

## Technical Section

Evaluation of depth perception in crowded volumes<sup>☆</sup>Žiga Lesar<sup>\*</sup>, Ciril Bohak, Matija Marolt

University of Ljubljana, Faculty of Computer and Information Science, Večna pot 113, 1000 Ljubljana, Slovenia

## ARTICLE INFO

## Keywords:

Volume rendering  
Crowded volumes  
Depth perception  
User study

## ABSTRACT

Depth perception in volumetric visualization plays a crucial role in the understanding and interpretation of volumetric data. Numerous visualization techniques, many of which rely on physically based optical effects, promise to improve depth perception but often do so without considering camera movement or the content of the volume. As a result, the findings from previous studies may not be directly applicable to crowded volumes, where a large number of contained structures disrupts spatial perception. Crowded volumes therefore require special analysis and visualization tools with sparsification capabilities. Interactivity is an integral part of visualizing and exploring crowded volumes, but has received little attention in previous studies. To address this gap, we conducted a study to assess the impact of different rendering techniques on depth perception in crowded volumes, with a particular focus on the effects of camera movement. The results show that depth perception considering camera motion depends much more on the content of the volume than on the chosen visualization technique. Furthermore, we found that conventional non-photorealistic rendering techniques, which have often performed poorly in previous studies, showed comparable performance to modern photorealistic techniques in our study. The source code for the visualization system, survey, and analysis, as well as the data set used in the study and the participants' responses, have been made publicly available.

## 1. Introduction

Volumetric data are ubiquitous in various scientific fields. They often originate from imaging techniques such as X-ray scans, computed tomography and magnetic resonance imaging, or from simulations in different domains. Effective visualization of such data is paramount for understanding its characteristics. Typically, this is achieved through volume rendering, where a transfer function is used to map raw volume samples to optical properties such as color and opacity. Depending on the chosen rendering technique, the rendering process may additionally account for transparency, illumination, occlusion, and projection to produce a 2D image. This is done to ensure adequate spatial perception of the viewer since mimicking or simulating realistic optical effects has been shown to enhance this aspect [1–3]. In essence, the primary objective of scientific volume visualization is to convey spatial information, including depth, distance, size, shape, and orientation as clearly as possible while providing users with the flexibility to interactively adjust certain visualization parameters, such as the transfer function and illumination setup.

Many visualization techniques promise to enhance spatial perception, a benefit that has undergone rigorous testing and validation in numerous studies. To ensure precision and repeatability, experiments in these studies are often conducted within highly controlled environments, with researchers making extensive efforts to isolate variables.

As a result, such experiments typically involve only static images and a limited set of independent variables. Conventional volume rendering techniques with no global illumination features, such as the emission–absorption model and isosurface raycasting, tend to underperform in these constrained experiments due to their simplicity and lack of depth-enhancing features. The emission–absorption model [4] is an illustrative example of such a technique. While it is a rather simple technique, its widespread adoption in diverse domains attests to its ability to handle various visualization tasks.

Furthermore, previous experiments have focused on volumes representing recognizable structures, such as CT scans of the human body or angiographic images of blood vessels. These images predominantly feature familiar elements and can be rendered less visually dense by an appropriate transfer function. Both familiarity and visual density have been shown to profoundly affect spatial perception [5,6]. This presents a major challenge when visualizing *crowded volumes*, which are characterized by an overwhelming amount of mutually occluding structures. A significant difference between crowded and non-crowded volumes is that structures in non-crowded volumes tend to be larger and more interconnected throughout the volume, whereas structures in crowded volumes are typically smaller and scattered across the volume. Existing volume rendering techniques prove to be highly ineffective in a crowded scenario due to the inherently complex structure of crowded

<sup>☆</sup> This article was recommended for publication by Stefan Bruckner.

<sup>\*</sup> Corresponding author.

E-mail addresses: [ziga.lesar@fri.uni-lj.si](mailto:ziga.lesar@fri.uni-lj.si) (Ž. Lesar), [ciril.bohak@fri.uni-lj.si](mailto:ciril.bohak@fri.uni-lj.si) (C. Bohak), [matija.marolt@fri.uni-lj.si](mailto:matija.marolt@fri.uni-lj.si) (M. Marolt).

volumes, making it imperative to use specialized tools [7–9]. In this case, interactivity plays a much more important role, as users benefit from interactive grouping, colorization, and sparsification of the contained structures to lower their visual density. On the other hand, the substantial amount of occlusion observed in crowded volumes also has the potential to enhance spatial perception. According to Englund and Ropinski [10], occlusion stands out as one of the strongest monocular depth cues, and the parallax effect induced by the camera movement can further amplify its significance. Therefore, when evaluating spatial perception in crowded volumes, it is essential to incorporate camera motion as an integral component of the visualization.

Previous studies have recognized the importance of camera motion in improving spatial perception. However, its evidently dominant positive effect has often led to its exclusion from evaluations. As a result, rendering techniques with no global illumination features have often underperformed due to the lack of essential depth cues. Furthermore, while studies have shown that the content of a volume has a significant impact on spatial perception, they have not considered crowded volumes or the degree of crowdedness as variables. In light of this, we conducted a study on depth perception in crowded volumes, in which we included camera motion to reflect a typical use case scenario. We employed two photorealistic rendering techniques in addition to two non-photorealistic techniques for visualization. We visualized three crowded volumes from different domains exhibiting varying degrees of crowdedness. Our goal was to find out how crowdedness and occlusion affect depth perception in crowded volumes in the presence of camera motion. We also wanted to evaluate the performance of different rendering techniques and lighting setups under such conditions. Based on previous research and the above arguments, we developed the following hypotheses:

- **(H1)** Crowdedness has a negative impact on depth estimation accuracy in crowded volumes.
- **(H2)** Depth estimation accuracy in crowded volumes increases with the size of the instances.
- **(H3)** Camera motion plays a more crucial role in depth estimation accuracy than the rendering technique and lighting setup.

By understanding how depth estimation accuracy varies with respect to the size, shape, and crowdedness of structures within a crowded volume, we are making a step toward understanding the perception of a volume as a whole. As a main contribution, our study highlights the importance of camera motion and volume content in the perception of crowded volumes, while demonstrating that the rendering technique plays a lesser role. The results of the study will help to identify suitable visualization approaches for visual inspection and analysis of crowded volumes.

The rest of the paper is organized as follows. Section 2 lists the most important works and outcomes in volume visualization and perception. Section 3 describes our experimental setup, the volumes and the techniques used. We report the results in Section 4 and discuss their implications in Section 5. Finally, we present the conclusions in Section 6.

## 2. Related work

We divide the related work into two subsections, covering visualization techniques and perception studies.

### 2.1. Visualization techniques

As per Viola et al. [11], visualization techniques can be broadly divided into photorealistic and non-photorealistic techniques.

Non-photorealistic techniques include early optical models such as isosurface rendering [12] and the emission-absorption model [4]. Despite their age and limitations, both approaches are still well established in the visualization community. The non-photorealistic category

of visualization techniques includes illustrative techniques [11] and various smart visualization techniques [13]. Bruckner and Gröller [14] developed volumetric halos based on established medical illustration techniques, where partially occluded structures are darkened to create the illusion of depth. This idea has been later extended by Schott et al. [15] into a physically-based technique called directional occlusion shading (DOS). Several studies ranked DOS as one of the best techniques for enhancing depth perception [1,3,16]. In addition to illumination, researchers have explored leveraging other visual phenomena, such as depth of field [5], aerial perspective [17], and chromadepth [10,18–21], to improve depth perception. The motivation behind such a wide array of techniques that rely on chromadepth or pseudo-chromadepth is the observation that visualization techniques generally perform best when tasks can be reduced to simple comparisons [21]. In this manner, depth-related tasks are simplified to straightforward hue, brightness, or saturation comparisons. For a comprehensive overview of rendering techniques that are based on a close understanding of the human visual system to improve the images (so-called perception-driven techniques), we refer the reader to a survey by Weier et al. [22].

The photorealistic category includes techniques that aim to simulate light behavior in order to improve spatial perception, the reason being that the human visual system has evolved to interpret realistic illumination phenomena resulting from the underlying physics [3]. This category includes physically-based techniques, such as path tracing [23], diffusion simulation [24], combined volume and surface rendering [25–27], as well as approximations of global illumination, such as spherical harmonic lighting [28] and half-angle slicing [29]. A comprehensive survey of illumination techniques in volume rendering is provided by Jönsson et al. [2].

Both photorealistic and non-photorealistic techniques benefit from visibility optimization techniques. This is especially true for crowded volumes. Chan et al. [30] presented transparency optimization to improve visibility of certain parts of a volume. Correa and Ma [31] introduced a similar idea based on visibility histograms. LeMuzic et al. [8] extended the technique to crowded mesoscopic biological models by utilizing automatic sparsification. Lesar et al. [9] adapted the visibility management concepts for the visualization of crowded volumes.

### 2.2. Perception studies

Perception in volume rendering has been the subject of extensive research. This is mainly due to the fact that perception is a complex process and is influenced by a wide range of factors, which are often hard to measure. Moreover, the importance of factors depends on the task at hand [1,21]. For example, no single visualization technique is optimal for both shape and depth perception [3].

Human vision relies on various depth cues to deduce depth information from a scene. These can be categorized as either binocular, which are based on information from both eyes, and monocular, which can be observed with just one eye [32]. Computer graphics literature primarily focuses on monocular cues since binocular cues cannot be utilized when viewing images on a single 2D screen. Depth cues vary in their impact on depth perception, and they can be combined to enhance the effect [21,33]. For example, many visualization techniques rely on shadows and shading, and multiple studies have demonstrated their crucial role in depth perception [34]. Interestingly, motion parallax and depth from motion are typically excluded from studies due to their substantial influence on depth perception, leaving only a handful of studies that include motion [21,33,35,36]. Díaz et al. [16] even state that it is disruptive. Some studies acknowledge the significance of motion as an important depth cue and suggest that further research is warranted in this area [5,6].

Boucheny et al. [35] noted that, ideally, one would see at a glance the whole volume and have a clear image in mind of the spatial layout. However, they also acknowledged that such representations of data are

not natural because the human visual system usually experiences only solid surfaces. In their study, they demonstrate a significant preference for interpreting size as the main indicator of depth, even in the presence of conflicting information from other depth cues. Additionally, the strength of perspective has a notably positive impact on depth perception. Their study, however, had limitations due to the restricted range of rendering techniques and the use of very simplistic scenes.

Ropinski et al. [37] compared their shadow volume propagation technique to gradient-based shading and found significant improvements in speed and accuracy of depth perception. Lindemann and Ropinski [1] extended these results by conducting a larger user study in which they compared 7 state-of-the-art illumination models. They found that DOS [15] significantly outperformed other illumination models in ordinal depth perception and relative size perception tasks and performed very well in absolute depth perception tasks. In contrast to the previous study by Ropinski et al. [37], participants expressed a preference for gradient-based shading over more advanced illumination models. The authors attribute this outcome to brighter and more colorful images produced by gradient-based shading. Interestingly, the study suggests that shadows may negatively affect size and depth perception.

Englund and Ropinski [10] assessed the perception of images produced with non-photorealistic volume rendering techniques and discovered that volumetric halos and depth-darkening techniques excel in absolute depth perception tasks. These techniques were also favored by the users. Ordinal depth perception, on the other hand, requires qualitative techniques, such as pseudo-chromadepth. The study was carried out using a crowdsourcing solution [38], which was shown to produce evaluation results comparable to those obtained with controlled experiments. The authors later expanded the study with feedback from three experts who stressed the significance of motion, clutter avoidance, and reduced transparency [6].

Grosset et al. [5] investigated the influence of depth-of-field on depth perception. Although they did not manage to validate any of their hypotheses, their study highlights substantial variations in depth perception across datasets. The research also indicates that familiarity has a notably positive effect on perception. Furthermore, anchoring proves to be beneficial, as objects suspended in mid-air pose challenges for spatial perception.

A great deal of research has already been carried out in the field of virtual reality (VR) and augmented reality (AR). Diaz et al. [39] evaluated depth perception in AR with different settings for shadows, shading, perspective, dimensionality and texture. Their study showed that users' performance improved significantly with the use of shadows. In addition, their study showed that depth is consistently underestimated in AR. Díaz et al. [16] measured ordinal depth perception in stereoscopic displays of static images. They compared four techniques, with DOS performing the best and the emission-absorption model performing the worst. Interestingly, their study revealed no correlation between user accuracy and the on-screen distance between the estimation points. Hu et al. [33] investigated ordinal depth perception in VR on a neurological use case. They recognized the importance of shadows and motion parallax. However, their study found no significant differences in user performance when either of these depth cues was used. Heinrich et al. [40] conducted a study in a projective AR environment in which pseudo-chromadepth and support lines provided the best results.

In the field of medical visualization, considerable attention has been devoted to angiographic images. The study by Ropinski et al. [18] compared several visualization methods and revealed that angiography benefits from color coding depth information. This result was subsequently validated by Kersten-Oertel et al. [21], who found that pseudo-chromadepth and aerial perspective resulted in the best user performance and were also favored by the users. Lesar et al. [41] assessed user preferences for various angiography visualization methods and found a strong inclination toward isosurface raycasting with Phong shading and ambient occlusion. Drouin et al. [42] explored dynamic

visualization techniques, in which the visualization parameters are modified based on the motion of the VR controllers, applied in a surgical scenario. Their results unveiled significantly improved accuracy, albeit at a slower pace compared to static methods. Titov et al. [36] conducted a similar study in VR and found that the advantages of binocular disparity and motion parallax outweighed the benefits of dynamic techniques. Kreiser et al. [43] explored the use of empty space surrounding vessels to convey depth information and observed substantial improvements in ordinal depth perception.

In summary, although there are numerous volume rendering techniques that enhance depth perception, selecting a specific technique largely depends on the task at hand. Motion has received little attention in previous studies, and to the best of our knowledge, no study has yet explored depth perception in crowded volumes.

### 3. Experiment setup

The aim of the study was to investigate the effect of crowdedness, occlusion and camera motion on depth perception in crowded volumes using different rendering techniques and lighting setups. The study was conducted in the form of an online crowdsourced survey whose target group comprises the general population. We did not focus exclusively on people who have experience with volume rendering, visualization or 3D modeling. Participants were asked to estimate the relative depth between structures after being shown a short video of a 3D volume. Since the study was conducted in the form of an online crowdsourced survey, the experimental environment was largely uncontrolled. The only condition we enforced was the use of a desktop device under the assumption that such a device would have a sufficiently large screen. The experiment consisted of 128 test cases in which four volumes were rendered using four different rendering techniques and lighting setups and shown from eight different perspectives. The videos used in the test cases have been made publicly available [44]. The source code for the visualization system and the survey is available on GitHub.<sup>1</sup>

#### 3.1. Evaluation interface

Each evaluation session began by showing participants background information about the study and the crowded volumes. Participants were asked to enter basic demographic data: gender, age, occupation, and experience with 3D environments (VR, gaming, and 3D modeling). They were then given instructions on how to complete each task. Specifically, for each task, they were shown a 5-s video of a volume in which the camera performed a subtle circular motion while keeping the volume in focus (see Section 3.4 for details). When the video stopped playing, three points lying on the structures in the volume were shown, and participants were instructed to estimate the relative depth of the *estimation point* between the two *bounding points* using a slider (see Fig. 1). The points were marked with color-coded circles: the point closest to the camera was marked with a green circle, the point furthest away from the camera with a purple circle and the estimation point with a gray circle. The colors were chosen to be complementary and approximately equally bright, but at the same time as neutral as possible to avoid conveying unintended information to the participants. In particular, we tried to avoid the combination of green and red, which could be misinterpreted as “right” and “wrong”. Participants were allowed to replay a video once to get a better sense of the spatial relationships between the structures in the volume. During video playback, the three points were hidden to avoid shifting the participants' focus toward the parallax effect rather than the actual rendering, as it is possible to estimate the relative depth of the three points based solely on their apparent movement due to the parallax effect.

The positions of the points were manually chosen to adhere as closely as possible to the following criteria:

<sup>1</sup> <https://github.com/UL-FRI-LGM/depth-perception-evaluation>



Fig. 1. A screenshot of the user interface during the evaluation. The slider below the rendering was used to estimate the depth of the gray estimation point between the green and purple bounding points.

- the relative depth of the estimation point should be approximately uniformly distributed,
- the depths of the bounding points should be approximately uniformly distributed inside the bounds of the volume,
- the points should be approximately uniformly distributed across the display, and
- the points should not be occluded by structures in the volume when the video stops.

The statistics for the above metrics across all test cases are shown in Fig. 2. The distances involved in the above metrics are sketched in Fig. 3.

### 3.2. Volumes

We included four volumes in the study, three of which were crowded and the remaining one served as a control volume. The volumes are shown in Fig. 4. They come from different domains and their crowdedness varies considerably.

The first volume is from material science [7] and represents an X-ray computed tomography image of a fiber-reinforced polymer that belongs to a family of materials that are in high demand in the automotive and aerospace industries due to their strength, durability, elasticity and light weight. Experts working with such materials analyze their properties by examining the length, orientation and distribution of the fibers. Due to the high fiber density, the tasks associated with the analysis of fiber-reinforced polymer materials require the use of special tools such as Fiber Scout [7] or Volume Conductor [9]. The volume we used for the user study measures  $400 \times 401 \times 800$  voxels and contains 3828 instances of glass fibers. We will refer to this volume as *fibers*.

The second volume comes from the biological domain. It is a  $1280 \times 1024 \times 1024$  focused ion beam scanning electron microscopy (FIB-SEM) image capturing a segment of a mouse's bladder cell. The image was segmented to highlight the different intracellular organelles [45,46]. In addition to the numerous fusiform vesicles, Golgi apparatus, endosomes and lysosomes, there are 3051 mitochondria in the volume.




Fig. 2. Distribution of distances and depths between evaluation points among test cases: relative depth of the estimation point, on-screen (2D) distance between the bounding points, volumetric (3D) distance between the bounding points, and depth difference between the bounding points. The on-screen distance has been computed in normalized device coordinates, and the volumetric distance has been computed such that the volume was a cube with a unit side length.




Fig. 3. The three distances between bounding points: on-screen (2D) distance between the bounding points (A), volumetric (3D) distance between the bounding points (B), and depth difference between the bounding points (C).

Analyzing their distribution is of crucial importance for understanding the biological processes within a cell. We will refer to this volume as *mitos*.

The third volume is synthetic and was created to contain a large number of instances. The volume of size  $512 \times 512 \times 512$  contains 22670 instances of small convex polytopes resembling pebbles. The instances were physically simulated to achieve the densest possible packing within the contained volume. We will refer to this volume as *pebbles*.

The fourth volume is a non-crowded control volume that was included in the study to allow comparison with previous studies. The volume is a  $512 \times 512 \times 460$  CT image of the upper torso of a human.<sup>2</sup> We will refer to this volume as *manix*.

<sup>2</sup> Manix dataset from <http://www.osirix-viewer.com/datasets/>



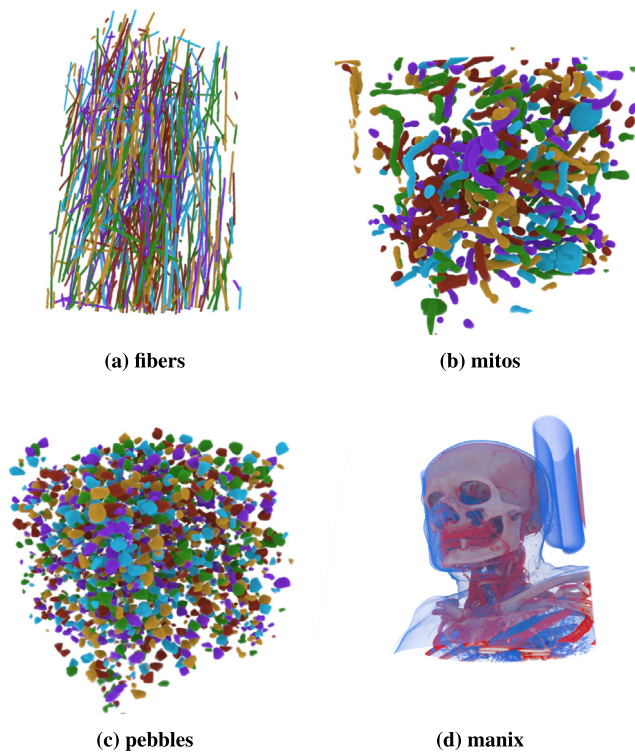


Fig. 4. The test volumes used in the experiment: fibers, mitos, pebbles, and manix. Instances in the crowded volumes have been divided into 5 groups and colored. The transfer function for manix has been constructed manually. The images were rendered with volumetric path tracing.

To reduce crowdedness to a level where the interior of the volume is at least partially visible, the pebbles were uniformly sparsified to 10% density, and the fibers were uniformly sparsified to 20% density. To represent a typical usage scenario, where the instances are grouped according to some criterion, the instances in all crowded volumes were divided into 5 groups of approximately equal size, and the colors have been assigned to the groups according to the golden ratio sequence [47], which guarantees an approximately uniform distribution of hues. Since the manix volume is not crowded and could not be assigned colors automatically, the transfer function has been constructed manually so that important structures inside the volume were clearly visible.

### 3.3. Rendering techniques and lighting setups

For the visualization we used 4 widely used volume rendering techniques, namely the Emission-Absorption Model (EAM) [4], Directional Occlusion Shading (DOS) [15], Isosurface Raycasting (ISO) [12] and Volumetric Path Tracing (VPT) [23], each with a specific lighting setup. Fig. 5 shows a visual comparison of the four rendering techniques. In all renderings, we used a white illuminant. The isosurfaces were shaded using the Disney’s physically-based model [48] with a directional light source and a fully rough material. For the directional occlusion model, an illumination cone with an apex angle of  $60^\circ$  was used. For volumetric path tracing, the multiple scattering model was used with a slightly backscattering medium modeled with a Henyey-Greenstein phase function [49] with the anisotropy parameter of  $-0.5$ . The volume was illuminated with an infinite environmental isotropic white light.

These techniques were implemented on top of the platform-independent volume rendering framework VPT [50] and extended by the Volume Conductor interface [9] with the capabilities to render and




Fig. 5. The rendering techniques and lighting setups used in the experiment: emission-absorption model (EAM) without shading, directional occlusion shading (DOS) with a  $60^\circ$  illumination cone, isosurface raycasting (ISO) with local surface shading and a directional light source, and volumetric path tracing (VPT) with a Henyey-Greenstein phase function with anisotropy of  $-0.5$  and an infinite environmental isotropic light source.

sparsify crowded volumes. No additional depth-enhancing techniques, such as aerial perspective, pseudo-chromadepth, or volumetric halos, were used to improve the renderings. This decision was made to limit the number of independent variables examined in the study.

### 3.4. Views

To ensure comparability between images produced with different rendering techniques and lighting setups, we generated 8 camera views that captured the data as uniformly as possible from different angles. We used spherical Fibonacci mapping [51] to generate the camera positions on the bounding sphere surrounding the volume and aimed the camera at the center of the volume so that the up vectors were aligned in all configurations. The vertical field of view of the camera was set to  $90^\circ$ .

The camera performed a small circular motion around the base of a cone with the apex at the center of the volume and an apex angle of approximately  $11^\circ$ . The apex angle was chosen so that the parallax effect was sufficiently clear but not excessive.

## 4. Results

The experiment was online for 46 days. During that time, we received 4340 depth estimates, which have been made publicly available [44]. The estimates were provided by 47 participants, 29 of whom completed the entire set of 128 depth estimation trials. 33 participants reported demographic data. 21 were women, and 12 were male. Their ages ranged from 14 to 74, with a mean age of around 38 (see Fig. 6). The participants reported their prior experience with 3D environments (games, 3D modeling, and VR) on a 5-point Likert scale ranging from 1 (no experience) to 5 (a lot of experience). The average gaming experience reported by the participants was 2.64, with a standard deviation of 1.41. 3D modeling experience was reported with a mean




Fig. 6. Distribution of ages among 33 out of 47 study participants who reported demographic data.




Fig. 7. Distribution of the participants' prior experience with gaming, 3D modeling, and VR on a 5-point Likert scale ranging from 1 (no experience) to 5 (a lot of experience).

of 2.30 and a standard deviation of 1.38. VR experience had a mean of 2.06 and a standard deviation of 1.09. Histograms illustrating the distributions of these experiences are shown in Fig. 7.

To evaluate depth perception, we measured the difference between the correct and estimated relative depths. The estimated depth is given by the participants using a slider (see Fig. 1), while the correct relative depth is calculated using a depth buffer associated with the rendering. Both numbers fall within the range of 0 to 1 due to the nature of their acquisition. To test our hypotheses, we relied on the mean absolute difference (MAD) between the correct and estimated depths. In theory, if the correct and estimated depths are independent and uniformly distributed on the unit interval, then the expected MAD is 1/3. We use this number as a reference for interpreting the measured data.

Since the experiment was uncontrolled, the participants could pause the evaluation or refresh the page at any time during the experiment. As a result, we could not reliably measure the response time, which is often analyzed as an important performance metric in controlled experiments.

First, we measured MAD by grouping the trials by rendering technique and volume. When grouping the trials by rendering technique (see Fig. 8), the ISO rendering technique exhibited the lowest MAD (MAD = 0.175, SD = 0.153), followed by DOS (MAD = 0.180, SD = 0.152), VPT (MAD = 0.181, SD = 0.156), and EAM (MAD = 0.186, SD = 0.151). We tested the null hypothesis that the MAD was equal between different techniques. The normality and homoscedasticity assumptions were tested with Bartlett's and Shapiro-Wilk's tests, respectively. The normality assumption was unmet; therefore, we used the Kruskal-Wallis test. The test showed no significant differences in MAD between the techniques ( $H(3) = 0.85406, p = 0.8365$ ). Furthermore, we found no statistically significant bias of the mean depth difference.

When grouping the trials by volume (see Fig. 9), the lowest MAD was achieved on the mitos volume (MAD = 0.154, SD = 0.135), followed by fibers (MAD = 0.164, SD = 0.151), manix (MAD = 0.198, SD = 0.163), and pebbles (MAD = 0.210, SD = 0.157). We tested the




Fig. 8. Mean absolute difference (MAD) with respect to the rendering technique and lighting setup.




Fig. 9. Mean absolute difference (MAD) with respect to the test volume.

null hypothesis that the MAD was equal between different volumes. The normality and homoscedasticity assumptions were tested with Bartlett's and Shapiro-Wilk's tests, respectively, and have not been met. Therefore, we used the Kruskal-Wallis test. The test showed a significant difference in MAD between the volumes ( $H(3) = 55.262, p < 0.001$ ). Furthermore, Dunn's post-hoc test with Holm's adjustment showed significant differences between all pairs of volumes ( $p < 0.05$ ) except between the manix volume and the pebbles volume. Furthermore, we found no statistically significant bias of the mean depth difference.

In order to address the potentially overbearing influence of test volumes when testing for differences in MAD between rendering techniques, we additionally ran the Kruskal-Wallis test for each volume




Fig. 10. Mean absolute difference (MAD) with respect to the participants' prior experience with gaming, 3D modeling, and VR expressed on a 5-point Likert scale ranging from 1 (no experience) to 5 (a lot of experience).

separately. Once more, our analysis revealed no statistically significant differences in MAD between the evaluated techniques.

Next, we measured MAD by grouping the trials by the participants' prior experience with 3D environments (based on the participants who provided these data). A boxplot of the data is shown in Fig. 10. Since prior experience is a discrete variable expressed on a Likert scale, we used the Kruskal–Wallis test on a null hypothesis, stating that the MAD was equal between different experience levels. The test showed significant differences in all three experience categories: gaming ( $H(4) = 64.84, p < 0.001$ ), 3D modeling ( $H(4) = 59.12, p < 0.001$ ), and VR ( $H(4) = 65.02, p < 0.001$ ). We used Dunn's post-hoc test with Holm's adjustment to find the experience levels that exhibited significant differences in MAD. In all three experience categories, the test showed that the participants who reported substantial experience (level 5 on the Likert scale) performed significantly better ( $p < 0.001$ ) than those who reported less experience (levels 1–4 on the Likert scale). We found similar significant differences between experience levels 2 and 3, but we believe that this result might be skewed by psychological factors other than experience (see Section 5). Furthermore, we found weak negative correlations between MAD and experience levels in the VR category ( $r = -0.01, F(1, 2957) = 16.77, p < 0.001$ ) and the 3D modeling category ( $r = -0.006, F(1, 2957) = 7.30, p < 0.01$ ). We found no statistically significant correlation between MAD and age of the participant.

Additionally, we found statistically significant correlations between MAD and the positions of the bounding points. We found a positive correlation between MAD and 2D distance between bounding points on the screen ( $r = 0.14, F(1, 4338) = 43.52, p < 0.001$ ), a slight negative correlation between MAD and 3D distance between bounding points in the volume ( $r = -0.07, F(1, 4338) = 29.65, p < 0.001$ ), and a negative correlation between MAD and depth difference between bounding points in the volume ( $r = -0.26, F(1, 4338) = 51.75, p < 0.001$ ).

## 5. Discussion

The results of our study highlight several important aspects of participants' performance in a depth estimation task within crowded volumes. All participants achieved a MAD lower than the theoretical 1/3 derived from uniformly random distance sampling and estimation. In general, MAD ranged consistently between 0.15 and 0.25 (see Fig. 11).

Our analysis revealed no significant differences in MAD when grouping across different rendering techniques and lighting setups (Fig. 8). Conventional non-photorealistic techniques (EAM and ISO)

yielded a MAD similar to that of modern photorealistic techniques (DOS and VPT), which supports hypothesis (H3). This outcome aligns with the claims of several papers, which state that monocular cues resulting from camera movement, such as motion parallax, tend to outweigh the cues from shading and illumination [6,10,35]. In practical scenarios, where users typically control camera movement through direct interaction, we expect this effect to be even more pronounced, given the close connection between the human visual and motor systems [42,52]. We believe that the enduring popularity of conventional non-photorealistic volume rendering techniques, which has persisted for decades since their introduction, may be partly attributed to the influence of interaction.

We found significant differences in MAD when we grouped across different volumes (Fig. 9), suggesting that the content of a volume has a greater impact on depth perception than the rendering technique or lighting setup used for visualization. Interestingly, the manix volume, which does not exhibit crowdedness, ranked third among the evaluated volumes, with only the pebbles volume achieving a worse MAD. Familiarity with the shapes of human body parts does not appear to have contributed to a better score. The lower ranking may also be due to the transparency causing confusion about which structures were highlighted by the on-screen markers. Results from grouping across individual test cases (Fig. 12) show that the manix volume ranked among the best and worst in terms of MAD, depending on the camera view. This suggests that certain camera views might introduce bias due to familiarity with the internal structures. In contrast, crowded volumes showed more consistent MAD values between camera views, indicating the absence of familiar features within these volumes.

The two top-ranked volumes were mitos and fibers, which is consistent with hypotheses (H1) and (H2). We attribute this ranking to the larger size of the structures they contain, leading to increased occlusion between them. Notably, the mitos volume achieved a lower MAD than the fibers volume, possibly due to the larger surface area of the structures, which made them more easily visible and identifiable within the crowded volume. This provides compelling evidence in support of hypothesis (H2). Conversely, the pebbles volume attained the highest MAD due to its substantial crowdedness and the small average size of the structures, which made it considerably more challenging to identify and visually track structures during camera movement. Participants informally confirmed this difficulty during post-experiment discussions. Indeed, when a structure is small, its visibility essentially becomes binary, eliminating partial occlusion as one of the strongest monocular depth cues [10] from the visualization. Furthermore, small structures do not significantly contribute to the overall illumination of the scene. They often produce small and soft shadows that have been shown to be confusing and detrimental to depth perception [1]. In accordance with hypothesis (H1), this can be overcome by strong sparsification to reduce the number of visible instances and by introducing specific anchoring objects to create harder shadows. From the data we have available, we cannot determine whether it is the small size of the structures or their crowdedness that contributes most to the high MAD. This aspect remains open for future exploration.

Our analysis revealed significant correlations between MAD and reported experience with 3D environments. Specifically, participants with the highest reported experience levels consistently achieved higher accuracy in their depth estimates. This might stem from familiarity with 3D environments in general or with common biases inherent in 3D rendering, such as size bias [35], depth bias [32,39], or shading bias [53,54]. Setting aside the highest experience levels, the slight negative correlation between MAD and experience level aligns with our expectation that prior experience with 3D environments positively influences depth perception. This is most evident when examining the performance of participants who reported no prior experience, as they consistently achieved worse MAD than more experienced participants. Additionally, we observed a significant difference in MAD between groups of participants that reported experience levels 2 and 3, which is




Fig. 11. Mean absolute difference (MAD) for each participant sorted in ascending order. Note that not all participants evaluated all test cases.




Fig. 12. Mean absolute difference (MAD) for each test case sorted in ascending order. Test cases involving the manix volume are among the best and worst in terms of MAD. Test cases involving the fibers volume exhibit high variance of the estimate errors.

likely related to the larger number of participants in the latter group. It is possible that participants with uncertain knowledge about their prior experience simply selected the middle option (level 3) instead of indicating no experience (level 1). Furthermore, our findings indicate that the age of the participants had no significant impact on performance. We had expected that older participants would either perform better due to their potentially greater experience with 3D images, or worse due to slower and less precise interaction with the computer and potential vision issues. The lack of age-related effects could be attributed to the absence of time constraints during the experiment that could have allowed for longer response times, or the artificial nature of the task itself, which would nullify any additional experience.

The analysis of correlations between MAD and the positions of the bounding points yielded the expected results. Participants achieved better performance when the bounding points were closely positioned on the screen. In contrast, focusing on two distant areas of the screen simultaneously means that the participants probably had to alternate their focus, which is difficult at best. During post-experiment discussions, participants noted that they often lost track of the bounding points during the initial animation, which made depth estimation considerably more difficult. The ability to replay the video slightly alleviated the issue. Furthermore, participants performed better when the depths of the bounding points differed substantially. This could result from a more pronounced motion parallax effect. Additionally, we believe that given a larger depth difference between the bounding points, participants could have fine-tuned their answers since they could leverage the contextual information provided by a larger neighborhood of structures surrounding the bounding points and the estimation point.

The study also has certain limitations, primarily due to the lack of control over the experiment. To ensure comparability among the participants, there was no direct interaction with the visualization involved in the experiment. Having direct control over the camera view, sparsification, grouping, and colorization of the structures, as well as the rendering technique and its parameters, would undoubtedly improve depth estimation accuracy. From the available data we also cannot assess the effect of rendering techniques and lighting setups on depth perception in dynamic scenarios without carefully isolating these variables, although they clearly have a profound effect on depth perception in static images. Furthermore, a major result of our study highlights the effect of the size, shape, and crowdedness of structures in a crowded volume. However, since we only had a limited number of test volumes available, a follow-up study is warranted to investigate the effect of these variables in more detail, perhaps using a synthetic data set.

## 6. Conclusion

Depth perception is a challenging trait to measure, as it depends on the user’s prior experience, environment, lighting setup, interaction, and the data itself. Every experiment aiming to measure depth perception must, therefore, clearly and precisely define all the above parameters. In contrast to previous studies, our study emphasizes the significance of camera movement as one of the most important features of visualization. While previous studies identified significant differences in user performance between different rendering methods with static images, our results demonstrate that even non-photorealistic rendering techniques without global illumination features can perform



as well as advanced methods when camera movement is included. This provides strong evidence supporting hypothesis (H3).

Additionally, we showed that the volume's content affects depth perception more than the rendering method, which is in line with previous research [1,3,5]. Most importantly, we used crowded volumes where we expected no familiarity with the structures contained in the volumes. The degree of crowdedness appears to negatively impact user performance due to the excessive amount of detail in the images, which is consistent with previous research [1,3] and hypothesis (H1). We showed that instances with a large screen surface area exhibit complex mutual occlusion, which positively affects depth perception since in this case the visibility of a single instance is not binary. This outcome is in line with hypothesis (H2). The worst-case scenario seems to be an extremely crowded volume with small instances, such as the pebbles volume in our study, where user performance was considerably worse than with other volumes.

Furthermore, we discovered correlations between user performance and the on-screen distance between the bounding points, and also between the depth difference between the bounding points. This finding underscores the importance of direct interaction, especially when the user can roughly align the structures of interest with the viewing direction. Such alignment minimizes the on-screen distance and maximizes the depth difference, allowing for better utilization of spatial context (the neighborhood of the structures of interest) and motion parallax.

Since our study was conducted on the general population, prior experience with 3D environments was found to be a contributing factor to the performance in the experiment. As expected, we observed a positive correlation between accuracy of depth estimation and the reported amount of prior experience.

The main limitation of our study is, of course, the lack of direct interaction between the participant and the visualization. This includes adjusting the camera view, the degree of sparsification, and the grouping and coloring of instances in the crowded environment. The reason for not including these aspects in the study is that the results would not be comparable between participants, as each participant interacts with the visualization in their own way. This opens up many possibilities for future research, especially in the context of direct interaction. Furthermore, having uncovered the significance of the size of structures and the degree of crowdedness, a closer examination of these variables is warranted in a follow-up study.

#### CRediT authorship contribution statement

**Žiga Lesar:** Conceptualization, Data curation, Methodology, Software, Visualization, Writing – original draft, Writing – review & editing. **Ciril Bohak:** Writing – review & editing. **Matija Marolt:** Supervision, Writing – review & editing.

#### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Data availability

The data has been made publicly available at [44].

#### Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.cag.2024.103918>.

#### References

- [1] Lindemann F, Ropinski T. About the influence of illumination models on image comprehension in direct volume rendering. *IEEE Trans Vis Comput Graphics* 2011;17:1922–31. <http://dx.doi.org/10.1109/TVCG.2011.161>, URL: <http://ieeexplore.ieee.org/document/6064955/>.
- [2] Jönsson D, Sundén E, Ynnerman A, Ropinski T. A survey of volumetric illumination techniques for interactive volume rendering. *Comput Graph Forum* 2014;33:27–51. <http://dx.doi.org/10.1111/cgf.12252>, URL: <http://doi.wiley.com/10.1111/cgf.12252>.
- [3] Preim B, Baer A, Cunningham D, Isenberg T, Ropinski T. A survey of perceptually motivated 3D visualization of medical image data. *Comput Graph Forum* 2016;35:501–25. <http://dx.doi.org/10.1111/cgf.12927>, URL: <http://doi.wiley.com/10.1111/cgf.12927>.
- [4] Max N. Optical models for direct volume rendering. *IEEE Trans Vis Comput Graphics* 1995;1:99–108. <http://dx.doi.org/10.1109/2945.468400>, URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=468400>.
- [5] Grosset AVP, Schott M, Bonneau G-P, Hansen C. Evaluation of depth of field for depth perception in DVR. In: 2013 IEEE Pacific visualization symposium (Pacificvis). IEEE; 2013, p. 81–8. <http://dx.doi.org/10.1109/PacificVis.2013.6596131>.
- [6] Englund R, Ropinski T. Quantitative and qualitative analysis of the perception of semi-transparent structures in direct volume rendering. *Comput Graph Forum* 2018;37:174–87. <http://dx.doi.org/10.1111/cgf.13320>, URL: <http://doi.wiley.com/10.1111/cgf.13320>.
- [7] Weissenböck J, Amirkanov A, Li W, Reh A, Amirkanov A, Gröller E, Kastner J, Heinzl C. FiberScout: An interactive tool for exploring and analyzing fiber reinforced polymers. In: 2014 IEEE Pacific visualization symposium. IEEE; 2014, p. 153–60. <http://dx.doi.org/10.1109/PacificVis.2014.52>, URL: <http://ieeexplore.ieee.org/document/6787162/>.
- [8] Le Muzic M, Mindek P, Sorger J, Autin L, Goodsell DS, Viola I. Visibility equalizer cutaway visualization of mesoscopic biological models. *Comput Graph Forum* 2016;35:161–70. <http://dx.doi.org/10.1111/cgf.12892>, URL: <https://onlinelibrary.wiley.com/doi/10.1111/cgf.12892>.
- [9] Lesar Ž, Alharbi R, Bohak C, Strnad O, Heinzl C, Marolt M, Viola I. Volume conductor: interactive visibility management for crowded volumes. *Vis Comput* 2023. <http://dx.doi.org/10.1007/s00371-023-02828-8>.
- [10] Englund R, Ropinski T. Evaluating the perception of semi-transparent structures in direct volume rendering techniques. In: SIGGRAPH ASIA 2016 symposium on visualization. New York, NY, USA: ACM; 2016, p. 1–8. <http://dx.doi.org/10.1145/3002151.3002164>.
- [11] Viola I, Gröller E, Hadwiger M, Bühler K, Preim B, Ebert D. Illustrative visualization. In: Lin M, Loscos C, editors. Eurographics 2005 - tutorials. The Eurographics Association; 2005, <http://dx.doi.org/10.2312/egt.20051052>.
- [12] Levoy M. Display of surfaces from volume data. *IEEE Comput Graph Appl* 1988;8:22–37. <http://dx.doi.org/10.1109/38.511>.
- [13] Viola I, Gröller E. Smart visibility in visualization. In: Neumann L, Sbert M, Gooch B, Purgathofer W, editors. Computational aesthetics in graphics, visualization and imaging. The Eurographics Association; 2005, p. 209–16. <http://dx.doi.org/10.2312/COMPAESTH/COMPAESTH05/209-216>.
- [14] Bruckner S, Gröller E. Enhancing depth-perception with flexible volumetric halos. *IEEE Trans Vis Comput Graphics* 2007;13:1344–51. <http://dx.doi.org/10.1109/TVCG.2007.70555>.
- [15] Schott M, Pegoraro V, Hansen C, Boulanger K, Bouatouch K. A directional occlusion shading model for interactive direct volume rendering. *Comput Graph Forum* 2009;28:855–62. <http://dx.doi.org/10.1111/j.1467-8659.2009.01464.x>, URL: <https://onlinelibrary.wiley.com/doi/10.1111/j.1467-8659.2009.01464.x>.
- [16] Díaz J, Ropinski T, Navazo I, Gobetti E, Vázquez P-P. An experimental study on the effects of shading in 3D perception of volumetric models. *Vis Comput* 2017;33:47–61. <http://dx.doi.org/10.1007/s00371-015-1151-6>.
- [17] Kersten M, Stewart J, Troje N, Ellis R. Enhancing depth perception in translucent volumes. *IEEE Trans Vis Comput Graphics* 2006;12:1117–24. <http://dx.doi.org/10.1109/TVCG.2006.139>.
- [18] Ropinski T, Steinicke F, Hinrichs K. Visually supporting depth perception in angiography imaging. In: International symposium on smart graphics. 2006, p. 93–104. [http://dx.doi.org/10.1007/11795018\\_9](http://dx.doi.org/10.1007/11795018_9).
- [19] Bruckner S, Gröller E. Instant volume visualization using maximum intensity difference accumulation. *Comput Graph Forum* 2009;28:775–82. <http://dx.doi.org/10.1111/j.1467-8659.2009.01474.x>, URL: <https://onlinelibrary.wiley.com/doi/10.1111/j.1467-8659.2009.01474.x>.
- [20] Díaz J, Vázquez P-P. Depth-enhanced maximum intensity projection. In: Westermann R, Kindlmann G, editors. IEEE/EG symposium on volume graphics. The Eurographics Association; 2010, <http://dx.doi.org/10.2312/VG/VG10/093-100>.
- [21] Kersten-Oertel M, Jy-Shyang Chen S, Collins DL. An evaluation of depth enhancing perceptual cues for vascular volume visualization in neurosurgery. *IEEE Trans Vis Comput Graphics* 2014;20:391–403. <http://dx.doi.org/10.1109/TVCG.2013.240>.
- [22] Weier M, Stengel M, Roth T, Diddy P, Eisemann E, Eisemann M, Grogoric S, Hinkenjann A, Kruijff E, Magnor M, Myszkowski K, Slusallek P. Perception-driven accelerated rendering. *Comput Graph Forum* 2017;36:611–43. <http://dx.doi.org/10.1111/cgf.13150>, URL: <http://doi.wiley.com/10.1111/cgf.13150>.

- [23] Lafortune EP, Willems YD. Rendering participating media with bidirectional path tracing. In: Proceedings of the eurographics workshop on rendering techniques '96. 1996, p. 91–100. [http://dx.doi.org/10.1007/978-3-7091-7484-5\\_10](http://dx.doi.org/10.1007/978-3-7091-7484-5_10), URL: [http://link.springer.com/10.1007/978-3-7091-7484-5\\_10](http://link.springer.com/10.1007/978-3-7091-7484-5_10).
- [24] Koerner D, Portsmouth J, Sadlo F, Ertl T, Eberhardt B. Flux-limited diffusion for multiple scattering in participating media. *Comput Graph Forum* 2014;33:178–89. <http://dx.doi.org/10.1111/cgf.12342>, URL: <http://doi.wiley.com/10.1111/cgf.12342>.
- [25] Kroes T, Post FH, Botha CP. Exposure render: An interactive photo-realistic volume rendering framework. In: Zuo X-N, editor. *PLoS One* 2012;7:e38586. <http://dx.doi.org/10.1371/journal.pone.0038586>, URL: <https://dx.plos.org/10.1371/journal.pone.0038586>.
- [26] Bohak C, Aleksandrov J, Marolt M. Collaborative web-based merged volumetric and mesh rendering framework. In: Paolis LTD, Bourdot P, editors. *Augmented reality, virtual reality, and computer graphics*. Cham: Springer International Publishing; 2019, p. 36–42.
- [27] Šmajdek U, Lesar Ž, Marolt M, Bohak C. Combined volume and surface rendering with global illumination caching. *Vis Comput* 2023. <http://dx.doi.org/10.1007/s00371-023-02932-9>.
- [28] Kronander J, Jönsson D, Low J, Ljung P, Ynnerman A, Unger J. Efficient visibility encoding for dynamic illumination in direct volume rendering. *IEEE Trans Vis Comput Graphics* 2012;18:447–62. <http://dx.doi.org/10.1109/TVCG.2011.35>.
- [29] Kniss J, Premoze S, Hansen C, Ebert D. Interactive translucent volume rendering and procedural modeling. In: *IEEE visualization 2002*. IEEE; 2002, p. 109–16. <http://dx.doi.org/10.1109/VISUAL.2002.1183764>.
- [30] Chan M-Y, Wu Y, Mak W-H, Chen W, Qu H. Perception-based transparency optimization for direct volume rendering. *IEEE Trans Vis Comput Graphics* 2009;15:1283–90. <http://dx.doi.org/10.1109/TVCG.2009.172>, URL: <http://ieeexplore.ieee.org/document/5290740/>.
- [31] Correa CD, Ma K-L. Visibility histograms and visibility-driven transfer functions. *IEEE Trans Vis Comput Graphics* 2011;17:192–204. <http://dx.doi.org/10.1109/TVCG.2010.35>, URL: <http://ieeexplore.ieee.org/document/5416704/>.
- [32] Howard IP. Depth perception. In: *Stevens' handbook of experimental psychology*. John Wiley & Sons, Ltd; 2002. <http://dx.doi.org/10.1002/0471214426.pas0103>, URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/0471214426.pas0103>.
- [33] Hu P, Boorboor S, Jadhav S, Marino J, Mirhosseini S, Kaufman AE. Spatial perception in immersive visualization: A study and findings. In: *2022 IEEE international symposium on mixed and augmented reality adjunct (ISMAR-adjunct)*. IEEE; 2022, p. 369–72. <http://dx.doi.org/10.1109/ISMAR-Adjunct57072.2022.00080>.
- [34] Wanger L. The effect of shadow quality on the perception of spatial relationships in computer generated imagery. In: *Proceedings of the 1992 symposium on interactive 3D graphics - SI3D '92*. New York, New York, USA: ACM Press; 1992, p. 39–42. <http://dx.doi.org/10.1145/147156.147161>.
- [35] Boucheny C, Bonneau G-P, Droulez J, Thibault G, Ploix S. A perceptive evaluation of volume rendering techniques. *ACM Trans Appl Percept* 2009;5:1–24. <http://dx.doi.org/10.1145/1462048.1462054>.
- [36] Titov A, Kersten-Oertel M, Drouin S. The effect of interactive cues on the perception of angiographic volumes in virtual reality. In: *Comput Methods Biomech Biomed Eng: Imaging Vis.* 2022;10:357–65. <http://dx.doi.org/10.1080/21681163.2021.1999332>.
- [37] Ropinski T, Doring C, Rezk-Salama C. Interactive volumetric lighting simulating scattering and shadowing. In: *2010 IEEE Pacific visualization symposium (Pacificvis)*. IEEE; 2010, p. 169–76. <http://dx.doi.org/10.1109/PACIFICVIS.2010.5429594>.
- [38] Englund R, Kottraval S, Ropinski T. A crowdsourcing system for integrated and reproducible evaluation in scientific visualization. In: *2016 IEEE Pacific visualization symposium (Pacificvis)*. IEEE; 2016, p. 40–7. <http://dx.doi.org/10.1109/PACIFICVIS.2016.7465249>.
- [39] Diaz C, Walker M, Szafir DA, Szafir D. Designing for depth perceptions in augmented reality. In: *2017 IEEE international symposium on mixed and augmented reality*. ISMAR, IEEE; 2017, p. 111–22. <http://dx.doi.org/10.1109/ISMAR.2017.28>.
- [40] Heinrich F, Bornemann K, Lawonn K, Hansen C. Depth perception in projective augmented reality: An evaluation of advanced visualization techniques. In: *25th ACM symposium on virtual reality software and technology*. New York, NY, USA: ACM; 2019, p. 1–11. <http://dx.doi.org/10.1145/3359996.3364245>.
- [41] Lesar Ž, Bohak C, Marolt M. Evaluation of angiogram visualization methods for fast and reliable aneurysm diagnosis. In: *Mello-Thoms CR, Kupinski MA, editors. SPIE medical imaging*. Vol. 9416, 2015, 94161D. <http://dx.doi.org/10.1117/12.2082179>, URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.2082179>.
- [42] Drouin S, Di Giovanni DA, Kersten-Oertel M, Collins DL. Interaction driven enhancement of depth perception in angiographic volumes. *IEEE Trans Vis Comput Graphics* 2020;26:2247–57. <http://dx.doi.org/10.1109/TVCG.2018.2884940>.
- [43] Kreiser J, Hermosilla P, Ropinski T. Void space surfaces to convey depth in vessel visualizations. *IEEE Trans Vis Comput Graphics* 2021;27:3913–25. <http://dx.doi.org/10.1109/TVCG.2020.2993992>.
- [44] Lesar Ž. Data from: Evaluation of depth perception in crowded volumes. 2024. <http://dx.doi.org/10.5281/zenodo.10555909>.
- [45] Žerovnik Mekuč M, Bohak C, Hudoklin S, Kim BH, Romih R, Kim MY, Marolt M. Automatic segmentation of mitochondria and endolysosomes in volumetric electron microscopy data. *Comput Biol Med* 2020;103693. <http://dx.doi.org/10.1016/j.compbiomed.2020.103693>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0010482520300792>.
- [46] Žerovnik Mekuč M, Bohak C, Boneš E, Hudoklin S, Romih R, Marolt M. Automatic segmentation and reconstruction of intracellular compartments in volumetric electron microscopy data. *Comput Methods Programs Biomed* 2022;223. <http://dx.doi.org/10.1016/j.cmpb.2022.106959>.
- [47] Schretter C, Kobbelt L, Dehaye P-O. Golden ratio sequences for low-discrepancy sampling. *J Graph Tools* 2012;16:95–104. <http://dx.doi.org/10.1080/2165347X.2012.679555>, URL: <http://www.tandfonline.com/doi/abs/10.1080/2165347X.2012.679555>.
- [48] McAuley S, Hill S, Hoffman N, Gotanda Y, Smits B, Burley B, Martinez A. Practical physically-based shading in film and game production. In: *ACM SIGGRAPH 2012 courses*. New York, NY, USA: ACM; 2012, p. 1–7. <http://dx.doi.org/10.1145/2343483.2343493>.
- [49] Toublanc D. Henyey–greenstein and mie phase functions in Monte Carlo radiative transfer computations. *Appl Opt* 1996;35:3270. <http://dx.doi.org/10.1364/AO.35.003270>, URL: <https://opg.optica.org/abstract.cfm?URI=ao-35-18-3270>.
- [50] Lesar Ž, Bohak C, Marolt M. Real-time interactive platform-agnostic volumetric path tracing in webgl 2.0. In: *Proceedings of the 23rd international ACM conference on 3D web technology - web3D '18*. New York, New York, USA: ACM Press; 2018, p. 1–7. <http://dx.doi.org/10.1145/3208806.3208814>, URL: <http://dl.acm.org/citation.cfm?doi=3208806.3208814>.
- [51] Keiner B, Innmann M, Sängner M, Stamminger M. Spherical fibonacci mapping. *ACM Trans Graph* 2015;34:1–7. <http://dx.doi.org/10.1145/2816795.2818131>.
- [52] Wexler M, van Boxtel JJA. Depth perception by the active observer. *Trends in Cognitive Sciences* 2005;9:431–8. <http://dx.doi.org/10.1016/j.tics.2005.06.018>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S1364661305002159>.
- [53] Mamassian P, Kersten D. Illumination, shading and the perception of local orientation. *Vis Res* 1996;36:2351–67. [http://dx.doi.org/10.1016/0042-6989\(95\)00286-3](http://dx.doi.org/10.1016/0042-6989(95)00286-3), URL: <https://linkinghub.elsevier.com/retrieve/pii/0042698995002863>.
- [54] Langer MS, Bühlhoff HH. Depth discrimination from shading under diffuse lighting. *Perception* 2000;29:649–60. <http://dx.doi.org/10.1068/p3060>.