Perception-Based Illumination Information Measurement and Light Source Placement

Pere-Pau Vázquez† and Mateu Sbert‡

†Dept. LSI - Universitat Politècnica de Catalunya,
Campus Sud - Ed. ETSEIB, Diagonal 647, 8ena Planta, E - 08028 Barcelona, Spain.
ppau@lsi.upc.es
‡IIiA, Universitat de Girona
Campus Montilivi, EPS, E-17071 Girona, Spain
mateu@ima.udg.es

Abstract. The automatic selection of good viewing parameters is very complex. In most cases, the notion of *good* strongly depends on the concrete application. Moreover, when an intuitive definition of good view is available, it is often difficult to establish a measure that brings it to the practice. Commonly, two kind of viewing parameters must be set: the position and orientation of the camera, and the ones relative to light sources. The first ones will determine how much of the geometry can be captured and the latter will influence on how much of it is revealed (i. e. illuminated) to the user. In this paper we will define a metric to calculate the amount of information relative to an object that is communicated to the user given a fixed camera position. This measure is based on an information-based concept, the Shannon entropy, and will be applied to the problem of automatic selection of light positions in order to adequately illuminate an object.

1 Introduction

The selection of the adequate viewing parameters is a very complicated problem that is usually solved through a large process of test and error that requires long-time human dedication. Obviously, the necessity of human intervention becomes impractical for large collections of models. Moreover, many applications from scientific visualization are often employed by users with little or null experience in Computer Graphics. It is desirable then to find a way to automatically setting adequate viewing parameters. Recently, research on the automatic placement of cameras has attracted the attention of the Computer Graphics community, but only few papers focus on the important problem of correctly illuminating a scene. See

Fig. 1. Good and bad lighting of a chess queen.

for instance Figure 1, where a good and a bad illumination of a chess queen are shown. Our aim in this paper is to present a perception-based metric that evaluates the amount of illumination information contained in a view and we apply it for the automatic selection of light positions that adequately illuminate an object or scene.

The rest of the paper is organized as follows. In Section 2 we review previous work on parameter tweaking and the related work of automatic camera placement, Section 3 presents our perception-based measure of illumination information. In Section 4 a method for the automatic positioning of a single light source is developed. We also show how the extension to several light sources is straightforward. Finally, Section 5 concludes our work and points out some lines for future work.

2 Previous Work

The selection of viewing parameters for scene rendering is a complex and tedious process. Parameter tweaking can be divided into two phases: a) Camera position and direction setting, and b) Light source selection and positioning. The problem of good camera positioning has become an active field of research mainly due to the emergence of the so-called Image-Based Rendering methods. Light source selection has usually been treated from the point of view of inverse lighting. Now we proceed to review the previous work on these two areas of research.

2.1 Camera Placement

Colin [1] presents a method to select a good view to observe a scene modeled with an octree. Kamada and Kawai [2] define a criterion for the quality of a view for orthogonal projections. Plemenos and Benayada [3] extend Kamada's definition, considering the amount of detail shown in a view as the number of visible faces. Barral *et al.* [4] present a method for the automatic exploration of objects or scenes. In this case, the quality of a view is computed by defining a new importance function that depends on the visible pixels of each polygon.

Vázquez *et al* [5] have presented a new measure based on the Shannon entropy [6], *viewpoint entropy*, to evaluate the amount of *geometric* information seen from a point. It has been successfully applied to some scientific visualization problems such as automatic selection of good views of molecular models [7]. Takeuchi and Onishi [8] measure the entropy of an image based on histograms of intensities in order to find the complex parts of a scene.

2.2 Light Source Selection

Adequate lighting selection research can be divided in two subfields: inverse lighting, and maximum information communication. Here we review recent work on these fields.

In inverse lighting, the user specifies how the scene should look like and the adequate parameters are searched. Therefore it is assumed that the user has a knowledge on the object shape and material properties. Although there is a broad bibliograpy in inverse lighting (for a survey see [9]), we only cite here some examples. Schoeneman *et al.* [10] describe an interactive system that, given a set of lights with fixed positions, determines their colours and intensities in order to match a target image painted by the designer. Kawai *et al.* [11] control light emissions and directions, as well as surface reflectances for designing the illumination in an environment rendered with a radiosity based method. As in the former case, the user has to specify how the final image should be illuminated. None of these methods automatically sets the light source positions. Costa *et al.* [12] have implemented an automatic method of light placement and intensity selection. Their objective is to obtain a configuration that determines a given radiance distribution. Although it is a powerful approach, the objective function needs to be specified by the user with a scripting language, and therefore is not easily applicable when the user is not expert. Poulin and Fournier [13] and Poulin *et al.* [14] manipulate highlights and shadows in order to define a resulting illumination. These modifications are translated to the corresponding changes to the positions of the light sources. Jolivet *et al.* [15] present a Monte-Carlo method for the selection of light positions in direct lighting. They use a declarative paradigm in order to help the users to describe in a more intuitive way the lighting goal.

Opposite to inverse lightning problems, some methods seek the adequate light sources configuration that reveals the maximum of information to the user, by means of adequately placing light sources, no matter which object or scene is inspected. The Design Galleries^{TM} (DG) system is a method to automatically set parameters for computer graphics and animation. They automatically compute and organize sets of views or animations which are perceptually different from each other. The resulting images are presented to the user to choose among them. Apart from some parameters concerning to material properties, they also study light selection and placement [16]. Gumhold has also explored the problem of automatic parameter setting [17]. He has presented a method for the automatic light source placement which also uses an entropy-based function, the *lighting entropy*. He defines the unit of information $-\log p_i$ as a function of the measured brightness of the visible pixels. The brightness of a pixel is computed as the Y tristimulus value of the CIE 1931 standarized colour model. It is calculated with the following formula: $Y = 0.21262 \cdot R + 0.71514 \cdot G + 0.07215 \cdot B$. Then, the lighting entropy is defined as $H(X) = -\sum_{i=1}^{m} p_i \log p_i$, where the probability p_i is defined as the number of pixels whose brightness falls into interval i (the logarithms are taken in base 2 and $0 \log 0 = 0$ for continuity). The author defines the unit of information $-\log p_i$ as the number of pixels that fall into interval i where index i is computed as $i = \lceil m \cdot (Y + \frac{1}{2}) \rceil$. Therefore, their measure is maximum when the number of different brightness values in an interval is uniform across the scene. The number m of intervals chosen by the author is 30. The *lighting entropy* is similar to the *viewpoint entropy* presented by Vázquez *et al.* [5]. The latter uses as unit of information the relative projected area of each face, while the former use the normalized number of pixels that have the same brightness (in an interval). Surprisingly, the tests with users revealed that the best views as selected by his method were discarded by them due to the fact that they presented too large specular regions. With the results of the study he has improved the method taking into account their comments. Some fast methods for light positioning are also presented.

Shacked and Lischinski [18] propose a quality function formed by six terms that are weighted by the user. Their system optimizes these parameters based on a perceptual quality metric. Their objective is to effectively communicating information on the scene: shapes, materials, and their relationships. The quality metric they build is composed by six factors. Each of them is devoted to a different kind of information (such as edge detection or variance reduction), and some of them may have contradictory effects such as the histogram equalization term that, when applied, tends to increase the variance (which is controlled by the variance reduction term). Therefore, manual calibration is necessary for every scene whereas they have empirically found some weights that perform well for a certain number of models.

3 Illumination Information Measurement

3.1 Introduction

From the previous work, only Gumhold [17] has defined a fully automatic method for lighting parameter tweaking which needs no user intervention. The probability distribution of the entropy measure presented is based on the normalized number of pixels whose brightness falls in an interval. Unfortunately, such a distribution function may cause some problems. First, the use of brightness values does hide the information concerning to the colour that is perceived by the human visual system. Consequently, some colours that appear different to us are measured as the same one (see Figure 2*a*). If only a single material and a single light source (or several light sources with the same emission colour) is present in the scene, this would not be important. However, scenes with two differently coloured light sources might yield contradictory results, as an illumination which is best according with this method may be communicating a low amount of information than another one. This is exactly what happens in Figure 2*b*. The left cone is illuminated with a pink and a green light sources while the right one is is illuminated with a blue and a green light sources. When displayed, our visual system distinguishes easier between green and pink than between blue and green, and therefore the left cone communicates a higher amount of information. As blue and green on the right cone have similar brightnesses, with Gumhold's method the right cone would be chosen. Second, the normalization by the number of projected pixels does produce a lateral effect: a scaling of a model, under the same lighting conditions does yield the same entropy (see the Section 3.2).

We propose in this paper a new metric for illumination information measurement. It differs from previous approaches in that we measure the perceptual-based spectrum of LUV-colours. Moreover, we take into account spatial issues, that is, all the pixels with the same LUV colour will accumulate if they form a single continuous region (see Section 3.2), as otherwise they are perceived as separate sources of information by an observer.

(a) (b)

Fig. 2. In (a) we can see two different colours which have the same Y value. (b) shows two different illuminations of a cone. The first one, with a pink and a green light sources provides a higher amount of information because those colours are more easily distinguisible for us than blue and green in the right image. However, as in the right cone blue and green have similar brightnesses and therefore some regions that appear different to us, are summed together, this one would be chosen by Gumhold's system.

3.2 A Perception Based Measurement of the Illumination Information

(a) $H_L = 0.08865$ (b) $H_L = 0.06389$

Fig. 3. In (a) we can see a zoom-in of (b). If the background was not used, both would yield roughly the same quality value. With illumination entropy we have a noticeable difference.

An accurate positioning of a light source must reveal as much information as possible of a scene. The problem is how to define a formula that indicates the quantity of information present in an image. Vázquez *et al.* [5] have developed an entropy based measure that can be used to determine the amount of geometric information captured from a point. This value is calculated by projecting the scene onto a bounding sphere of the viewpoint and using as distribution function the relative projected areas of each face. Therefore its value is maximum when all the faces are seen with the same relative projected area. Although this effectively solves the problem of measuring geometric information, when we render a scene the current shading may hide some faces of the object. Thus, to measure the illumination information that arrives to a point, it is necessary to build a distribution function based on the visual stimulus that arrive at this point. Moreover, as the human visual system is limited, the measure must be sensitive to this and only take into account perceptible information. With these conditions we can build an entropy by using as unit of information the relative area of each region *whose colour is different from its surrounding*. The expression of this measure will be:

$$H_L(X) = -\sum_{i=0}^{m} A_i \log A_i,$$

where m is the number of different regions, A_i is the relative area of the region i.

The background is represented by region index 0. Remark that it is important to take the background into account because the probability distribution must add up to 1. Some approaches do not use the background information and instead they normalize the probabilities dividing by the number of valid pixels (see Gumhold [17]). But this normalization may hide some information, as a zooming-in of a certain view (provided that the object still remains inside the viewing frustum) will give the same value despite we are really seeing the object better and therefore this fact should be detected (see Figure 3).

Observe that the background plays another important role. Usually, the background has a different colour than the rest of the object, nevertheless, if the scene does not contain an ambient term, it might happen that the object is completely black under a certain light position. If the background is black, the information present is zero, as we are not seeing anything. On the other hand, if the background is not black, it will help us to see the silhouette of the object, and this must be somehow taken into account. In our examples we considered the background white and therefore the silhouette of a completely dark object is not zero. In any case, the colour of the background must be taken into account when measuring the illumination entropy, as it might be necessary to add its area if some parts of the scene have the same colour in any lighting configuration. Another important problem is the colour regions measurement. Gumhold measures the lighting entropy by adding the number of pixels with the same (in an interval) brightness. However, the distribution of the luminances across the scene may also be informative. So, instead of accumulating the relative area of the pixels which have a colour that can not be distinguished by the human visual system, our method takes as information unit, the relative area of each isolated region of the same colour.

The colours are transformed to CIE LUV format. In order to detect if two neighboring colours are the same, we use the CIE LUV colour difference formula: $\Delta E^* = \sqrt{(\Delta L^*)^2 + (\Delta u^*)^2 + (\Delta v^*)^2}$. It is known that colours that have a CIE LUV difference of less than 1 appear to be the same [19], so this is the criterion we have used to distinguish between regions. For a real complex environment, the difference may raise to 6 or even more, but we are dealing here with a simpler scenario, a window in a computer screen. Note that with this criterion a region labeled with the same colour could have some pixels that compared to each other yield a difference value of more than one. In order to avoid this, when a colour is compared to a neighbor to determine if it belongs to the same region than the previously computed, the current colour is assigned the colour of the neighbor to which it has been compared, if its difference is less than one. This would prevent a very smooth colour gradient traversing a big scene to be labeled as a single region. The use of such a probability function avoids the problem of mixing perceptible different colours in a same measure, as it happens in Figure 2. Gumhold's method would select right cone (lighting entropy of 0.1229 versus 0.1216) while our method yields a smaller value for the right cone (illumination entropy of 0.1127 versus a value of 0.1625 for the left one).

The neighborhood criterion that has been applied is the following one: Two pixels are neighbors and therefore belong to the same region if their colour difference is below one and if these pixels have an edge in common. With this method a typical checkerboard texture will not be considered as two regions but at a higher number of them, depending on the number of squares (we consider here a situation of constant shading along the texture, otherwise, the illumination will also introduce a higher number of different regions). In Figure 4 we can see this with an example. Figure 4a has a higher illumination entropy than Figure 4b, which corresponds to our perception of four regions in the first case and two in the second.

(a) $H_L = 0.1423$ (b) $H_L = 0.1308$

Fig. 4. (a) shows a checkerboard of four squares. Note that we intuitively identify four squares, while in (b) only two regions are perceived. Therefore, (a) is more informative than (b).

4 Perception Based Automatic Light Source Placement

The selection of the best position for a single light source was fully implemented as a brute-force algorithm. We place the light source at a set of different positions on a bounding sphere of the object and measure its illumination entropy. The position with the highest quality is selected. This method is general and can be applied whichever the shading algorithm is used, as the calculations are performed on the resulting image. In our case we have used OpenGL's default lighting mode. The highest cost of our method is incurred by the rendering tool and the capture of the data. Each time a new illumination has to be analyzed, the scene must be rendered and the image must be read back to main memory. However, several accelerations can be added, depending on the rendering tool. Gumhold [17] has presented a fast lighting scheme suitable for an OpenGL lighting, but other techniques can be applied if the rendering is not OpenGL based, as the quality criterion works for any kind of shading. In order to accelerate the computation we reduce the size of the window read back to main memory by reading the depth buffer at the beginning and inspecting it in order to obtain the bounding box of the object projection. Moreover, the lighting space can be restricted to the hemisphere where the user is placed as the lights placed on the other hemisphere will not illuminate most of the polygons facing the camera. This allows to obtain a proper light position in two to ten seconds for scenes of several thousands of polygons. We have also implemented a software renderer which uses Phong lighting. Initially, the object is projected and the depth buffer is read. Then, the illumination calculations are performed only on these captured pixels. This results in a reduction of 4:1 in computation time.

As seen in [18], it is difficult to visually evaluate the quality of a rendering. On the other hand, the optizimed images with their method also yield high values of illumination entropy with our system, with the advantage that we do not need a preprocess of calibration and different tuning for each model or view. In Figure 5 we can see some examples of adequately illuminated objects.

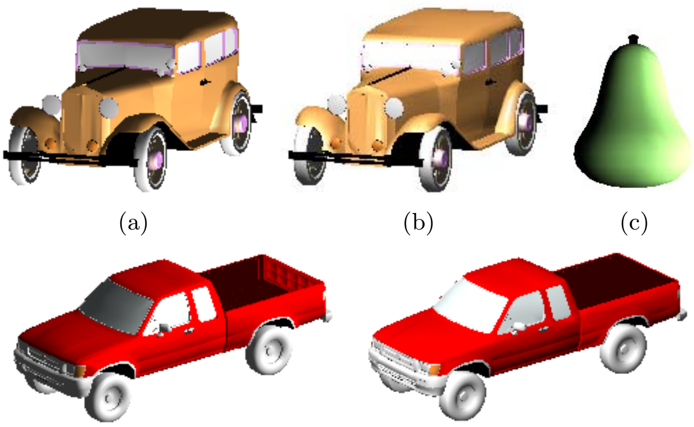

(a) (b) (c)

(d) (e)

Fig. 5. Comparison of our method with Gumhold's system. (a) and (d) were generated with our method, and (b) and (e) with Gumhold's measure. Notice the highly specular regions that hide some shape detail (note the back of the car in (d) and (e)), as reported by authors. (c) shows the optimal configuration of two light sources for a pear.

4.1 Discussion

(a) $H_L = 0.6117$ (b) $H_L = 0.5979$

Fig. 6. Two glossy spheres rendered using Phong prefiltered environment maps with exponents of 100 and 25. The higher the exponent the closer to totally specular, and the higher the entropy.

The measure we have presented here tries to maximize the information revealed to the user given an illumination of the scene. Its value is the highest when all the regions of the same colour are of the same relative area. In contrast to this, large regions of the same colour give low entropy, because of the logarithmic nature of the entropy measure. This is intuitively correct as a large flat polygon will be perceived better if there is a gradient in the illumination that shades it. The six quality terms employed by Shacked and Lischinski focus similar goals [18]. For instance, the histogram term seeks to equalize the amount of quantities of each luminance value appearing on the scene, although without taking into consideration the part of the scene where they appear. Hence, it is not surprising to discover that the images the authors select as the best ones are often the same that with our method.

On the other hand, our measure has several advantages. It is compact and general and thus, it must not be manually calibrated for every scene, because it gives a quantity that can be compared with successive renderings of the same scene. As it is a measure of the amount of information present in an image, it can also be used to compare between similar images. See for instance Figure 6. It shows two parabolic environment maps (actually only one of their paraboloid parts). The one on the left is filtered with an exponent of 100 while the second is been filtered with an exponent of 100. They are used to simulate different glossy objects [20]. Intuitively, we can see that the amount of information present on the second one is lower than in the former, but how different, we do not know. With our method we can measure the amount of information of both maps and prove that the result corresponds with intuition.

4.2 Several Light Sources

We have tested the selection of a good illumination with more than one light source. The extension of our algorithm to several light sources is straightforward. It has the disadvantage that it scales badly with the number of light sources. The bottleneck is at the rendering process and reading the generated information back to main memory. An increase in the number of light sources increases exponentially the number of views that have to be analyzed. With a single light source, it requires several seconds to find the adequate position. Therefore, for the case of several light sources, it becomes necessary to accelerate it by means of a global optimization method or an fast adaptive strategy. In Figure 5 right we can see an example of a pear optimally illuminated by two light sources.

5 Conclusions and Future Work

In this paper a perceptual based measure of the illumination information of a view has been developed. It is simple and robust and has a mathematical foundation on Information Theory. An extension to a set of light sources is also presented. Although it does not scale well with the number of light sources as the number of positions to analyze grow exponentially, some accelerations can be foreseen. In particular as the bottleneck is on the rendering process, and the viewpoint is fixed, a software renderer of the visible region reduces the computation to a 25%. In the future we will do a deeper research on this problem. Moreover, a study with trained and not trained users is necessary to evaluate the suitability of our method for applications such as scientific visualization.

Acknowledgments

This work has been supported by TIC-2001-2416-C03-01 of the Spanish government, and SGR2001-00296 and ACI2002-29 grants from Catalan Government.

References

1. C. Colin. Automatic computation of a scene's good views. In *Proc. MICAD*, February 1990.
2. T. Kamada and S. Kawai. A simple method for computing general position in displaying three-dimensional objects. *Computer Vision, Graphics, and Image Processing*, 41(1):43–56, January 1988.
3. D. Plemenos and M. Benayada. Intelligent display in scene modeling. new techniques to automatically compute good views. In *Proc. International Conference GRAPHICON'96*, July 1996.

4. P. Barral, G. Dorme, and D. Plemenos. Scene understanding techniques using a virtual camera. In A. de Sousa and J.C. Torres, editors, *Proc. Eurographics'00, short presentations*, 2000.
5. P.-P. Vázquez, M. Feixas, M. Sbert, and W. Heidrich. Viewpoint selection using viewpoint entropy. In T. Ertl, B. Girod, G. Greiner H. Niemann, and H.-P. Seidel, editors, *Proceedings of the Vision Modeling and Visualization Conference (VMV-01)*, pages 273–280, Stuttgart, November 21–23 2001. IOS Press, Amsterdam.
6. R. E. Blahut. *Principles and Practice of Information Theory.* Addison-Wesley, Cambridge, MA, 1987.
7. P.-P. Vázquez, M. Feixas, M. Sbert, and A. Llobet. Viewpoint entropy: A new tool for obtaining good views for molecules. In D.Ebert, P.Brunet, and I.Navazo, editors, *Data Visualisation 2002 (Eurographics/IEEE TCVG Symposium Proceedings).* Eurographics/IEEE, May 27-29 2002.
8. Y. Takeuchi and N. Ohnishi. Active vision system based on information theory. *Systems and Computers in Japan*, 29(11):31–39, 1998.
9. G.Patow and X.Pueyo. A survey on inverse rendering problems. Technical Report IIiA 00-07-RR, IIiA, University of Girona, 2000.
10. C. Schoeneman, J. Dorsey, B. Smits, J. Arvo, and D. Greenberg. Painting with light. In James T. Kajiya, editor, *SIGGRAPH 93 Conference Proceedings*, Computer Graphics Proceedings, Annual Conference Series, pages 143–146. ACM SIGGRAPH, ACM Press, August 1993.
11. J. K. Kawai, J. S. Painter, and M. F. Cohen. Radioptimization - goal based rendering. In *Proc. of SIGGRAPH-93: Computer Graphics*, pages 147–154, Anaheim, CA, 1993.
12. A. Cardoso Costa, A. Augusto de Sousa, and F. Nunus Ferreira. Lighting design: A goal based approach using optimization. In *Rendering Techniques '99*, pages 317–328, New York, NY, 1999. Springer Wien.
13. P. Poulin and A. Fournier. Lights from highlights and shadows. *Computer Graphics*, 25(2):31–38, March 1992.
14. P. Poulin, K. Ratib, and M. Jacques. Sketching shadows and highlights to position lights. In *Proceedings of Computer Graphics International 97*, pages 56–63. IEEE Computer Society, June 1997.
15. V. Jolivet, D. Plemenos, and P. Poulingeas. Inverse direct lighting with a Monte Carlo method and declarative modelling. 2330:3–??, 2002.
16. 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 T. Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 389–400. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7.
17. S. Gumhold. Maximum entropy light source placement. In *Proc. of the Visualization 2002 Conference*, pages 275–282. IEEE Computer Society Press, October 2002.
18. R. Shacked and D. Lischinski. Automatic lighting design using a perceptual quality metric. In *Computer Graphics Forum (Proceedings of Eurographics 2001)*, volume 20, pages C–215–C–226, September 2001.
19. L. Neumann, K. Matkovic, and W. Purgathofer. Perception based color image difference. *Computer Graphics Forum*, 17(3):233–241, September 1998.
20. J. Kautz, P.-P. Vázquez, W. Heidrich, and H.-P. Seidel. A Unified Approach to Prefiltered Environment Maps. In B. Peroche and H. Rushmeier, editors, *Eleventh Eurographics Rendering Workshop 2000*, pages 185–196. Eurographics, June 2000.