same analysis for the depth maps of the left (T1: $r = 0.08$, $p = 0.46$, T2: $r = 0.05$, $p = 0.63$) and right images (T1: $r = 0.08$, $p = 0.46$, T2: $r = 0.05$, $p = 0.64$). No linear correlations were neither found. These results point out that the quantity of information provided by the color images and the depth maps of both eyes seem to have no real impact on user's performance when judging depth relations of volumetric scenes.

5.3 Luminance difference of the marked points

Although our study was not designed to this end, we found interesting to check if the luminance difference of the marked points (5x5 pixels region around them) would affect depth perception. To do this, we extracted the luminance channel from the images showed in tasks 1 and 2 (where the four compared shading models where used) and we tested if there was a linear correlation between luminance differences and the average correctness of the answers. The results showed that there was no linear correlation neither in task 1 ($r = -0.02$, $p = 0.83$) nor task 2 ($r = 0.20$, $p = 0.08$). This suggests that the election of the closest point does not rely on luminance differences. It does not mean that luminance does not affect depth perception in stereo images. To know the role of luminance variations, a specific study should be conducted where not just the difference of luminance between the marked points but also the luminance of the whole scene should be analyzed.

5.4 Relationship between depth map and shading.

As we have seen in the previous sections, no other variable was detected as influencing the results of the experiment. Therefore, we have performed another extra study to determine how the shading itself is related to the geometry of the scene. In order to perform this study, we have isolated the changes produced by the shading to the images that would be generated by using only the transfer function. Since the previous results showed that let the TF alone, it does not seem to convey enough information on the geometry of the models, we have also analyzed what differences the applied shading present over the non-shaded images.

Although it is obvious from the previous experimental results that the adequate selection of the shading method leads to an improved perception of the model, the reasons are not clear. Some authors have argued that Ambient Occlusion is better to communicate shape because it darkens farther regions. However, to the authors' knowledge, no previous theoretical analysis supports this idea. To determine if the shading is key to improve perception, we have isolated the effect that produces on the image, besides the transfer function, and measured the added information.

To determine if there is some correlation between the contribution of the shading to the shape of the scene, we analyze the normalized mutual information [27] between the depth image and the shading contribution. Normalized Mutual Information is an information theoretic measure that is able to find correlations between images of different sources and has been successfully applied to register medical images from different sources (e.g. CT with MRI). To analyze our scenario, we generated images that contain the contribution of the shading model to the color as would be determined by the transfer function. In Figure 9 we show this information next to the depth map. So, for the NO shading method, the contribution would be null, and thus, the image does not contain information (Figure 9 (b)). On the opposite side, it is commonly argued that Ambient Occlusion provides better depth-related information. This can be seen in Figure 9 (e), where depth changes seem somewhat related to the changes in Directional Occlusion Shading contribution. The idea with this last experiment is to determine if we can measure such correlation.

It would be difficult to search for this correlation with the images of experiments 1 to 3, because all images are captured from different views to avoid learning effects. So we generated a new set of images. In this case, we simulated 12 different viewpoints, and for each viewpoint, we captured 4 images, each showing one of the shading methods. With this information, we analyzed the normalized mutual information between the shaded pixels and the corresponding depth values of those in the depth map, using Normalized Mutual Information (NMI).

The ANOVA analysis ($alpha = 0.05, p < 0.0001$) of the NMI values shows that there is a significant difference between the images shaded with DOS with respect to the images shaded using HA or PH. A further Bonferroni's test revealed that DOS provides a significantly higher NMI (avg $NMI = 3.327$) than HA (avg $NMI = 1.84$) and PH (avg $NMI = 1.88$). Instead, there's no significant difference between the NMI means of HA and PH.

Although these results are promising, there is still a way to go. To assess these findings, we plan to perform another experiment where participants perform depth-sensitive tasks. Then, we will analyze whether the differences in depth perception under varying illumination techniques correlate with the values found here.

6 Discussion

It is well-known that lighting and shadows provide monocular depth cues that enhance human perception of depth, distance and shapes. Previous studies [21] show that advanced volume illumination techniques may improve depth perception when visualizing volumetric data sets in traditional desktop setups. The main goals of our study were to check the effectiveness of these techniques and to understand the influence of incoming real-world lighting on stereoscopic desktop environments.

6.1 Summary of results

Tasks 1 and 2 compared four different shading models ranging from simple to advanced illumination methods: NO, PH, HA and DOS (see Section 3). If depth perception was just related to stereoscopy, we would have expected a similar performance among the compared shading approaches. Instead, results of both tasks show significant differences between them. Thus, we can state that depth perception is not just due to stereoscopy but also the shading plays a role, which leads us to accept *[H1]*. Furthermore, DOS (one of the advanced illumination techniques) performed significantly better than NO (simple shading model), so we can conclude that advanced volume illumination may improve depth perception with respect to more simple shading models. Though HA presents a better performance than NO, no significant difference was found during the analysis of each task. Being HA and DOS global illumination techniques based on

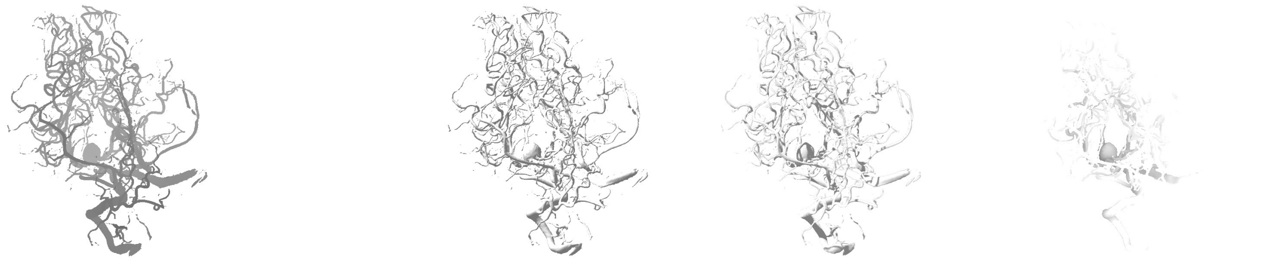

(a) Depth map (b) No shading (c) Phong shading (d) Half-angle slicing (e) Directional occlusion shading

Fig. 9 The images containing the shading contribution of the illumination models compared in the study were analyzed along with their associated depth maps to look for a possible relation between them. As one may expect, no shading contribution is present when using the NO shading model (b).

simulating shadows, one may expect a similar behavior between them. The fact that directional occlusion shading is independent on the light source direction (producing similar results to ambient occlusion methods) while half-angle slicing shadows totally depend on it, may explain the difference in performance between them. Regarding the influence of real illumination, the obtained results show that depth perception is not affected by the external light when using NO, HA and DOS. Which leads us to reject *[H2]* for these shading models. Instead, user's performance with PH was significantly better with respect to NO in the presence of real illumination. For this reason, PH is evaluated in the third task, where we analyzed the influence of having real and the virtual light directions aligned (EQ) and disaligned (DIFF). Results show that having inconsistent external lighting does not affect depth perception when using Phong shading in SDEs.

Concerning the elapsed time to complete the tasks it was interesting to know if any shading model allowed the users to perform the selection of the closest point faster than the others, and if there was any relation between the time and the correctness of the answers. The results showed no significant differences between the average time spent with the different shading models, which means that none of them seems to provide a clear advantage with respect to the others as a matter of elapsed time. Regarding the relation between the average correctness and the average time, a moderate negative linear correlation is present in the absence of real light. This means that users tend to spend more time to select the closest point in the images with higher average error, which were pressumably the more difficult choices. On the contrary, no linear correlation was found when real illumination was present. Could this fact give some clue about the different behaviour of PH in both scenarios? It may suggest that users in the first task were more focused on the test to solve than users on the second task, but if this was the reason, one would have expected a different behaviour also in the other shading techniques and not just when applying Phong shading.

In any case, the better performance of PH when real illumination is present deserves special attention. Which is the reason for this different behavior, whereas the other techniques are not influenced by the external light? In such immersive environments as the ones reproduced in our study, where the real light does not disturb the user when visualizing stereoscopic images, we would expect no real impact of external lighting on depth perception. Because the images shown in both tasks were the same and most of the participants where different in each one (just 3 users took part in both of them), we analyzed in detail the personal information provided by the users (stereoscopic perception, previous experience with VREs, etc.) to find a reasonable explanation of the PH results. Instead, no significant differences that could explain the different performance were found. We think that increasing the number of participants may reveal the same behavior for PH with and without real illumination, but further research is needed in this direction.

Regarding the relative depth between the points to compare and independently on shading models and lighting scenarios, the results of the three tasks showed a lower number of correct answers when the points were placed at similar depths, whereas the performance improved when the relative depth was higher. To prove the effectiveness of volume illumination techniques to enhance depth perception in stereo images, future studies should focus on evaluating the influence of lighting for points with similar depths.

6.2 Extension of monoscopic findings

As stated before, one of the motivations of the presented study was to know if the results obtained in monoscopic environments could be transferred to stereoscopy. To this end, the presented study was specifically designed to allow a direct comparison with the previous monoscopic evaluation of volume illumination techniques carried out by Lindemann and Ropinski [21]. For instance, PH, HA and DOS were considered in both studies, users had to judge the relative depth between two points of the scene, datasets and transfer functions were chosen to provide representative images of volume visualization, that were rendered using similar viewing and lighting conditions and finally, the same statistical analysis was performed.

The results of both studies show the same tendency in mono and stereo environments, being directional occlusion shading the technique that provides better depth perception. As well, Phong shading and half-angle slicing present a similar behaviour in both scenarios. The main difference between mono and stereo is the significance of the results. Whereas DOS (avg. correctness = 64.8%) behaves significantly better than PH (45.1%) and HA (44.8%) in mono environments, there are no significant differences between them in stereo. We believe that this is due to the strong depth cues that provide stereoscopy itself, which also explains the considerably better performance when having stereo vision (above 73% in the worst case). Concerning the average time employed to do the task, there are no significant differences between shading techniques neither with mono images nor stereo ones. Furthermore, no linear correlation between the average correctness and the average time is found in both cases with the same lighting conditions, that is in the presence of real light. In conclusion, we can state that the findings of previous evaluations in monoscopic environments related to volumetric illumination techniques can be extended to stereoscopic setups, but taking into account that the impact on depth perception of this models seems to be higher in monoscopic environments, since there are no strong additional depth cues like the ones provided by stereoscopy.

6.3 Derived guidelines

Based on the results of our study, we would recommend the following guidelines to improve the exploration of volumetric data sets in stereoscopic desktop-based environments:

- Using advanced volumetric shading improves depth perception: among the tested shading models, we would recommend the simulation of soft shadows by using directional occlusion shading for SDEs.
- Real illumination does not affect depth perception when using advanced volume illumination techniques. However, external lighting may be carefully controlled in order to provide a pleasant environment while exploring the data (specular highlights on the screen, reflections or over-illuminated areas will certainly affect the correct perception of the data).
- When trying to judge depth in volume models, the X/Y relative position of the markers or the luminance of the points to classify seems to have no importance. No specific considerations have to be taken on this behalf.

7 Conclusions

This paper presents an experimental study to evaluate the influence of advanced volume illumination and the impact of real lighting to depth perception on stereoscopic desktop-based environments. In order to do this, we designed three tasks involving different well-known shading methods (no shading (i.e. emission-absorption model), Phong shading, half-angle slicing and directional occlusion shading) and lighting scenarios (with and without real lighting). Participants were asked to judge the relative depth between two points in stereo visualizations. Static images were chosen to avoid increasing the depth cues from e. g. motion, which would make to isolate the effect of shading quite difficult. Results showed that volume illumination improves depth perception in stereoscopic visualizations, being the simulation of soft shadows (direct occlusion shading) the shading model that provides better results. As well, we demonstrated that real-world lighting in a controlled environment does not affect depth perception when using advanced shading techniques. Furthermore the lack of coherence between real and virtual illumination neither affects depth perception when using the Phong shading model.

We believe that the presented study can be the starting point for a further evaluation of volume illumination techniques in VR environments. More concretely, we would like to develop a method that is able to accurately predict the goodness of a certain shading model or viewpoint, with respect to the proper communication of shape and depth of the scene. The initial experiments could be carried out in the line of searching for correlation between shading and depth maps, although this might be not the only approach. Future work may consider the evaluation of other global illumination effects like ambient occlusion or scattering. As well, future studies may check how luminance and color variations affect to depth perception. An interesting question that has not been covered in this study is the evaluation of individual light directions, instead of just considering the cases where virtual lighting is coherent or not with the real one. Previous research states that light direction affects the perception of the scene [34]. Is there any dominant direction which provides better depth cues than the others? Since volumetric models tend to be quite complex, cast shadows may highly vary from one angle to another. Therefore, the analysis seems not straightforward. We consider that further research in this direction could be useful to improve the exploration of volume data sets in VR environments.

Acknowledgements. The authors want to thank J.L. Díaz-Barrero and M. Fairén for their valuable contributions. As well, we thank the volunteers who took part in the study. This work was partially supported by the EU FP7 Program under the DIVA project (290277), the Eurostar CAMILIS project, the TIN2013-47137-C2-1-P and TIN2014-52211-C2-1-R projects of the Spanish Government.

References

1. Ayten, H., Herder, J., Vonolfen, W.: Visual acceptance evaluation of soft shadow algorithms for virtual tv studios. In: 13th International Conference on Humans and Computers, pp. 66–71 (2010)
2. Bach, C., Scapin, D.L.: Comparing inspections and user testing for the evaluation of virtual environments. International Journal of Human-Computer Interaction **26**(8), 786–824 (2010)
3. Baer, R., Adler, F., Lenz, D., Preim, B.: Preim b. perception-based evaluation of emphasis techniques used. In: in 3D Medical Visualization. In: Vision, Modeling, and Visualization Workshop, pp. 295–304 (2009)
4. Boucheny, C., Bonneau, G.P., Droulez, J., Thibault, G., Ploix, S.: A perceptive evaluation of volume rendering techniques. In: APGV: 4th Symp. on Applied perception in graphics and visualization, pp. 83–90 (2007)
5. Bowman, D.A., Gabbard, J.L., Hix, D.: A survey of usability evaluation in virtual environments: Classification and comparison of methods. Presence: Teleoper. Virtual Environ. **11**(4), 404–424 (2002)
6. Díaz, J., Ropinski, T., Navazo, I., Gobbetti, E., Vázquez, P.P.: Perceptual effects of volumetric shading models in desktop-based stereoscopic environments. In: Computer Graphics International. Strasbourg, France (2015)
7. Díaz, J., Vázquez, P.P., Navazo, I., Duguet, F.: Real-Time Ambient Occlusion and Halos with Summed Area Tables. Computers and Graphics **34**, 337–350 (2010)
8. Garau, M., Slater, M., Vinayagamoorthy, V., Brogni, A., Steed, A., Sasse, M.A.: The impact of avatar realism and eye gaze control on perceived quality of communication in a shared immersive virtual environment. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 529–536 (2003)
9. Germani, M., Mengoni, M., Peruzzini, M.: Metrics-based approach for VR technology evaluation in styling product design. In: International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, pp. 1325–1339. American Society of Mechanical Engineers (2009)
10. Grosset, A., Schott, M., Bonneau, G.P., Hansen, C.: Evaluation of depth of field for depth perception in dvr. In: Proceedings of the 2013 IEEE Pacific Visualization Symposium (PacificVis), pp. 81–88 (2013)
11. Hernell, F., Ljung, P., Ynnerman, A.: Local Ambient Occlusion in Direct Volume Rendering. IEEE Transactions on Visualization and Computer Graphics **15**(2), 548–559 (2009)
12. Ijsselsteijn, W., de Ridder, H., Hamberg, R., Bouwhuis, D., Freeman, J.: Perceived depth and the feeling of presence in 3dtv. Displays **18**(4), 207 – 214 (1998)
13. Jönsson, D., Sundén, E., Ynnerman, A., Ropinski, T.: A Survey of Volumetric Illumination Techniques for Interactive Volume Rendering. Computer Graphics Forum **33**(1), 27–51 (2014)
14. Kniss, J., Premoze, S., Hansen, C., Shirley, P., McPherson, A.: A Model for Volume Lighting and Modeling. IEEE Transactions on Visualization and Computer Graphics **9**(2), 150–162 (2003)
15. Laha, B., Bowman, D.A.: Identifying the benefits of immersion in virtual reality for volume data visualization. In: Immersive Visualization Revisited Workshop of IEEE VR conference (2012)
16. Laha, B., Bowman, D.A., Schiffbauer, J.D.: Validation of the mr simulation approach for evaluating the effects of immersion on visual analysis of volume data. Visualization and Computer Graphics, IEEE Transactions on **19**(4), 529–538 (2013)
17. Laha, B., Sensharma, K., Schiffbauer, J.D., Bowman, D.A.: Effects of immersion on visual analysis of volume data. Visualization and Computer Graphics, IEEE Transactions on **18**(4), 597–606 (2012)
18. Langer, M.S., Bülthoff, H.H.: Depth discrimination from shading under diffuse lighting. Perception **29**(6), 649–660 (2000)
19. Lee, C., Rincon, G.A., Meyer, G., Hollerer, T., Bowman, D.A.: The effects of visual realism on search tasks in mixed reality simulation. Visualization and Computer Graphics, IEEE Transactions on **19**, 547–556 (2013)
20. Levoy, M.: Display of Surfaces from Volume Data. IEEE Computer Graphics and Applications **8**, 29–37 (1988)
21. Lindemann, F., Ropinski, T.: About the Influence of Illumination Models on Image Comprehension in Direct Volume Rendering. IEEE Transactions on Visualization and Computer Graphics **17**(12), 1922–1931 (2011)
22. Liu, B., Todd, J.T., et al.: Perceptual biases in the interpretation of 3d shape from shading. Vision research **44**(18), 2135–2146 (2004)
23. Livatino, S., Koffel, C.: Handbook for evaluation studies in virtual reality. In: Virtual Environments, Human-Computer Interfaces and Measurement Systems, 2007. VECIMS 2007. IEEE Symposium on, pp. 1–6 (2007)
24. Max, N., Chen, M.: Local and Global Illumination in the Volume Rendering Integral. In: Scientific Visualization: Advanced Concepts, pp. 259–274 (2010)
25. McMahan, R.P., Bowman, D.A., Zielinski, D.J., Brady, R.B.: Evaluating display fidelity and interaction fidelity in a virtual reality game. Visualization and Computer Graphics, IEEE Transactions on **18**(4), 626–633 (2012)
26. Phong, B.T.: Illumination for computer generated pictures. Commun. ACM **18**(6), 311–317 (1975)
27. Pluim, J.P., Maintz, J.A., Viergever, M.A.: Image registration by maximization of combined mutual information and gradient information. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2000, pp. 452–461. Springer (2000)
28. Ropinski, T., Meyer-Spradow, J., Diepenbrock, S., Mensmann, J., Hinrichs, K.H.: Interactive Volume Rendering with Dynamic Ambient Occlusion and Color Bleeding. Computer Graphics Forum (Eurographics 2008) **27**(2), 567–576 (2008)
29. Ruiz, M., Boada, I., Viola, I., Bruckner, S., Feixas, M., Sbert, M.: Obscurance-based Volume Rendering Framework. In: Int. Symp. on Volume and Point-Based Graphics, pp. 113–120 (2008)
30. Schott, M., Pegoraro, V., Hansen, C., Boulanger, K., Bouatouch, K.: A directional occlusion shading model for interactive direct volume rendering. In: Proceedings of EuroVis'09, pp. 855–862 (2009)
31. Sielhorst, T., Bichlmeier, C., Heining, S., Navab, N.: Depth perception a major issue in medical ar: Evaluation study by twenty surgeons. In: Medical Image Computing and Computer-Assisted Intervention MICCAI 2006, *Lecture Notes in Computer Science*, vol. 4190, pp. 364–372 (2006)
32. Starck, J.L., Murtagh, F., Querre, P., Bonnarel, F.: Entropy and astronomical data analysis: Perspectives from multiresolution analysis. Astronomy & Astrophysics **368**(2), 730–746 (2001)
33. Stewart, A.J.: Vicinity Shading for Enhanced Perception of Volumetric Data. In: IEEE Visualization, pp. 355–362 (2003)
34. Sun, J., Perona, P.: Where is the sun? Nature Neuroscience (1), 183–184 (1998)
35. Sutcliffe, A., Gault, B.: Heuristic evaluation of virtual reality applications. Interacting with Computers **16**(4), 831 – 849 (2004)
36. Sutcliffe, A.G., Kaur, K.D.: Evaluating the usability of virtual reality user interfaces. Behaviour and Information Technology **19**(6), 415–426 (2000)
37. Thompson, W.B., Willemsen, P., Gooch, A.A., Creem-Regehr, S.H., Loomis, J.M., Beall, A.C.: Does the quality of the computer graphics matter when judging distances in visually immersive environments? Presence: Teleoperators and Virtual Environments **13**(5), 560–571 (2004)
38. Vázquez, P.P.: Automatic light source placement for maximum visual information recovery. In: Computer Graphics Forum, vol. 26, pp. 143–156. Wiley Online Library (2007)

39. Šoltészová, V., Patel, D., Viola, I.: Chromatic Shadows for Improved Perception. In: Proceedings of Non-Photorealistic Animation and Rendering (NPAR), pp. 105–115 (2011)
40. Šoltészová, V., Turkay, C., Price, M.C., Viola, I.: A Perceptual-Statistics Shading Model. IEEE Transactions on Visualization and Computer Graphics **18**(12), 2265–2274 (2012)
41. Wanger, L.C., Ferwerda, J.A., Greenberg, D.P.: Perceiving Spatial Relationships in Computer-Generated Images. IEEE Comput. Graph. Appl. **12**(3), 44–51, 54–58 (1992)