

Structure-Aware Lighting Design for Volume Visualization

Yubo Tao, Hai Lin, Feng Dong, Chao Wang, Gordon Clapworthy, and Hujun Bao

Abstract—Lighting design is a complex, but fundamental, problem in many fields. In volume visualization, direct volume rendering generates an informative image without external lighting, as each voxel itself emits radiance. However, external lighting further improves the shape and detail perception of features, and it also determines the effectiveness of the communication of feature information. The human visual system is highly effective in extracting structural information from images, and to assist it further, this paper presents an approach to structure-aware automatic lighting design by measuring the structural changes between the images with and without external lighting. Given a transfer function and a viewpoint, the optimal lighting parameters are those that provide the greatest enhancement to structural information - the shape and detail information of features are conveyed most clearly by the optimal lighting parameters. Besides lighting goodness, the proposed metric can also be used to evaluate lighting similarity and stability between two sets of lighting parameters. Lighting similarity can be used to optimize the selection of multiple light sources so that different light sources can reveal distinct structural information. Our experiments with several volume data sets demonstrate the effectiveness of the structure-aware lighting design approach. It is well suited to use by novices as it requires little technical understanding of the rendering parameters associated with direct volume rendering.

Index Terms—Automatic lighting design, structural dissimilarity, lighting similarity, lighting stability, volume rendering.

1 INTRODUCTION

In volume visualization, direct volume rendering has been widely used for the analysis and presentation of volumes in a variety of fields, ranging from medicine and engineering to geophysics and biology. Meanwhile, many classification methods based on the information extracted from the volume have been proposed to identify and analyze features of interest. However, the perception of complex structures inside the volume depends not only on the classification, but also on the rendering parameters, most particularly the viewpoint and lighting parameters. The viewpoint determines how many features can be presented in the rendered image. The lighting has a major impact on the viewer's perception of the shape and detail of features [3]. For example, the effect of specular reflections [8] assists the judgement of convexity/concavity and the perception of the smoothness of structural surfaces.

While automatic viewpoint selection has received wide attention in volume visualization in recent years [2, 12, 25, 29], relatively little research has considered the lighting design problem in direct volume rendering. During volume rendering, once the transfer function and the viewpoint are specified, the visual perception of features in the volume is mainly decided by the lighting parameters. A result of poor lighting design is that meaningful information is not conveyed effectively to the users, which undermines the fundamental goal of visualization. Unfortunately, the specification of appropriate lighting parameters is a very complex and labor-intensive process, especially for novices due to their limited technical understanding of the rendering process.

Automatic lighting design has been well investigated for polygon rendering in computer graphics and isosurface rendering in scientific visualization. A perceptual quality metric [24] has been proposed to evaluate the ability of a rendered image to communicate information about the scene, and information entropy-based metrics (e.g., lighting entropy [10], illumination entropy and multi-scale entropy [26]) have also been developed for automatic lighting design.

Despite the fact that structural information plays a significant role in the human visual system (HVS) [27], none of these existing methods have considered the use of structural information to evaluate the lighting parameters. They evaluate images rendered individually under different lighting configurations, but then fail to assess how well these configurations improve the presentation of structural information.

In fact, evaluation of the enhancement of structural information that results from different lighting configurations can be implemented naturally for direct volume rendering. Direct volume rendering treats each voxel as a radiance emitter, and the basic information of features can be shown in a rendered image even without applying any external lighting. Hence, the quality of an external lighting configuration can be assessed by the amount of additional structural information conveyed by the lighting.

Based on this observation, this paper presents a novel method to evaluate lighting quality based on the structural information in the rendered images, which is in line with the basic mechanism of the HVS. An approach is proposed to support highly automated configuration of the lighting for direct volume rendering. The main contributions are summarized as follows:

1) We propose a multi-scale homogeneity-weighted structural dissimilarity metric (MS-HWSD), which captures the difference in structural information between two images. When the MS-HWSD is applied to measure the dissimilarity between images rendered with and without external lighting (the illuminated and unilluminated images, respectively), it indicates the quality of the lighting configuration in terms of its ability to provide structural information.

2) We also put forward three lighting characteristics in terms of *lighting goodness*, *lighting similarity* and *lighting stability* in a similar fashion to the view characteristics [1, 2]. Lighting goodness measures the lighting quality based on the MS-HWSD from (1); light similarity measures whether a light source is highly representative, and evaluates the complementarities between multiple light sources; lighting stability provides an extra measure on the light sources by considering the structural similarity of images rendered from neighboring light sources.

3) An automatic lighting design approach has been developed, based on the lighting characteristics proposed in (2). It supports the design of single and multiple light sources and can cope with different illumination models and find optimal solutions for various lighting parameters.

To the best of our knowledge, this is the first paper to address structure-aware lighting design for volume visualization.

The paper is structured as follows. Related work is discussed in Section 2, while Section 3 explains the motivation for the work, with

- Yubo Tao, Hai Lin, and Hujun Bao are with State Key Lab of CAD&CG, Zhejiang University, E-mail: {taoyubo, lin, bao}@cad.zju.edu.cn.
- Feng Dong, Chao Wang, and Gordon Clapworthy are with Centre for Computer Graphics & Visualisation, University of Bedfordshire, UK, E-mail: {Feng.Dong, Chao.Wang, Gordon.Clapworthy}@beds.ac.uk.

Manuscript received 31 March 2012; accepted 1 August 2012; posted online 14 October 2012; mailed on 5 October 2012.

For information on obtaining reprints of this article, please send e-mail to: ivcg@computer.org.

further details. Section 4 describes the MS-HWSD, and in Section 5, we propose three lighting characteristics. Section 6 describes the automatic lighting design techniques. In Section 7, we present and discuss the results obtained by applying the proposed techniques to various volumes and validate their effectiveness by user studies. Finally, we draw conclusions in Section 8.

2 RELATED WORK

Lighting design has long been an active research topic in computer graphics, and many lighting design methods have been proposed over the years. Traditional methods usually directly manipulate the lighting parameters, such as the light position and intensity. The specification of appropriate lighting parameters in volume visualization is generally a trial-and-error data-exploration process, which is time consuming and labor intensive, especially for novice users.

A great deal of attention has been paid to inverse lighting design, which is a major category of semi-automatic lighting design in computer graphics; it is surveyed in [19]. The general idea of inverse lighting design is to allow users to interact with the desired lighting effects through indirect user interfaces or to generate a desired lighting image through painting, based on which the lighting parameters are solved via optimization in an automatic manner. Poulin and Fournier [21] suggested direct manipulation of highlights and shadows to infer the optimal lighting parameters implicitly. Schoeneman et al. [22] optimized the lighting intensities at given positions to minimize the difference between the user-painted image and the rendered image. Pellacini et al. [20] proposed a non-linear optimization technique for a general painting interface to derive all the lighting parameters. However, the effectiveness of a painting interface is significantly less than direct and indirect interfaces, according to a later evaluation in [14], mainly because the desired goal images are usually obtained by sketching rather than painting. Furthermore, without a clear definition of features at the data exploration stage, it is difficult to directly manipulate the lighting effects or to paint the desired lighting images in order to derive appropriate lighting parameters for volume visualization.

Besides indirect and painting interfaces, several metrics have been proposed for automatic lighting design. Kawai et al. [13] presented subjective impressions of rendering qualities, such as pleasantness and privateness, which are maximized by optimizing lighting parameters. A perceptual quality metric involving six terms has been investigated for automatic lighting design by Shackled and Lischinski [24]. The goal is to evaluate the effectiveness in communicating information about the scene, such as shape, detail and relationships. No user study has been performed to validate this quality metric. Lighting entropy, which measures the brightness distribution of the visible pixels, has been used by Gumhold [10] to automatically determine the optimal light position, while Vázquez [26] developed illumination entropy and multi-scale entropy for lighting design by looking into color distribution and pixel correlation. However, all of these metrics are derived using the illuminated image only and do not take into account the presentation of structural information. In contrast, the proposed MS-HWSD in this paper follows the basic principles of the HVS by measuring structural differences between the unilluminated and illuminated images. Consequently, it can achieve considerably better results than those produced by existing methods.

There are several methods for optimizing lighting parameters in visualization. Marks et al. [18] introduced Design Galleries to explore the parameter space for light selection and placement. It automatically generates and organizes rendered images under various lighting parameters, and allows users to select the lighting parameters by choosing satisfying images. LightKit developed by Halle and Meng [11] suggests the default light parameters of the three-point lighting model, while Light Collages presented by Lee et al. [16] places light sources for surfaces to support better shape perception. In direct volume rendering, Chan et al. [4] adaptively refined the ambient, diffuse and specular coefficients to enhance the image quality. The lighting parameters are adjusted by minimizing the information deviation between the image and ray at each pixel.

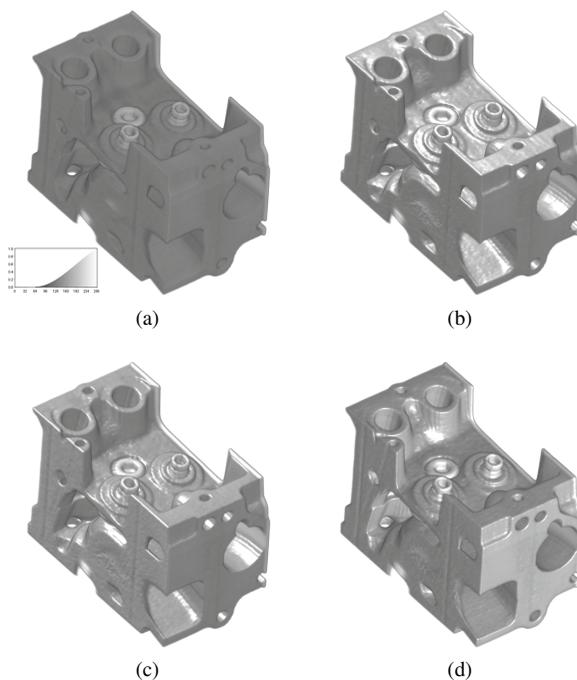


Fig. 1. Lighting quality evaluation based on the MS-HWSD. All three illuminated images have the same lighting entropy 2.41, but different MS-HWSD values for different lighting parameters. Higher MS-HWSD values mean the lighting reveals structural information more clearly compared to the unilluminated image. (a) The unilluminated image. (b) Illuminated image with MS-HWSD = 0.324. (c) Illuminated image with MS-HWSD = 0.249. (d) Illuminated image with MS-HWSD = 0.234.

3 MOTIVATION

This work is motivated by one of the basic assumptions of the HVS - that it is highly sensitive to structural information [27]. From this, the effectiveness of lighting can be assessed by the amount of structural information that is conveyed by the rendered image. Structural dissimilarity between two images can be quantified by calculating the variation of pixel colors in local regions.

Direct volume rendering usually regards each voxel as an illuminating particle with a certain level of density, and the mapping from the density to the optical property is specified by the classification. Each feature is assigned a distinct optical property, so some basic structural information of features can be revealed expressively without using external lighting. The purpose of applying external lighting is to emphasize shape and to highlight details on the surfaces of structures. The proposed approach for structure-aware lighting design searches for lighting parameters that will make interesting structures more perceptible.

An illuminated image can be thought of as the sum of an unilluminated image and the lighting effect, so the evaluation of the lighting effectiveness can be reasonably approximated by measuring the structural changes (i.e., dissimilarity) between the unilluminated and illuminated images. To this end, we introduce a structural dissimilarity metric (SDIM) to measure these changes.

On the other hand, over-strong or improper illumination may lead to the loss of structural details, so some structural information may be reduced by the introduction of external lighting. We therefore introduce a homogeneity metric to distinguish between structural enhancement and structural weakening.

By combining these two elements, this paper proposes a homogeneity-weighted structure dissimilarity metric (HWSD) to estimate the enhancement of structural information produced by the lighting. As the viewing conditions (e.g., image resolution and viewing distance) also have an important impact on the perceivability of struc-

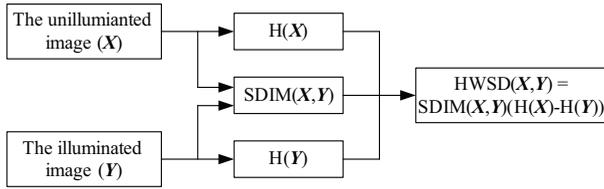


Fig. 2. The pipeline of homogeneity-weighted structural dissimilarity metric (HWSD). The structural dissimilarity metric (SDIM) measures perceptual structural differences between the unilluminated and illuminated images. The homogeneity metric (H) estimates how uniform the pixels are in a local region within an image, and is used to distinguish structural enhancement and degradation due to the lighting.

tural information from the image, we extend the proposed metric to the multi-scale HWSD (MS-HWSD) to incorporate structural changes at different resolutions. A clear advantage of the MS-HWSD is that it reflects the characteristics of the HVS. Fig. 1 illustrates the power of the MS-HWSD. The three images in Fig. 1(b)-(d) are rendered from different light positions. They have the same lighting entropy, as the information entropy-based metric is not sensitive to the HVS. The shape on the top of the engine and the detail on the structural surface are more clearly conveyed in Fig. 1(b), and the MS-HWSD can generate a consistently perceived difference of visual structures similar to the HVS.

Further, as in work on viewpoint selection, we propose three lighting characteristics: lighting goodness, lighting similarity and lighting stability, to assist the evaluation of different lighting configurations. While lighting goodness is purely based on the structure dissimilarity (i.e., the MS-HWSD), lighting similarity looks into the structural information revealed by different lights, and lighting stability helps to select stable light positions - these three elements form the basis of the proposed method for automatic lighting design.

4 MULTI-SCALE HOMOGENEITY-WEIGHTED STRUCTURAL DISSIMILARITY METRIC

The homogeneity-weighted structure dissimilarity metric (HWSD) is composed of the structural dissimilarity metric and the Homogeneity metric. The high-level procedure of the HWSD is shown in Fig. 2.

Structural Dissimilarity Metric. The Structural DISSIMilarity metric (SDIM) is based on the Structural SIMilarity (SSIM) index measure introduced by Wang et al. [27]. The SSIM was originally used to evaluate the degradation of visual quality between a reference image and a distorted image; in this, it regarded image degradations as the perceived structural changes. In the context of lighting design, the unilluminated image and the illuminated image are considered as the reference image and the distorted image, respectively, and the SSIM is used to measure perceived structural changes between them.

For completeness, we briefly review the SSIM. The SSIM measures the local structural similarity for each corresponding pair of pixels in their local neighborhoods. It separates the structural similarity measurement into three comparisons: luminance, contrast and structure.

Suppose x and y represent the local neighborhoods from the unilluminated and illuminated images, respectively. The number of the pixels in each neighborhood is N ($N = 11 \times 11$ is suggested in [27]). The luminance comparison is performed by the Gaussian weighted mean intensities $\mu_x = \sum_{i=1}^N w_i x_i$ and $\mu_y = \sum_{i=1}^N w_i y_i$. The luminance function $l(x, y)$ is defined as follows:

$$l(x, y) = \frac{2\mu_x \mu_y + c_1}{\mu_x^2 + \mu_y^2 + c_1}, \quad (1)$$

where c_1 is a constant based on the range of the pixel values and is included to avoid the introduction of a singularity. The luminance function gradually declines as the intensity of y increases. This formulation is qualitatively consistent with Weber's law, which states that a just-noticeable difference in the intensity change is approximately

proportional to the original intensity, as the HVS is more sensitive to a relative change in luminance, rather than an absolute one.

Similarly, the contrast function $c(x, y)$ is an estimate of the signal contrast and is defined as:

$$c(x, y) = \frac{2\sigma_x \sigma_y + c_2}{\sigma_x^2 + \sigma_y^2 + c_2}, \quad (2)$$

where $\sigma_x = (\sum_{i=1}^N w_i (x_i - \mu_x)^2)^{1/2}$ and $\sigma_y = (\sum_{i=1}^N w_i (y_i - \mu_y)^2)^{1/2}$, and the constant c_2 has the same functionality as c_1 . The contrast function monotonically decreases when the local variance of y becomes gradually different from that of x . This formulation also obeys the contrast masking feature of the HVS.

The SSIM assumes that the structural information in an image is determined by the structure of objects in the scene, and is independent of the average luminance and contrast. Thus, the structural similarity is quantified through a correlation between the images, as follows:

$$s(x, y) = \frac{\sigma_{xy} + c_3}{\sigma_x \sigma_y + c_3}, \quad (3)$$

where σ_{xy} defines the inner product between the vectors $x - \mu_x$ and $y - \mu_y$, $\sigma_{xy} = \sum_{i=1}^N w_i (x_i - \mu_x)(y_i - \mu_y)$, and the constant c_3 has the same functionality as c_1 . The structure function decreases when the intensity in y increases in a manner uncorrelated to the intensity in x .

The overall similarity measure is the composition of the three components, and the formulation of the SSIM is as follows:

$$\begin{aligned} SSIM(x, y) &= l(x, y) \cdot c(x, y) \cdot s(x, y) \\ &= \frac{(2\mu_x \mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)}, \end{aligned} \quad (4)$$

where $c_3 = c_2/2$ to simplify the formulation. It is worth pointing out that the three components are relatively independent. For example, a change of luminance or contrast does not influence the structure of the image. Fig. 3 shows two illuminated images of the vismale data set, together with their SSIM images. Fig. 3(c)-(e) show the luminance, contrast and structure comparisons between (a) and (b), respectively.

The SDIM proposed in this work is a distance metric for measuring structural differences between the unilluminated and illuminated images. The SDIM is the converse of SSIM, defined as follows:

$$SDIM(x, y) = 1 - SSIM(x, y). \quad (5)$$

As the SSIM value lies in $(-1, 1]$, the valid range of the SDIM is $[0, 2)$, with higher values indicating higher structural dissimilarity. When the lighting has no effect on features in the illuminated image (i.e., $x = y$), the SDIM reaches its minimum 0. Loza et al. [17] introduced another structural dissimilarity definition $DSSIM = \frac{1}{SSIM} - 1$. However, this is non-monotonic for the SSIM range $(-1, 1]$, so it is inappropriate for lighting quality evaluation.

The SDIM measures structural differences in terms of the just-noticeable intensity, local contrast and structures of the features between the unilluminated and illuminated images. For instance, the specular reflection enhances shape perception of features [8], and the SDIM can estimate these differences by the just-noticeable intensity change on the surfaces of features. Edges and contours contain a wealth of information about the structures of features [15], and the SDIM can identify such information by evaluating the changes in local contrast due to the lighting. The SDIM can also capture the increasing of local variances of the details on the surface revealed by the lighting.

Homogeneity Metric. By the symmetry of Equation 5, the SDIM measures only structural differences between the two images, without indicating which contains the clearer structures. This is sufficient for the evaluation of image degradation, where the reference image is known to have satisfactory quality, so the degradation must lie in the other image. However, as previously mentioned, in lighting design, additional external lighting may either enhance or diminish structural perception, such as the highlights in Fig. 3(h), so we need to further

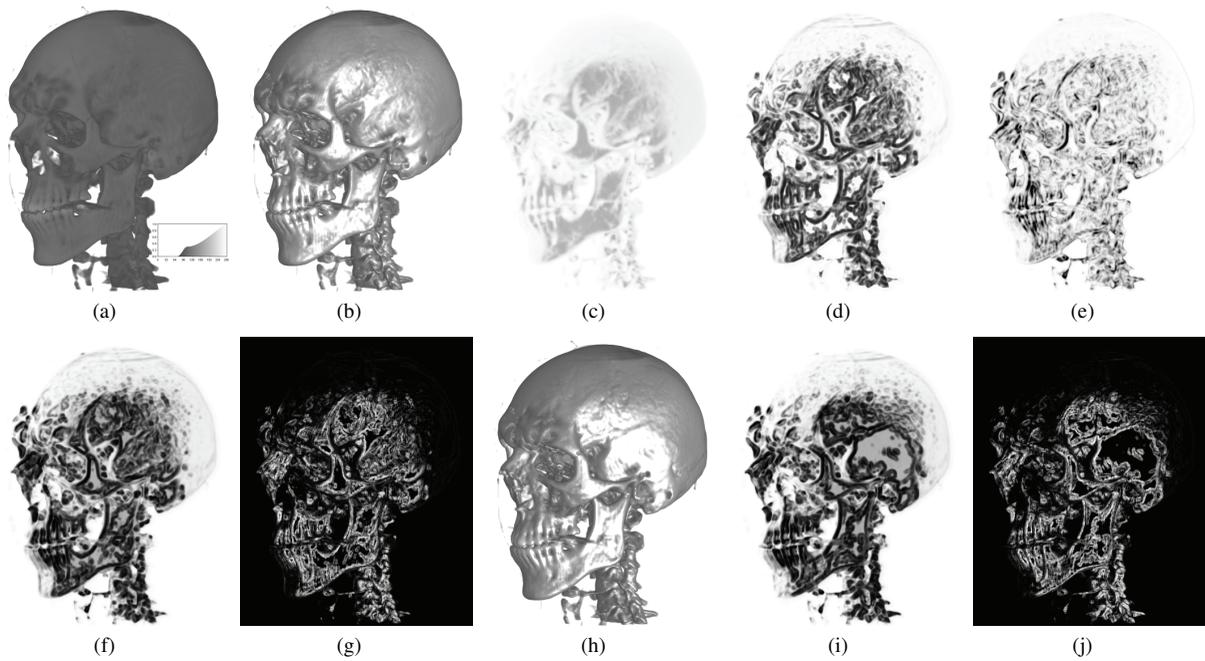


Fig. 3. SSIM and HWSD results of the vismale data set under two different light sources. Darker areas in SSIM results indicate larger structural differences between the unilluminated and illuminated images. Brighter areas in HWSD results mean greater structural enhancement between the unilluminated and illuminated images. The mean SDIM [= (1 - SSIM)] of the illuminated image (b) is lower than the value of the illuminated image (h), but the mean HWSD of (b) is larger than the value of (h). This is because the homogeneity metric successfully distinguishes enhancement and degradation in structural perception, and the HWSD values of the over-illuminated region in (h) are negative in the corresponding region in (j). (a) The unilluminated image. (b) An illuminated image. (c) The SSIM's luminance comparison of (a) and (b). (d) The SSIM's contrast comparison of (a) and (b). (e) The SSIM's structure comparison of (a) and (b). (f) The SSIM result of (a) and (b). (g) The HWSD result of (a) and (b). (h) Another illuminated image. (i) The SSIM result of (a) and (h). (j) The HWSD result of (a) and (h).

distinguish between these possibilities when considering the effect of using external lighting.

Insufficient and excessive illumination can both fail to present structural details effectively [10], but insufficient illumination is not applicable to our approach. This is because, compared to the unilluminated image, lighting never locally decreases the brightness of features in the illuminated image (if the ambient lighting coefficient is 1), so we need merely to avoid over-illumination.

In this work, the homogeneity metric is used to identify the over-lit areas in the illuminated image. Homogeneity is a metric indicating how uniform the pixels are in a local region within an image [5]. This is estimated by applying the standard deviation σ and the entropy e to a local neighborhood x of size N . The definition of the standard deviation σ_x is the same as that of the contrast function of the SSIM, and the formulation of the entropy e_x in the local region is defined as follows: $e_x = -\sum_{j=1}^J p_j \log p_j$, where $p_j = \frac{\#x_i=j}{N}$ is the probability of the pixel value at the j -th bin of the histogram of x , and the intensity level in the histogram is J . The maximum entropy occurs when all pixel values in x are different (at different bins of the histogram), and the entropy reaches its minimum when all pixel values are in the same bin of the histogram.

Although the original homogeneity definition in [5] further contains the edge value and the 4-th moment of the intensity distribution components, we found that, in practice, they make little contribution to the lighting design results, and a reduced form offers an equivalent capability of evaluating the uniformity and improves the computational efficiency. For example, the use of the original homogeneity generates the same optimal light source as the reduced form for the engine data set in Fig. 5(a). Hence, we define the homogeneity metric of a local region x as follows:

$$H(x) = \left(1 - \frac{\sigma_x}{\sigma_{max}}\right) \left(1 - \frac{e_x}{e_{max}}\right), \quad (6)$$

where σ_{max} and e_{max} are the maximum values of the standard deviation and the entropy in both unilluminated and illuminated images, respectively. The homogeneity increases when the intensities in the local region become much more similar. The valid range of the homogeneity is $[0, 1]$, with higher values indicating higher uniformity.

When the illuminated image is compared to the unilluminated image, the homogeneity metric provides a good measure of the structural weakening from over-illumination - a large increase of homogeneity implies the creation of more flat regions due to the use of lighting that is too strong, while a significant decrease suggests structural enhancement resulting from lighting that is more appropriate.

Homogeneity-Weighted Structural Dissimilarity Metric. We combine the structural-dissimilarity metric and the homogeneity metric to form a Homogeneity-Weighted Structural Dissimilarity metric (HWSD) as follows:

$$HWSD(x, y) = SDIM(x, y)(H(x) - H(y)), \quad (7)$$

where $H(x)$ and $H(y)$ are the homogeneity values of the local neighborhoods x and y from the unilluminated and illuminated images, respectively. The HWSD value grows quickly when the lighting reveals structures more clearly than the unilluminated image.

The HWSD utilizes the difference of the homogeneity values to distinguish structural enhancement and degradation. When the lighting enhances the structures in the illuminated image, the homogeneity value $H(y)$ decreases compared to $H(x)$, and this results in a positive HWSD value. Higher positive HWSD values represent more structural enhancement in the illuminated image. On the other hand, when the lighting weakens the structures in the illuminated image, for example, over-illumination produces more flat and less structural regions with the increased homogeneity value $H(y)$, and this leads to a negative HWSD value to penalize this situation. Negative HWSD values indicate that the external lighting actually diminishes the overall structural perception in the illuminated image. Figs. 3(g) and (j) show two

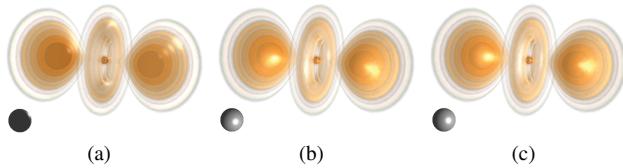


Fig. 4. Visual comparison of three lighting design methods for a hydrogen atom data set. The sphere in the bottom-left corner is used to illustrate the lighting effect. (a) Optimal light source using lighting entropy. (b) Optimal light source using multi-scale entropy. (c) Optimal light source using the proposed MS-HWSD.

HWSD results. Compared to Fig. 3(g), the HWSD values of the over-illuminated region in Fig. 3(h) are negative in the corresponding region in Fig. 3(j).

Since the SDIM and homogeneity metric are both defined in local neighborhood regions, the HWSD preserves the information of the pixels' correlation properly, which is an important factor in lighting design [26]. In contrast to the symmetry of SSIM and SDIM, the HWSD is not commutative, as $\text{HWSD}(x, y) = -\text{HWSD}(y, x)$. The valid range of the HWSD is $(-2, 2)$.

Multi-Scale HWSD. As the perceptibility of structural information in the image depends on its resolution and viewing distance, a single-scale method may be inappropriate for all lighting evaluations. Additionally, the multi-scale extension of SSIM (MS-SSIM) proposed in [28], has proven to produce better results than the single-scale SSIM for image quality assessment.

We propose a multi-scale HWSD metric (MS-HWSD) to evaluate the structural enhancement generated by the lighting at different resolutions. A low-pass filter is employed to downsample the unilluminated and illuminated images successively by a factor of 2, and the HWSD is applied to the filtered images at each scale. The MS-HWSD is obtained by averaging all single-scale HWSD values as follows:

$$\text{MS-HWSD}(X, Y) = \frac{1}{N_s} \sum_{j=1}^{N_s} \left(\frac{1}{M^j} \sum_{i=1}^{M^j} \text{HWSD}(x_i^j, y_i^j) \right), \quad (8)$$

where X and Y are the unilluminated and illuminated images, respectively, x_i^j and y_i^j are the local neighborhoods at the i -th pixel of the filtered unilluminated and illuminated images at the j -th scale, M^j is the total number of pixels in the filtered image at the j -th scale, and N_s is the number of scales used. The inner summation is a mean HWSD, which evaluates the overall structural enhancement at the j -th scale by taking the average over all pixels in the image. As the MS-HWSD incorporates all enhancements of the structural information at the multiple scales, it improves the efficiency and robustness of the HWSD.

5 LIGHTING CHARACTERISTICS

As mentioned in Section 3, to evaluate and compare different lighting configurations, we propose three lighting characteristics: lighting goodness, lighting similarity and lighting stability.

Lighting Goodness. Lighting goodness measures the effectiveness of different lighting configurations in terms of conveying structural information. External lighting may improve various aspects of structural perception, such as the shape of features, details of surfaces and feature occlusion relationships. A lighting configuration is good if it reveals a greater level of structural information. The SDIM is designed to measure the structural changes between unilluminated and illuminated images, and the MS-HWSD identifies the structurally enhanced regions in the illuminated image at different scales. Thus, the MS-HWSD can be used to measure lighting goodness - it is desirable that the lighting has a high MS-HWSD value. Fig. 4 shows three rendered results under different light configurations. In terms of the MS-HWSD, Fig. 4(c) has the highest value, while Fig. 4(a) has the lowest value.

Lighting Similarity. Lighting similarity estimates the similarities with regard to the structural information revealed by different lighting configurations. This is a very useful measure to see whether a light source is highly representative, and also to judge the complementarities between two light sources.

The multi-scale extension of SSIM (MS-SSIM) [28] is a good metric to evaluate the structural similarity of illuminated images under different lighting configurations. To compare two different lighting configurations, the MS-SSIM is calculated directly based on two illuminated images, in contrast to the MS-HWSD based on the unilluminated and illuminated images. Two lighting configurations with a high MS-SSIM value are similar, as they convey a great deal of similar structural information in their illuminated images. However, lighting similarity based on the MS-SSIM is not transitive, i.e., if lighting configuration A is similar to lighting configuration B and C, it does not mean light configuration B is similar to light configuration C, and actually the illuminated images under lighting configuration B and C may be significantly different. In Fig. 4, the lighting configuration in Fig. 4(c) is the most similar to that in Fig. 4(b), and the lighting configuration in Fig. 4(a) is the most dissimilar to the other lighting configurations.

Lighting similarity can be used to search for representative lighting configurations. If a lighting configuration is similar to many others, it is representative and hence is a good candidate for conveying the corresponding structures. Another application of lighting similarity is that the design of multiple light sources should consider lights that, individually, produce low similarity as they can enhance the perception of different structures.

Lighting Stability. Lighting stability indicates the amount of change of structural perception when the position of a light source is moved within a small neighborhood. It can be also evaluated through the MS-SSIM. The MS-SSIM is applied successively to the image illuminated from the current investigated position and the images illuminated from each of a set of neighbouring positions. These MS-SSIM values are summed to obtain the quantified stability value for the current lighting position; a larger stability value implies a more stable light source. In lighting design, lighting stability can be integrated with lighting goodness to help choose an optimal light in terms of both structural perception and stability.

6 AUTOMATIC LIGHTING DESIGN

Many illumination models can be used for direct volume rendering. Without loss of generality, let us firstly consider the Blinn-Phong local illumination model. This model can be calculated in real time, and hence is well suited to fast quality evaluation of the lighting. Formally, the color of a rendered voxel can be expressed as:

$$C = (k_a + k_d(\hat{N} \cdot \hat{L}))C_{TF} + k_s(\hat{N} \cdot \hat{H})^n, \quad (9)$$

where k_a , k_d , and k_s are the ambient, diffuse, and specular lighting coefficients respectively, n is the shininess exponent, C_{TF} is the color from the transfer function, \hat{N} is the normalized gradient direction of the voxel, \hat{L} is the light direction, and \hat{H} is the normalized half-way direction. In order to distinguish features in the visual representation, the color from the transfer function is used for the ambient and diffuse components. Light sources are chosen to be white to preserve the original hue for the visual labeling of features, and this is applied only to the specular component.

First, we consider the selection of the light position for a single point light source at a fixed viewpoint. The candidate positions for the light source are considered to be located on a fixed-distance bounding sphere in the volume, referred to as the *light-position sphere*. The HEALPix package [9] is used to generate evenly distributed candidate positions over the light-position sphere.

A number of candidate positions (we use 768) are generated as the initial candidates. The algorithm evaluates the MS-HWSD values of these candidate positions, and then generates sampling positions in regions where the MS-HWSD values are highest to further search for the optimal light position that is, the one with the highest value. The

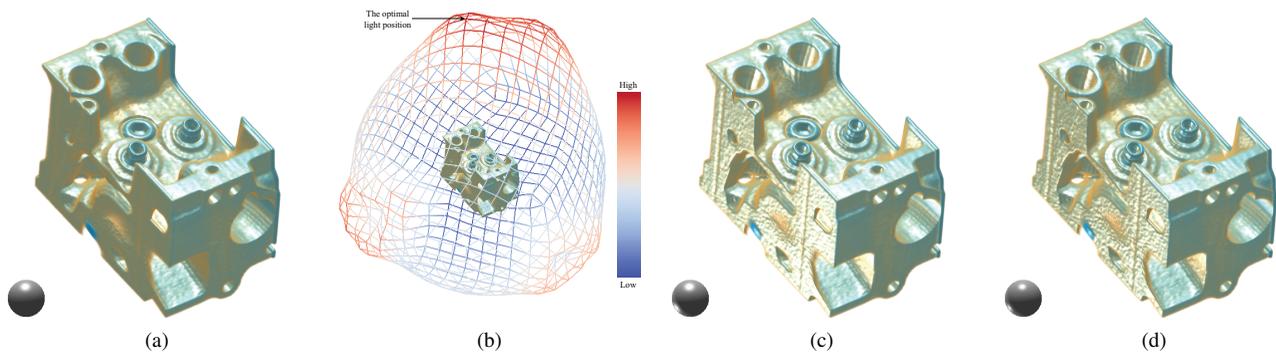


Fig. 5. Automatic lighting design for the engine data set. (a) The first optimal light source. (b) The deformed light-position sphere for the first optimal light source. The radius of each light position is proportional to the normalized lighting quality (the MS-HWSD) and the color mapping from blue to red corresponds to lighting quality from low to high. (c) Two optimal light sources ($\alpha = 0.4$ in Equation 10). The shininess exponent is 40. (d) The optimal shininess exponents for the two light sources in (c). The shininess exponent is searched from 20 to 80 with step size 1. The optimal shininess exponents are 47 and 74 for the two lights. Highlights are reduced to improve structure perception.

optimal light position for the engine data set is shown in Fig. 5(a), in which the shape of large structures and the detail on the top surface are clearly revealed. A deformed light-position sphere is illustrated in Fig. 5(b) for the lighting quality comparison (the MS-HWSD) of this selection. It is clear that, for the engine data set viewed from this point, there are three areas of high lighting quality.

Multiple Light Sources. The automatic design process for single light sources can be extended to consider the design when using multiple light sources. A possible approach would be to search for the optimal light positions by evaluating all possible combinations of the candidate positions, though this has the obvious shortcoming of poor scalability with the number of light sources. If M is the number of the candidate positions and K is the number of light sources, there are $O(M^K)$ lighting configurations.

In practice, users usually place light sources one by one. Although this may not lead to the globally optimal configuration, it is easy and flexible for users, and it allows straightforward use of the MS-HWSD. After the first optimal light source has been selected, the image rendered under the first light source serves as the unilluminated image for the next selection. Thus, light sources can be added successively until a satisfactory result is obtained. This approach scales well with the number of light sources, as the computational complexity is $O(MK)$.

During the design of multiple light sources, it is desirable that different light sources reveal different structural information. Two light sources positioned closely together will put light energy into similar regions; while this may make, for example, the areas near edges more perceptible by increasing the contrast, it does not introduce any additional structural information. For this, we use lighting similarity as a second criterion in the search for the next light source. This light source should not only further enhance structural information in the rendered image but should also have low similarity in relation to each of the previously selected light sources. Thus, the measure for the next light source l is modified as follows:

$$E(l) = \alpha MS\text{-}HWSD(X, Y) + (1 - \alpha) \frac{1}{K_s} \sum_{i=1}^{K_s} MS\text{-}SDIM(Z_l, Z_i), \quad (10)$$

where X is rendered by all previously selected light sources, Y is rendered by adding the current light source l , Z_l and Z_i are rendered individually under the light source l and the i -th selected light source, respectively, and K_s is the number of previously selected light sources. The MS-SDIM is the multi-scale extension of SDIM as follows: $MS\text{-}SDIM(Z_l, Z_i) = 1 - MS\text{-}SSIM(Z_l, Z_i)$, where the MS-SSIM [28] is the multi-scale SSIM. The parameter α in the range $[0, 1]$ is used to balance the structural enhancement in the first term and the dissimilarity between the light sources in the second term. The user can adjust α according to their specific applications. If $\alpha = 0$, the user obtains

a light source which is most dissimilar to the previously selected light sources, while if $\alpha = 1$, the next light source has the highest lighting goodness value, but lighting similarity is not considered. From our experiments, $\alpha = 0.4$ offers a good balance.

Using multiple light sources significantly improves the results. Fig. 5(c) depicts the optimal configuration of two light sources for the engine data set. Compared to Fig. 5(a), more structural details can be perceived under two light sources, especially the details on the side of the engine.

Other Lighting Parameters. Apart from introducing additional light sources, other parameters of the light source can also be adjusted automatically to provide better structure perception, based on the MS-HWSD. These include the diffuse and specular coefficients and the shininess exponent. The optimal lighting parameter is sought in the parameter domain by iteratively evaluating the goodness measure of the current lighting parameters for each light position. An extension from using point light sources to directional light sources is also straightforward, as the light direction can be directly sampled on the bounding sphere of the volume as for the light position. Fig. 5(d) shows the optimal shininess exponents for the two optimal light sources in Fig. 5(c).

Other Illumination Models. As the lighting design method requires only unilluminated and illuminated images, it can easily be applied to other illumination models.

Local illumination models, such as Blinn-Phong, primarily reveal the structural shape and local details, while global illumination models further depict the occlusion relationships between structures through mutual shadowing. Thus, we integrate shadows into direct volume rendering and employ our approach to select the optimal lighting parameters. Shadows can increase the MS-HWSD value by introducing the shadowing edges, which is a clear occlusion aspect of the spatial relationships between features. On the other hand, as shadows remove lighting from features, the MS-HWSD value would be reduced compared to local illumination models, and this is helpful for avoiding the over-shadowing associated with complex geometries, which produces complex, highly structured cast shadows that reduce structural perception. Thus, the MS-HWSD favors a light position with a good balance between the benefits of the shadowing edges introduced and the drawbacks of over-shadowing. Fig. 6 shows the optimal results of the light position for a vortex data set without and with shadowing. Compared to the rendered result without shadowing in Fig. 6(a), the rendered result in Fig. 6(b) provides more information about the structure arrangement by the inclusion of shadows, but does not have so many shadow regions that the local contrast of structures is reduced.

Stable Light Source. In some applications, users require stable lighting conditions. For this, lighting stability can be integrated with

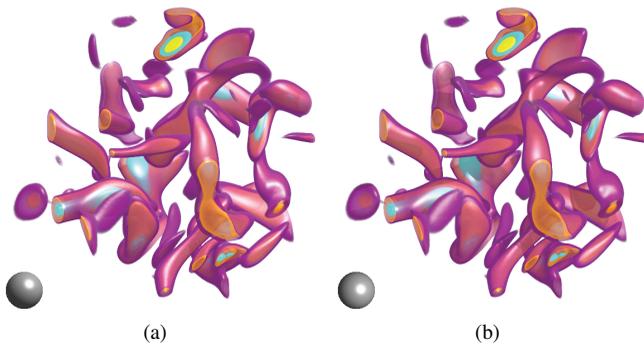


Fig. 6. Automatic lighting design without and with shadowing for the vortex data set. Shadowing provides more perceptual cues for the spatial relationships of features. (a) Optimal light source without shadowing. (b) Optimal light source with shadowing.

lighting goodness, with the modified lighting quality being defined by:

$$E(l) = \beta MS\text{-}HWSD(X, Y) + (1 - \beta) \frac{1}{K_n} \sum_{i=1}^{K_n} MS\text{-}SSIM(Y, Z_i), \quad (11)$$

where X is the unilluminated image, Y is an illuminated image under the light source l , and Z_i is an illuminated image under the i -th neighbor light position, K_n is the number of neighbor light positions ($K_n = 8$ is satisfactory), and the parameter β is in the range $[0, 1]$. The lighting quality is a weighted combination of the structural enhancement in the first term, and the stability between the light being considered and its neighbors in the second term. Users can adjust β according to their needs; if β is reduced, the selected light source becomes more stable.

7 RESULTS AND DISCUSSION

We have implemented the automatic lighting design discussed above. The rendered results are generated through pre-integrated direct volume rendering [7], which is implemented in a GPU-based ray-casting volume renderer. The shadow ray technique is integrated into the volume renderer to compute shadow effects.

The unilluminated and illuminated images are two 512×512 floating-point textures. The MS-HWSD is applied to the intensity channel of gray images, while for the color images, it is obtained by summing the MS-HWSD values calculated for the R, G, B components individually. As the HVS normally uses intensity variations to determine shape, the sum is a rough approximation of the intensity. On the other hand, as different features are often classified by distinct colors, the separate evaluation can preserve this visual labeling information of features. The number of scales used for results presented in this paper is 4. The MS-HWSD and MS-SSIM evaluation are performed on the GPU, which improves the computational efficiency by exploiting the massive parallelism of the GPU and avoiding texture read-back. Although the computational times for the automatic lighting design from a fixed viewpoint depend greatly on the data set due to the need to render the illuminated images, the computational times for the MS-HWSD evaluation are relatively fixed. In our experiments with image sizes of 512×512 , the average evaluation was about 12 seconds for all 768 light positions on an Intel Core i5 760 (2.80 GHz) processor with an NVIDIA GeForce GTX 470.

The proposed automatic lighting design approach is used to seek an adequate configuration of the lighting parameters. We now consider the effectiveness of each of the three lighting characteristics described in Section 5.

Lighting Goodness. Figs. 5 and 6 show the lighting design results of the engine and vortex data sets based on the proposed lighting goodness criterion for single light source, two light sources and shadowing. The selected light sources improve shape and detail perception

by enhancing local contrast, and introduce shadows which assist the understanding of spatial occlusion relationships.

As can be seen from Fig. 5(b), the lighting qualities above the engine are better than those behind it. This is usually true for opaque objects, as the front surfaces receive less lighting from lights behind the object to enhance the visual contrast of the global structures and surface details, and this makes them less perceptible for scene interpretability. The selection of the first light source for opaque objects is similar to the key light in the three-point lighting model, which is a standard method used in visual media. However, there is a significant difference: each voxel of the volume is taken as a radiance emitter, while the subject in the three-point lighting model usually does not emit lights. Nevertheless, we can still follow the principles of three-point lighting, for example we can apply the MS-HWSD only to the light positions within the valid range of the key light and search for the optimal key light for opaque objects.

The light behind the scene, called the rim light in the three-point lighting model, is normally used to visually separate the subject from the background. In volume visualization, a light behind transparent structures can also produce good visual perception as the radiance will not be occluded by the back structures and can traverse the whole volume to reveal the internal structures clearly. In Fig. 4(a), the optimal light position selected by lighting entropy [10] is located at the back of the light-position sphere. Fig. 7(d) gives another example for the engine data set. The global structure of the engine in the unilluminated image is very clear due to the appropriate classification, but the light source selected by the MS-HWSD provides shape and boundary cues for the internal structures and it also conveys more structural details on the left and right surfaces of the engine.

As the radiance from the external light is slightly absorbed in the transparent structures, its contribution to the rendered result is limited. It is not easy to recognize the lighting effects on transparent structures. Thus, automatic lighting design approaches usually favor structures that are more opaque. This is reasonable in volume visualization, since the users often highlight important structures with a high opacity value to make them more visible during the classification process. Fig. 8 shows an example for the vismale head data set. The skin is set transparent to make the skull visible. Although there may be many structural details on the skin, the optimal light sources selected by lighting entropy, multi-scale entropy [26] and the MS-HWSD all improve structure perception much more for the skull than for the skin. In addition, it is difficult to enhance the perception of transparent structures against a white background, as white light enhances the luminance value of sampled points during the ray casting, and this makes the illuminated color and the background hard to distinguish.

Three methods have been used to select the optimal light sources for the seven volume data sets used in the visual comparisons in Fig. 4, Fig. 7 and Fig. 8. Lighting entropy prefers lighting effects with an even intensity distribution. As the unilluminated image of the hydrogen atom data set has many areas of high intensity, the optimal light source is selected by lighting entropy to retain the original dark areas, as shown in Fig. 4(a). This lighting makes little contribution to the internal shape perception compared to the other selected light sources. Multi-scale entropy is based on 4 levels of Haar wavelet decomposition and is calculated by the Shannon entropy of the wavelet coefficients. Thus, multi-scale entropy prefers an even distribution of the wavelet coefficients. As shown in Fig. 7(b), the result from multi-scale entropy is more easily affected by rendering artifacts, as these tend to make the intensity distribution of the wavelet coefficients more even. The MS-HWSD is derived from the principles of the HVS, and the optimal lighting parameters found by the MS-HWSD generate rendered images that convey the most significant structural information compared to the unilluminated image, such as the shape in Fig. 7(b) and the detail in Fig. 7(d). The MS-HWSD prefers lighting parameters that render images with better luminance, contrast and structure enhancements, as shown in Fig. 3.

Lighting Similarity. Fig. 5 shows the effectiveness of lighting similarity in the design of multiple light sources. Fig. 7(f) presents an example for a tomato data set. The two light sources selected

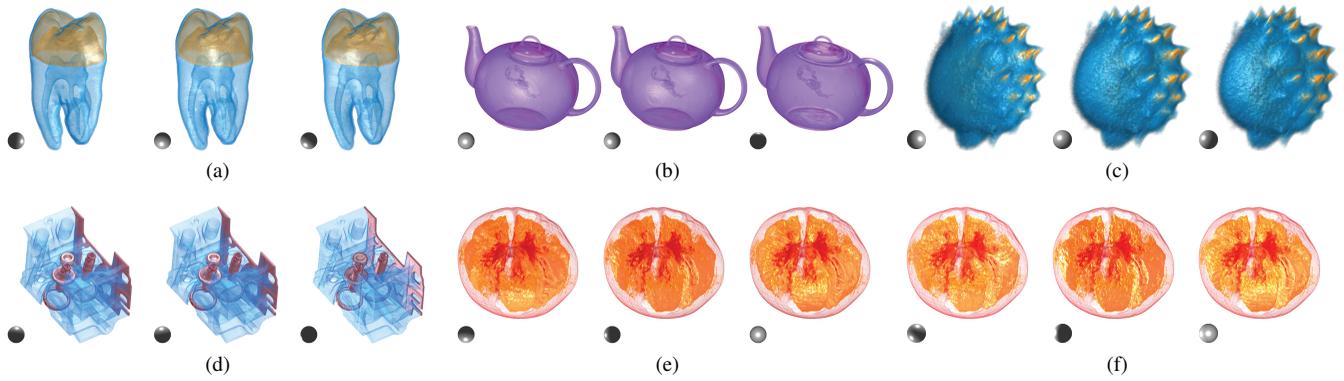


Fig. 7. Visual comparison of three lighting design methods for five volume data sets. The optimal lighting parameters from left to right in each sub-figure are selected using lighting entropy, multi-scale entropy and the proposed MS-HWSD, respectively: (a) tooth, (b) teapot, (c) daisy pollen grain, (d) engine, (e) tomato with one light source, (f) tomato with two light sources.

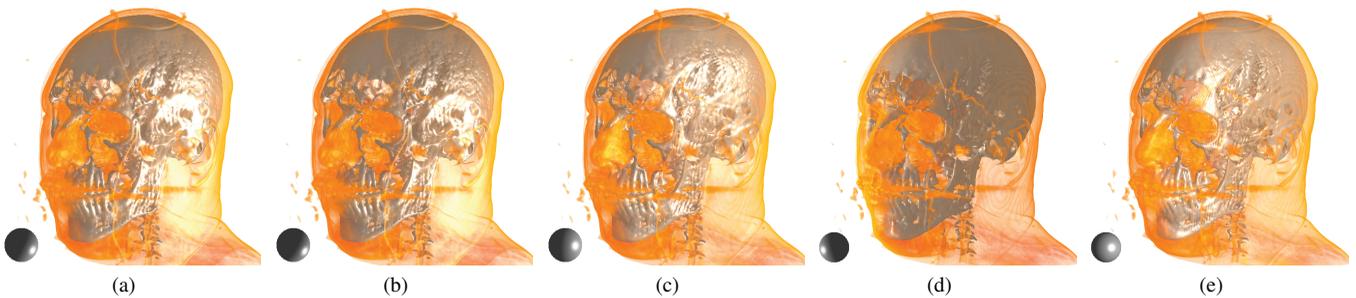


Fig. 8. Visual comparison of three lighting design methods and lighting stability analysis for the vismale data set: (a) optimal light source using lighting entropy, (b) optimal light source using multi-scale entropy, (c) optimal light source of the proposed MS-HWSD ($\beta = 1.0$ in Equation 11), (d) most stable light source ($\beta = 0.0$), (e) optimal light source based on lighting goodness and stability ($\beta = 0.4$).

by lighting entropy are spatially separated, as it also considers light source independence based on conditional entropy. However, multi-scale entropy does not take into account light source independence, and hence the selected light sources have a high overlap of illumination. In the proposed greedy approach, lighting similarity is used as the light source independence criterion to select the next light source. Although it may not find the same light sources as selected by the brute-force approach (Fig. 5 has the same light sources as selected by the brute-force approach), the proposed greedy approach has good scalability and can obtain a satisfactory result. Furthermore, the user can adjust the parameter α to balance the structural enhancement and the dissimilarity between the light sources.

An example in which representative light sources are selected is shown in Fig. 9. Fig. 5(b) is the deformed light-position sphere for the selection of the first representative light source, which is the one with the highest MS-HWSD value, as shown in Fig. 5(a). We then select as the next representative light source, the one with the smallest similarity to the previously selected light sources - this is expected to reveal structural information that was not presented by the previous light sources. This corresponds to the case $\alpha = 0$ in Equation 10. The deformed similarity sphere compared to the first representative light source is displayed in Fig. 9(a), and the second representative light source is selected from this sphere. As shown from the deformed light-position sphere, the worst light position in Fig. 5(b) does not necessarily have the smallest similarity with the optimal light position. Similarly, Fig. 9(b) shows the deformed similarity sphere compared to the first two representative light sources, and the third representative light source can be selected from this sphere. These three representative light sources collectively give a good visualization of the engine data set by revealing different structures through lighting effects.

Lighting Stability. Fig. 8(d) shows the most stable light position for the vismale data set, and the optimal light source based only on the MS-HWSD is shown in Fig. 8(c). The most unstable light position is usually opposite the viewpoint and produces few lighting effects on structures. In this case, a small deviation from the most unstable light position would generate significantly different lighting effects. Lighting goodness and lighting stability are combined in Equation 11 to search for a stable and good light source. As shown in Fig. 8(e) with $\beta = 0.4$, the shape and local details on the frontal skull are more easily perceptible.

User Study. A user study was performed to verify that the proposed MS-HWSD is an effective lighting goodness metric by comparing two images with the same volume and viewpoint, but with different light positions. We randomly selected 20 pairs of light positions for each of the vismale and engine data sets illustrated in Figs. 3 and 5, respectively. As in the viewpoint comparison, it is effective to compare pairs with nearby light positions as the lighting effects of the pairs are sufficiently different to allow comparisons but not so different that the comparisons become meaningless. We used an angular separation of $\frac{\pi}{8}$ radians [23]. In the random selection, we rejected light positions that were close to previously selected light positions, and discouraged pairs whose MS-HWSD values were too close for meaningful comparisons, such as light positions that were both behind the object. Fig. 10(a) shows the selected pairs of distributions.

As the definition of good lighting is highly subjective, and this judgement is hard to quantify, we use paired comparisons [6] and asked users: "Which of the two images reveals the object's structural information (shape and detail) better?". We randomly select 10 pairs from the 20 selected pairs for each data set, giving a total of 20 pair comparisons for each user. The pairs were presented in a random sequence, and the two images of each pair were shown in a random left-right order; 33 users were used, giving a total of 660 comparisons;

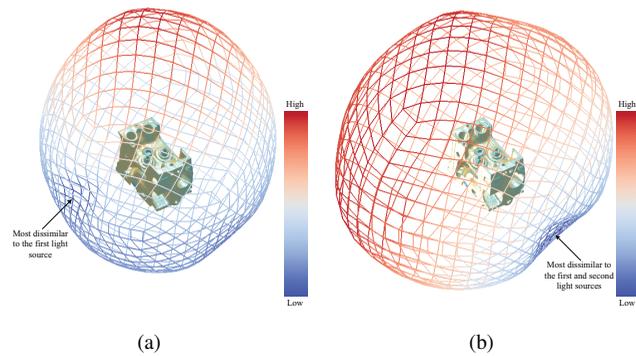


Fig. 9. Lighting similarity analysis for the engine data set. The lighting-similarity sphere is deformed and colored similar to the deformed light-position sphere. Blue light positions are more dissimilar to the previous selected light source(s) than red ones. (a) The lighting-similarity sphere compared to the optimal light source, as shown in Fig. 5(a). (b) The lighting-similarity sphere compared to the first two light sources.

each pair was compared by 9-23 users.

Fig. 10(b) shows, for each pair and each data set, the proportion of the tests in which the image with the higher MS-HWSD value was selected. For 35 of the 40 pairs (87.5%), the image with the higher MS-HWSD value was chosen in more than 50% of the selections; for 23 (more than half) of the pairs, the selection rate was greater than or equal to 80%. Overall, the user preferences were consistent with the MS-HWSD comparison in 78.9% of the pairs (521 comparisons).

The lowest percentage achieved was 33.3%, for the 15th pair of the engine data set. In this pair, the image with the overall higher MS-HWSD value had highlights in some regions, and it was found that users generally do not like highlights. However, with minor exceptions, the MS-HWSD proved itself a good lighting quality metric for measuring the enhancement of structural information.

Another user study was conducted to find out if the proposed approach identifies a better optimal light source in terms of structural information than other methods, from the user's point of view. Users viewed test sets and were asked to select which image reveals the object's structural information best. Lighting entropy, multi-scale entropy and the MS-HWSD were compared using the eight sets of images shown in Figs. 4, 7 and 8, which include seven with a single light source and one with two light sources. The test sets were presented in a random order, and the three images in a test set were also displayed randomly. The images were analyzed by 36 users, some of whom had previous knowledge of volume rendering.

Table 1 shows the statistics from the data collected. Users generally preferred the lighting parameters selected by the MS-HWSD over those from lighting entropy and multi-scale entropy. For a single light source, the selection percentages of the MS-HWSD were highest in six of the seven test sets, the exception being the pollen. In the tomato test set, 66.67% preferred the MS-HWSD an absolute majority of the users. For multiple light sources (tomato2), the proposed approach based on the MS-HWSD and lighting similarity was again found to provide better lighting parameters than other two.

For the pollen data set, the multi-scale entropy performed better than the MS-HWSD; this may be accounted for by the fact that the light sources selected by the multi-scale entropy and MS-HWSD are similar and the users found it difficult to decide which is the best. For the engine, while the light source selected by the MS-HWSD reveals more detail information on the surface, the selection percentage of the MS-HWSD is not much higher than that of the multi-scale entropy. This may be because users often prefer clear global structures to clear local details.

The users were also asked about the reasons for their selections. One noticeable comment for the teapot data set was that illumination effects make the teapot lid clearer.

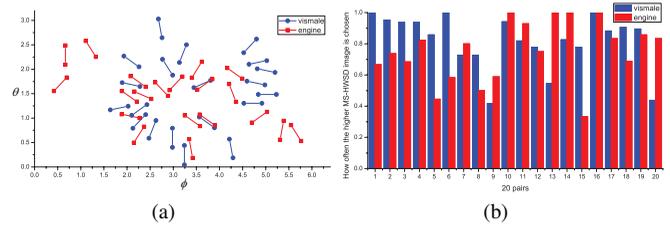


Fig. 10. (a) $\theta \times \phi$ distribution of the pairs of selected light positions for the vismale and engine data sets in our study (20 pairs per data set). (b) For each pair of images, the proportion of selections in which the image with the higher MS-HWSD value was chosen.

Table 1. User study statistics.

Test Sets	Selection Percentage (%)		
	Lighting Entropy	Multi-Scale Entropy	MS-HWSD
tooth	16.67	36.11	47.22
teapot	30.55	27.78	41.67
atom	22.22	33.33	44.44
pollen	11.11	47.22	41.67
engine	25.00	33.33	41.67
tomato	19.44	13.89	66.67
vismale	27.78	22.22	50.00
tomato2	30.55	27.78	41.67

Limitations. Since the definition of a good lighting configuration is highly subjective and depends greatly on the specific application, it is hard to generate absolutely optimal lighting parameters for all situations. This also explains the variations in the results from our user studies. For the MS-HWSD, we simply average all single-scale HWSD values. Although it worked well in our experiments, it would be better to use an image synthesis-based approach, similar to that used in MS-SSIM [28], to calibrate the weight of each scale, as the relative importance between different scales may be different.

As in lighting entropy and multi-scale entropy, the proposed MS-HWSD is an image-space-based metric. It may be beneficial to integrate object-space information for automatic lighting design, such as the importance of features, and domain knowledge may also be important for the lighting design in specific applications.

8 CONCLUSION

In this paper, we have presented a structure-aware automatic lighting design approach for volume visualization. The basic idea is to measure the structural differences between images with and without external illumination at different resolutions in order to quantify the quality of the lighting configurations. We proposed three lighting characteristics: lighting goodness, lighting similarity and lighting stability, on which to base automatic lighting design techniques. Our experiments, together with the user study, demonstrated the effectiveness of the proposed techniques, which release users from the intensive and laborious work needed to find the optimal lighting configuration. Since the proposed automatic lighting design searches only for the optimal lighting parameters for a fixed viewpoint, we plan to extend the proposed method to support interactive viewpoint changes. One possible solution is to optimize the pre-computed lighting parameters on the viewing sphere. In addition, we will extend the proposed lighting design techniques to time-varying volumes.

ACKNOWLEDGMENTS

The authors wish to thank all participants in the user study for their contributions to the evaluation, and the anonymous reviewers for their valuable comments. This work was partially supported by National Natural Science Foundation of China No. 60873122 and No. 60903133, and 863 Program Project 2012AA12A404.

REFERENCES

- [1] V. Blanz, M. J. Tarr, and H. H. Bühlhoff. What object attributes determine canonical views? *Perception*, 28(5):575–600, 1999.
- [2] U. D. Bordoloi and H.-W. Shen. View selection for volume rendering. In *Proceedings of IEEE Visualization '05*, pages 487–494, Los Alamitos, CA, USA, 2005. IEEE Computer Society.
- [3] F. Caniard and R. W. Fleming. Distortion in 3D shape estimation with changes in illumination. In *Proceedings of APGV '07*, pages 99–105, New York, NY, USA, 2007. ACM.
- [4] M.-Y. Chan, Y. Wu, and H. Qu. Quality enhancement of direct volume rendered images. In *Proceedings of Volume Graphics '07*, pages 25–32, 2007.
- [5] H. D. Cheng, M. Xue, and X. J. Shi. Contrast enhancement based on a novel homogeneity measurement. *Pattern Recognition*, 36(11):2687–2697, 2003.
- [6] H. A. Daivid. *The Method of Paired Comparison*. Halner Publishing, 1963.
- [7] K. Engel, M. Kraus, and T. Ertl. High-quality pre-integrated volume rendering using hardware-accelerated pixel shading. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware '01*, pages 9–16, New York, NY, USA, 2001. ACM.
- [8] R. W. Fleming, A. Torralba, and E. H. Adelson. Specular reflections and the perception of shape. *Journal of Vision*, 4(9):798–820, September 2004.
- [9] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622:759–771.
- [10] S. Gumhold. Maximum entropy light source placement. In *Proceedings of IEEE Visualization '02*, pages 275–282, Washington, DC, USA, 2002. IEEE Computer Society.
- [11] M. Halle and J. Meng. Lightkit: A lighting system for effective visualization. In *Proceedings of IEEE Visualization '03*, pages 363–370, Washington, DC, USA, 2003. IEEE Computer Society.
- [12] G.-F. Ji and H.-W. Shen. Dynamic view selection for time-varying volumes. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1109–1116, 2006. (Proc. Visualization'06).
- [13] J. K. Kawai, J. S. Painter, and M. F. Cohen. Radiooptimization: Goal based rendering. In *Proceedings of SIGGRAPH '93*, pages 147–154, New York, NY, USA, 1993. ACM.
- [14] W. B. Kerr and F. Pellacini. Toward evaluating lighting design interface paradigms for novice users. *ACM Trans. Graph.*, 28(3):26:1–26:9, 2009.
- [15] J. J. Koenderink and A. J. Doorn. The internal representation of solid shape with respect to vision. *Biological Cybernetics*, 32:211–216, 1979. 10.1007/BF00337644.
- [16] C. H. Lee, X. Hao, and A. Varshney. Light collages: Lighting design for effective visualization. In *Proceedings of IEEE Visualization '04*, pages 281–288, Washington, DC, USA, 2004. IEEE Computer Society.
- [17] A. Loza, L. Mihaylova, N. Canagarajah, and D. Bull. Structural similarity-based object tracking in video sequences. In *Information Fusion, 2006 9th International Conference on*, pages 1–6, July 2006.
- [18] J. Marks, B. Andalman, P. A. Beardsley, W. Freeman, S. Gibson, J. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Ruml, K. Ryall, J. Seims, and S. Shieber. Design galleries: A general approach to setting parameters for computer graphics and animation. In *Proceedings of SIGGRAPH '97*, pages 389–400, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.
- [19] G. Patow and X. Pueyo. A survey of inverse rendering problems. *Computer Graphics Forum*, 22(4):663–687, 2003.
- [20] F. Pellacini, F. Battaglia, R. K. Morley, and A. Finkelstein. Lighting with paint. *ACM Trans. Graph.*, 26(2):9:1–9:14, 2007.
- [21] P. Poulin and A. Fournier. Lights from highlights and shadows. In *Proceedings of I3D '92*, pages 31–38, New York, NY, USA, 1992. ACM.
- [22] C. Schoeneman, J. Dorsey, B. Smits, J. Arvo, and D. Greenberg. Painting with light. In *Proceedings of SIGGRAPH '93*, pages 143–146, New York, NY, USA, 1993. ACM.
- [23] A. Secord, J. Lu, A. Finkelstein, M. Singh, and A. Nealen. Perceptual models of viewpoint preference. *ACM Trans. Graph.*, 30(5):109:1–109:12, Oct. 2011.
- [24] R. Shacked and D. Lischinski. Automatic lighting design using a perceptual quality metric. *Computer Graphics Forum*, 20(3):215–226, 2001.
- [25] S. Takahashi, I. Fujishiro, Y. Takeshima, and T. Nishita. A feature-driven approach to locating optimal viewpoints for volume visualization. In *Proceedings of IEEE Visualization '05*, pages 495–502, Los Alamitos, CA, USA, 2005. IEEE Computer Society.
- [26] P.-P. Vázquez. Automatic lighting source placement for maximum visual information recovery. *Computer Graphics Forum*, 26(2):143–156, 2007.
- [27] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [28] Z. Wang, E. P. Simoncelli, and A. C. Bovik. Multi-scale structural similarity for image quality assessment. In *Proceedings of 37th IEEE Asilomar Conference on Signals, Systems and Computers*, pages 1398–1402, 2003.
- [29] Z. Zheng, N. Ahmed, and K. Mueller. iView: A feature clustering framework for suggesting informative views in volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17:1959–1968, 2011.