

# An Experimental Study of Light Source Determination for Computer Graphics \*

Yee-Hong Yang

The Scene Analysis and Modelling Group  
Computer Vision and Graphics Laboratory  
Department of Computer Science, University of Saskatchewan  
Saskatoon, Saskatchewan, Canada S7N 5A9

email: yang@cs.usask.ca

phone: (306) 966-4891

fax: (306) 966-4884

March 6, 1998

## Abstract

In combining synthetic and real imagery, the determination of light sources is an important step. In this paper, various issues of determining multiple light sources for computer graphics are explored. Two main approaches, the least squares minimization and the non-negative least squares minimization methods, are studied experimentally. The first approach was found to produce unsatisfactory results. Hence, the focus of this paper is on the second approach.

From the computer vision point of view, irrespective of the methods, the exact direction of light sources cannot be recovered. However, from the computer graphics point of view, the results are useful enough for regenerating an approximate shading.

Several placement strategies of initial candidates light sources are proposed and studied. The feasibility of the proposed algorithms in reconstructing the shading of an object in a real scene is also demonstrated.

*Keywords:* computer graphics, light source determination, realistic rendering.

## 1 Introduction

The light source determination problem can be formulated as follows:

$$E = \sum_{i=1}^N \lambda_i \mathbf{n} \cdot \mathbf{l}_i \quad (1)$$

where  $E$  is the observed intensity,  $\mathbf{n}$  is the surface normal at the point of observation,  $\lambda_i$  is the intensity and,  $\mathbf{l}_i$  is the unit direction vector of the  $i$ th source. Both  $\lambda_i$  and  $\mathbf{l}_i$  are unknown. If the direction of the sources are given, however, then it appears relatively simple to compute  $\lambda_i$  using the least-squares optimization technique.

The earliest work to determine the direction of a single light source using a Lambertian surface can be traced to a simple technique proposed by Pentland [3]. In Pentland's approach, the slant and the tilt are determined as follows:

$$\sigma_l = \text{acos} \left( \frac{1 - E^\Omega(I_x)^2 + E^\Omega(I_y)^2}{\text{var}^\Omega(I_s)^2} \right)^{\frac{1}{2}} \quad (2)$$

$$\tau_l = \text{atan} \frac{E^\Omega(I_y)}{E^\Omega(I_x)} \quad (3)$$

where  $E^\Omega$  denotes the expected value,  $I_x$  and  $I_y$  denote the partial derivatives of intensity with respect to  $x$  and  $y$ , respectively. Recently, Chojnacki et al. [1] have shown that the slant estimation is incorrect. The reason is because the above slant estimation diverges when pixels close to the boundary are used. Chojnacki et al. propose a revised

\*The author would like to acknowledge Alias|Wavefront and Natural Sciences and Research Council of Canada for financial support through grant number OGP 0000370.

version of the above expressions which can accommodate of the divergence problem. The approach proposed by Yang and Yuille [4] addresses the issue of multiple sources. The occluding contour is analyzed for maxima. Each separable source will appear as a local maximum. Unfortunately, such an approach is difficult to be applied in a practical environment because it depends critically on the requirement that the sources are to be separable. Hougen and Ahuja [2] approach the problem from the optimization point of view. Rather than determining the direction, they suggest fixing the direction of each source. Hence, equation 1 can be rewritten as:

$$\begin{pmatrix} E^1 \\ E^2 \\ \vdots \\ E^M \end{pmatrix} = \begin{pmatrix} \mathbf{n}^1 \cdot \mathbf{l}_1 & \cdot & \cdot & \cdot & \mathbf{n}^1 \cdot \mathbf{l}_N \\ \mathbf{n}^2 \cdot \mathbf{l}_1 & \cdot & \cdot & \cdot & \mathbf{n}^2 \cdot \mathbf{l}_N \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \mathbf{n}^M \cdot \mathbf{l}_1 & \cdot & \cdot & \cdot & \mathbf{n}^M \cdot \mathbf{l}_N \end{pmatrix} \cdot \begin{pmatrix} \lambda_1 \\ \cdot \\ \cdot \\ \cdot \\ \lambda_N \end{pmatrix}$$

If  $M > N$ , which is an easy to achieve condition, optimization techniques can be applied to compute the intensities,  $\lambda_i$  where ( $i = 1, \dots, N$ ). Both the least squares minimization and the non-negative least squares minimization techniques have been suggested [2]. In the least squares minimization approach, the resulting intensities can be less than 0, in which case, sources with negative intensities are called virtual sources. However, in the non-negative least squares minimization approach, the results are always larger than 0. Their experimental results suggest that the virtual source method performs better than the non-negative source method. However, contrary to findings as reported in [2], experiments performed by the author suggest that the least squares minimization approach does not produce usable results. Hence, in this paper, only results using the non-negative least squares approach are presented.

The organization of this paper is as follows. To gain an understanding of the behaviour and shortcomings of the non-negative least squares minimization techniques, experimental results in using the non-negative least squares are described in Section 2. In addition, initial placement and refinement strategies for the source locations using non-negative least squares are discussed in Section 3. The experimental results using the proposed approach strategies are presented in Section 4.

To enable the technique to be practical in real applications, a calibration object with known shape is required. In this paper, a spherical calibration




Figure 1: a) Original image, b) - d) reconstructed image using non-negative least squares. b) stepping interval = 10, error/pixel = 23.73, c) stepping interval = 50, error/pixel = 23.75, d) stepping interval = 90, error/pixel = 23.62.

object is used. The algorithm to locate such an object in an arbitrary scene is an interesting pattern recognition problem. A modified technique based on the Hough transform is employed. The experimental result in a real environment is presented in Section 5. The paper concludes in Section 6.


## 2 Empirical Findings of Non-Negative Least Squares

In this section, experimental results using non-negative least squares (NNLS) technique are discussed. In brief, NNLS incorporates the non-negative constraint for the intensity calculation. MatLab is used for the following experiments.


Because of the page limitations, only one set of images (sphere) are used. For each image, samples are taken at regular interval on a contour which is at a fixed distance from the boundary of the object. The sampling interval between samples varies from 10, to 90, in steps of 20. Only 49 fixed candidate light sources are used. The experimental results are shown in Figure 1. One can easily see that the reconstructed results are quite good.

### 2.1 Effects of Quantization

In general, as the number of candidate light sources goes higher, the resulting reconstructed image is



a) 65 predefined light sources.



b) 257 predefined light sources.

Figure 2: Effects of varying the number of predefined light sources. Circles denote original directions. Crosses denote recovered directions.

more accurate but it does not necessarily imply that the accuracy of the location of light sources is improved. In this experiment, the object (sphere) is illuminated with 10 randomly generated light sources.

The experimental results are shown in Figure 2. One can see that in all cases the directions of light sources are not recovered irrespective of the number of predetermined light sources.

### 3 The Proposed Approach

In the original NNLS approach [2], the locations of the light sources are fixed. The placement of the initial light sources was not considered. Furthermore, there was no discussion on refinements

or performance analysis. To simplify future discussions, these problems are, respectively, called: 1) the placement problem and 2) the refinement problem. The proposed solutions to these problems are studied experimentally.

One problem with the NNLS approach is the computational complexity. Experimentally, it was observed that the time increases drastically with the number of candidate sources. So, the placement of the initial set of sources is important. Since the placement cannot be exactly correct, refinement is required to improve the accuracy of the reconstruction results. Although the exact directions of the light sources cannot be determined, using the proposed refinement strategies, the error in the reconstructed images can be minimized.

The proposed approach consists of the following steps:

1. Place an initial set of light sources,
2. Solve the NNLS problem, i.e. determine the intensities of the candidate sources,
3. Threshold sources with intensities less than a predefined threshold, and
4. Refine the light source locations.

Steps 2 to 4 are iterated for a predefined number of times. In the current implementation, the threshold is selected to be 10. Contrary to solving the problem using the least-squares minimization approach, the intensities of a majority of light sources are zero using the NNLS approach. After thresholding, light sources with intensity value less than 10 will not participate in the refinement process.

#### 3.1 Placement of Candidate Light Sources

As indicated in the original paper [2] that it is desirable to place the sources uniformly. Unfortunately, on a sphere, it is found to be impossible. Since the initial placement will not be the final result, the problem one should ask instead is whether or not the initial placement has to be uniform. Another question one can raise is the definition of uniformity.

In the experiments, two different strategies are studied. The first strategy is called the square grid placement (see Fig. 3 a). In the figure, the projections of one of the ends of the light source vectors are shown. The other end is at the origin (not shown). It is apparent that on the projection plane, the placement is uniform. But with respect to the unit sphere, the placement is not uniform.

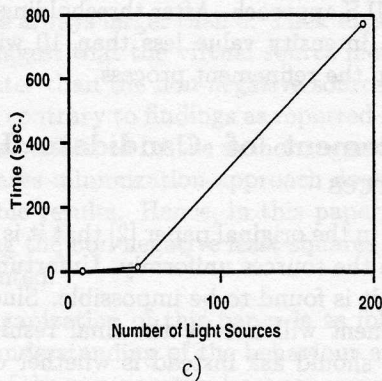




Figure 3: a) Uniform placement of light sources using a square grid. There are 16 light sources. b) Uniform angular placement of light sources on the unit sphere. c) Variation of computation time with the number of light sources. There are 17 light sources in this arrangement.

Another strategy is to place the sources uniformly with respect to the angular separation on the unit sphere (see Fig. 3 b)) and it is called the radial placement in this paper. Using radial placement, the angular separation between adjacent light sources maintains the same, but the projections of these light source vectors are not uniform. Furthermore, the solid angle subtended by four adjacent neighbours is dependent on their locations on the unit sphere. In the figure (Fig. 3 b), there is a singularity at the origin. Consequently, with respect to the solid angle measure, even on the unit sphere, the radial placement is not uniform.

In addition to the type of grid, the resolution of the grid has a significant impact to the computational speed as seen in Figure 3 c. It is apparent that by increasing the resolution, the time required increases drastically.

### 3.2 Refinements

Since the computational cost is high if the resolution of the initial placement is high. One may want to start with one which has a low resolution to get an initial guess and then refine the calculations. As a note, when the number of sources is over 200, the computation cannot be completed because of insufficient memory. This illustrates the significance of the proposed coarse-to-fine refinement strategies.

Three strategies are studied and to simplify future discussions, these refinements are, respectively, called: mesh, jittering, and shooting. Mesh involves generating a small local mesh of finer resolution (see Fig. 4a). Jittering involves moving a candidate light source by  $1/10$  the distance between the light source and the origin (see Fig. 4b). It is possible to implement a random displacement. Such an implementation is beyond the scope of this paper, however. The space we are using is the projection of the light source vector. Shooting involves generating a line of sources from the candidate to the origin (see Fig. 4c). Intuitively, both mesh and shooting refinements are the most expensive because they generate additional light sources. Each strategy has its advantages and disadvantages which are discussed in the following.

## 4 Experimental Results

### 4.1 Effects of Initial Placement Strategies

In the first set of experiments, the effects of the initial placement strategies are studied. The resolution




Figure 4: Different refinement strategies. a) mesh: around each candidate, a mesh of candidates are generated. b) jittering: the candidate is moved by  $1/10$  of the distance between the candidate and the origin. c) shooting: a line of candidates are generated from the original candidate.

is set at 10 for the case of the square placement and 16 for the case of the radial placement. The experimental results are shown in Figure 5. In the figure, one can observe the following:

- Using the square grid placement, the convergence is faster both in terms of the number of iteration (3 vs. 4) and the time (130 sec. vs. 350 sec.).
- The best error is lower using the square grid placement than that using the radial grid placement.

#### 4.2 Effects of Initial Resolution

The second set of experiments is to illustrate the effects of resolution for the square grid placement strategy. In this set of experiments, the mesh refinement strategy is employed. From Fig. 6, one can observe the following:

- Using a low resolution, although the initial computation time per iteration is low, the subsequent computation time per iteration increases.
- Using a low resolution takes more iterations as well as more time to reach the best solution than that uses a higher resolution initially.
- The best error is lower using a grid with a high initial resolution.

Based on the above observations, it suggests that one has to start with a reasonably good resolution to get a good result at the best possible time.

From the experimental results, it is noteworthy that using both resolutions, the final results are approximately the same in accuracy, which is a good indication that the final results are independent of

the initial placement or resolution. But using a finer resolution, the number of iterations required is lower than that of using a coarser resolution. The computation time required to arrive at the best solution is also lower. However, this result may be dependent on the object as well as on the number of sources.

#### 4.3 Effects of Refinement Strategies

In this set of experiments, the square grid with a resolution of 10 is used. The results of using the three refinement strategies are shown in Figure 7. Ignoring the first iteration, the following observations can be seen in the figure.

- The time per iteration is quite low as well as constant (2 sec.) using the jittering strategy.
- The shooting strategy appears to give the worst result in the reconstructed image.
- Both mesh and jittering give approximately the same result in the reconstructed image.
- The total time required using the jittering strategy is approximately 20 sec. whereas the mesh strategy is 90 sec.

Based on the above observations, the jittering strategy appears to be an attractive choice in terms of its speed as well as its accuracy.

### 5 Light Source Determination in a Real Environment

In the previous sections, the experimental results using synthetic objects demonstrate the feasibility of determining light sources for the purpose of rendering. Although the recovered light sources are not close to the original ones, the rendered image does look similar. In this section, the approach is applied to recover light sources in a real environment. However, before light sources can be recovered, a calibration object must first be located. In this paper, a sphere is used as the calibration object because its normals are known and also there are efficient algorithms to locate spherical objects, e.g. the Hough transform.

Figure 8 shows an input image.

Since the ball is a sphere, its normals are known. The algorithm to determine light sources are employed to compute the light source information. The computed light source information is then used to shade a synthetic sphere located at the location of the detected circle.

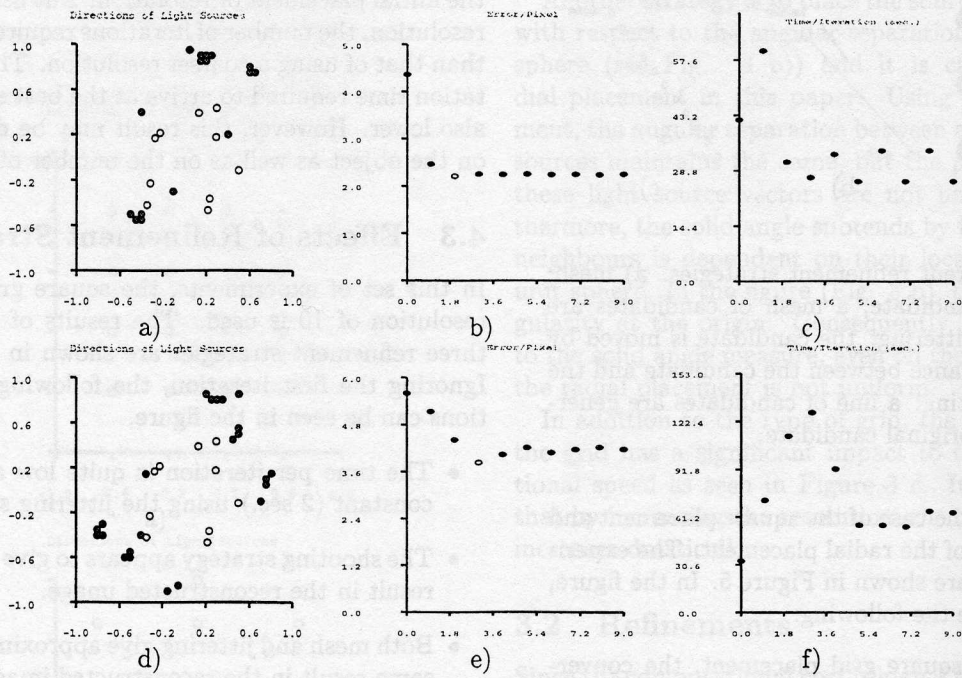


Figure 5: Effects of using initial placement strategies. a) and d) Computed light source locations. Open circles denote original directions. Closed circles denote computed directions. b) and e) Variation of error/pixel with respect to iteration. Open circle denotes the best result. c) and f) Time per iteration. First row: Square grid placement. Second row: Radial grid placement.

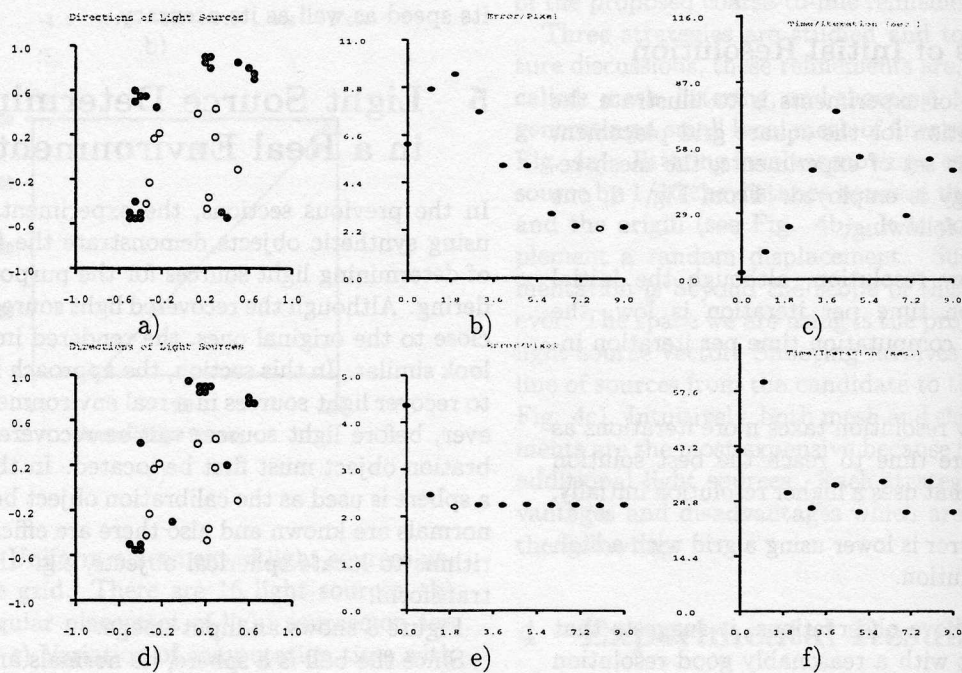


Figure 6: Effects of initial resolution. a) and d) Resulting light locations. Open circles denote original directions. Closed circles denote computed directions. b) and e) Variation of error/pixel with respect to iteration. c) and f) Computation time per iteration. First row: resolution = 5. Second row: resolution=10.




Figure 7: Effects of refinement strategies. a), d), and g) Resulting light locations. Open circles denote original directions. Closed circles denote computed directions. b), e), and h) Variation of error/pixel with respect to iteration. c), f), and i) Computation time per iteration. First row: mesh refinement. Second row: jittering refinement. Third row: shooting refinement.


After determining the light sources using the proposed approach, the synthesized result can be seen in Figure 8 b). The rendering is performed using a simple raytracing algorithm with supersampling.

## 6 Conclusions


In this paper, an algorithm to refine the placement of candidate light sources for computer graphics is proposed. The experimental results of a real scene with a calibration object is used to demonstrate the feasibility of the proposed approach. It was found that the shading of the regenerated object is a reasonable approximation to that of the original object. Future work will include the extension of the proposed approach to colour images as well as video sequences.

## References

- [1] W. Chojnacki, M.J. Brooks, , and D. Gibbins. Revisiting pentland's estimator of light source direction. *J. of Optical Society of America A*, 11(1):118-124, 1994.
- [2] D. H. Hougen and N. Ahuja. Estimation of the light source distribution and its use in integrated shape recovery from stereo and shading. In *IEEE 4th Int. Conf. on Computer Vision*, pages 148-155, Berlin, Germany, May 11-14, 1993.
- [3] A. Pentland. Finding the illuminant direction. *Journal of the Optical Society of America*, 74(4):448-455, April 1982.
- [4] Y. Yang and A. Yuille. Sources from shading. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 534-539, Lahaina, Maui, Hawaii, June 3-6, 1991.



a)



b)

Figure 8: a) Input image with the calibration object. b) Image with reconstructed synthesized object.