A Survey of
Light Source Detection Methods

Nathan Funk
University of Alberta
Mini-Project for CMPUT 603

November 30, 2003

Abstract

This paper provides an overview of the most prominent techniques for light source detection. The main goal of light source detection is to recover the location and intensity of light sources given a single image of a scene. The light source need not be visible in the image for most of the methods discussed in this paper.

Twelve methods are examined starting with Ullman's approach from 1975. Since then, a variety of techniques have been developed. Many of them are based on a number of assumptions that simplify the analysis. None of the current techniques are powerful enough to be used in any situation. For this reason a comparison of the approaches is also provided identifying the benefits of each individual method.

1 Introduction

Assuming you are given image of an arbitrary scene, how would you determine the origin of the light illuminating the objects in the image? How bright is each light source? These are the type of questions that light source detection techniques try to answer.

If all light sources are visible in the image, the problem may not be all too difficult. For example one could attempt finding the location of the sources by simply locating regions where the intensity is above a certain threshold.

But what if the light sources are not within the image region? Debevec [2] and others suggest making some of the light sources visible by placing a mirrored sphere in the scene. This can be very powerful when one has the option of placing an object in the scene, but unfortunately this is not always the case. There is undoubtedly a demand for techniques that can locate lights that are not visible in the image without manipulating the scene. All other methods discussed in this paper assist in dealing with this situation.

Light sources not being visible in the image is not the only challenge that needs to be dealt with. The second difficulty arises from the fact that it is often not known *what* is depicted in the scene. In other words, little or no information about the geometry, surface texture and other properties of the visible objects is available.

With all these unknowns, finding the exact locations and intensities of the sources is tremendously difficult. Hence, the term *illumination estimation* is a more appropriate description of the common techniques. Some techniques do not even attempt to recover multiple light sources, but instead give an approximate location of a single source. Even the methods that can detect multiple sources do not always do this high accuracy.

Since none of the current methods can reliably detect light sources in any arbitrary image, it is essential to know the benefits of the different approaches.

2 Background

The earliest attempt to detect light sources found in the research for this project was published in 1975 by Ullman [9]. However this approach is not a very general one since it only deals with light sources visible in the image. The first more general technique was developed by Pentland in 1982.

Since then, a variety of different approaches have been developed, each of which might be more applicable for a specific situation. This section will discuss the general applications and assumptions that are common to most techniques. Furthermore, the idea of how all techniques are associated with the field of computer vision is discussed.

2.1 Applications

Although location and intensity information of light sources might be used for a wide range of applications, only the two most important applications are discussed here.

Light source detection is tightly connected with the field of computer vision, where the one of the most important problems *scene reconstruction*. Scene reconstruction attempts to create a geometric model from a single image or multiple images of the same objects from different viewing positions.

Figure 1: Augmented vision results of Debevec [2]. The six spheres in this image are rendered using lighting information obtained from the real scene.

Many methods are available for this, and a large group of them is based on tracking *feature points* on the objects in a sequence of images. This can be very effective with textured objects where many such points are available.

However when an object's surface is smooth and has no texture, this type of an approach will likely fail. But *Shape from shading* techniques can be helpful in this situation. They approximate the geometry of an object from it's shading in the image. Of course, for this the lighting information is essential.

A second application is in *augmented vision*. An artificial object is inserted into a real scene, creating the illusion that it is actually part of the real scene. Without proper shading of the artificial object and shadow generation, anybody will immediately be able to see the inconsistencies. An excellent example of proper lighting is given in a paper by Debevec [2] and is shown in Figure 1.

2.2 Associations with Computer Vision

As mentioned in the applications, shape from shading methods require knowledge of the lighting conditions to be able to generate a model of the objects in the scene. These could be obtained from a light detection method, however it requires knowledge of the shape of the objects. So there is a definite problem in getting started with both of the necessary components.

A good way to approach this problem is through iteration. First an approximation of the lighting might be generated, then an estimation of the shape followed by a refinement of the lighting and so on. This type of approach was suggested by Brooks and Horn [1] in 1985.

2.3 Common Assumptions

Most of the methods make some basic assumptions about the scene to simplify the mathematics involved in the solution process. Some of the typical assumptions are:

- All light sources are directional.
- The object's surface reflection is lambertian.
- All surfaces are smooth.
- A specific object is shown in the image.
- The number of sources is known.

3 Techniques

The following are the most prominent techniques for light source detection. The techniques are listed in chronological order to provide an overview of how this field has developed over time.

The terms *tilt* and *slant* will be used to describe the position of the light sources in an angular coordinate system, similar to spherical coordinates. For the techniques of Pentland, Lee & Rosenfeld, and Brooks & Horn a copy of the original paper could not be obtained, so the summary of the method was based on Zhang's thesis [14] and references from other papers.

3.1 Ullman (1975)

This approach [9] is included mainly because it was one of the first attempts to solve the light source detection problem. It is not commonly referenced in light source detection papers. Unlike most future approaches, only light sources visible in the image can be detected with this method.

Ullman discusses why simplistic methods of finding light sources in an image such as finding the point of highest intensity are insufficient. He proposes a method that examines two adjacent patches in the image at a time, comparing the intensity-ratio to the gradient-ratio. If they are not equal, he claims one of the patches is a light source.

3.2 Pentland (1982)

Pentland's approach [8] to light source recovery is based on a statistical analysis of the image data. It assumes that the surface normals in the scene are not biased towards certain directions. For an image of a sphere this assumption is correct, and it will be approximately correct for a scene with many objects of random shape and orientation. But for shapes such as a cylinder the directions of the normals is strongly biased, making the results inaccurate.

The average intensity changes in the x and y directions (denoted as E_x and E_y) are used to determine the tilt and slant of a single directional light source. This is accomplished by first finding the least squares solution of

$$\begin{bmatrix} E_{d_0} \\ \vdots \\ E_{d_n} \end{bmatrix} = \begin{bmatrix} \cos\theta_0 & \sin\theta_0 \\ \vdots & \vdots \\ \cos\theta_n & \sin\theta_n \end{bmatrix} \begin{bmatrix} E_x \\ E_y \end{bmatrix}, \tag{1}$$

where E_d is the average intensity change in a specific direction θ. So each row represents the intensity change in a single direction. By sampling many directions, the values of E_x and E_y can be approximated, then finally the tilt is calculated using

$$\tau = \arctan(\frac{E_y}{E_x}) \pm \pi. \tag{2}$$

The slant is estimated from the E_x, E_y and E_d values as well. However this step is based on some assumptions that make the estimation quite inaccurate in most situations.

3.3 Lee and Rosenfeld (1985)

As another statistical technique similar to Pentland's, this approach [4] is targeted specifically at light source recovery from an image of a sphere. This is the first example in this paper of a method operating on a specific given geometry.

3.4 Brooks & Horn (1985)

In addition to recovering light source information, this iterative approach [1] also attempts to recover shape information. However it requires initial surface normal and light source estimates which in most cases are difficult to obtain.

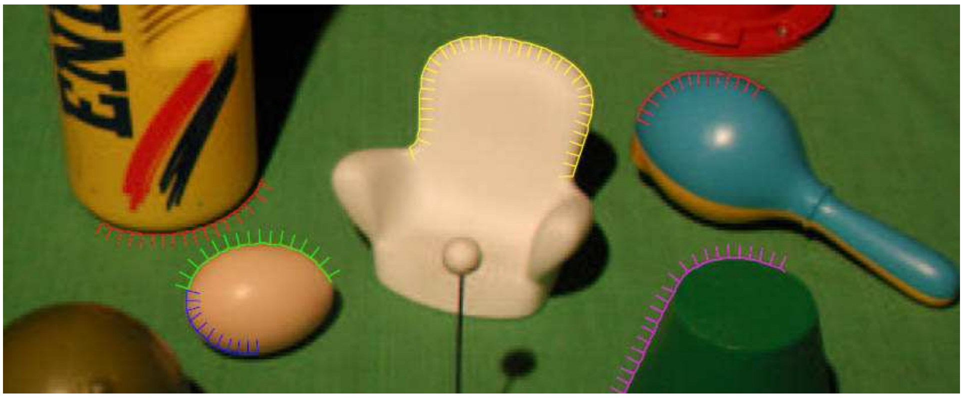

Figure 2: Highlighted occluding contours and normals [6]. Both Weinshall's and Yang & Yuille's approaches use these pieces of information in determining the location of light sources.

3.5 Weinshall (1990)

Similar to Brooks and Horn, Weinshall [12] proposed a combination of shape from shading and a light source detection method.

The light source detection method uses primarily the intensity information along occluding object boundaries to determine the tilt of the light. This is possible since the surface normals are known along the boundary. Figure 2 shows an example image with highlighted contours and normals. The tilt is estimated by searching for the points of maximum intensity along the boundary. Finally the slant is estimated by locating the brightest spot on the object.

3.6 Yang & Yuille (1991)

This is the first method [13] to attempt recovering multiple light sources. It also employs the analysis of intensity information along the occluding boundaries. At any point along the boundary, the surface normal can be obtained, and the intensity information is available. These two pieces of information can be used in the image irradiance equation to find the light source directions. It is not sufficient to only use a single point in this calculation, so a system of simultaneous equations is constructed from multiple points along the occluding boundary.

By solving these equations for the light source directions Yang & Yuille were able to recover up to three sources if they are separated far enough from each other.

3.7 Hougen & Ahuja (1993)

Another approach targeted for use in shape from shading methods is proposed by Hougen & Ahuja [3]. However it is based on initially acquiring depth information from stereo before estimating light sources. Predefined light source directions are used in a set of image irradiance equations. The light source intensities are found from this set of equations using a least squares technique.

3.8 Debevec (1998)

Debevec's motivation is different from the previous methods since he does not attempt to find position and intensity of distinct light sources, but rather the light field illuminating the area of interest. This is done by placing a mirrored sphere in the scene and obtaining high dynamic range images. The application this is applied to is augmented reality.

Okatani & Deguchi [7] also analyze this type of approach but look closer into how the image irradiance and the scene radiance are related. They also discuss how intensities are affected when an image is taken with a wide angle lens system (e.g. a fish-eye lens).

3.9 Zhang & Yang (2000)

This approach [14] [15] operates on an image of a sphere with a lambertian surface. Multiple sources can be recovered by locating cut-off curves on the sphere. The cut-off curves are great-circles around the sphere that separate an illuminated side from the dark side of the sphere for each individual light source.

Figure 3 shows the cut-off curves corresponding to a specific image. These curves are determined by searching for points along them, named *critical points*. The basic steps involved in determining the light sources in this technique are:

1. Find *critical points* on the sphere by searching along the visible sections of great circles, and comparing the local intensity properties to mathematical constraints.

2. Group the critical points into cut-off curves through use of the Hough transform. Each curve identifies the *pre-direction* of a single source.

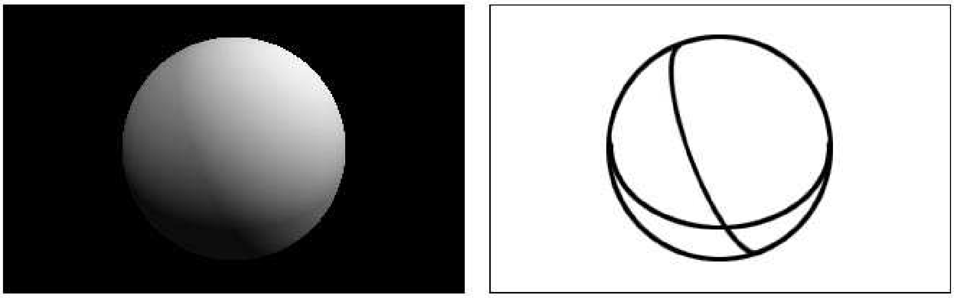

Figure 3: An example of the cut-off curves (right) corresponding to the image of a sphere (right) illuminated by two light sources. This type of image is analyzed in Zhang & Yang's approach to recover multiple light sources.

3. Determine the intensity (and direction) of each light source by finding the least squares solution to a set of irradiance equations.

Note that the differentiation between *pre-direction* and *direction* needs to be made since a light source can be located on either side of the plane that a cut-off curve lies in. The pre-direction only identifies the line normal to this plane, not the full direction.

3.10 Wei (2001)

This method [11] also recovers multiple light sources from the image of a sphere. The main difference from Zhang & Yang's approach is in the determination of the critical points. Rather than searching for critical points along great circles, a local light source constant constraint (LLSCC) is proposed. Under this constraint, local regions are analyzed and potentially categorized as containing candidate critical points. The final steps of grouping the critical points and determining the light source intensities are similar to Zhang & Yang's.

Wei claims that his approach is less computationally expensive and works with a more dense data set increasing the robustness of the technique. The method is tested on noisy images to prove the robustness.

3.11 Wang & Samaras (2002)

Similarly, this method also claims to outperform Zhang & Yang's approach. Although it is similar in many respects, it also extends the core ideas to allow analysis of arbitrary smooth objects of known geometry. The normals of the arbitrary shape are mapped onto a sphere, then after segmentation, the light source directions are determined from the mapped intensity information.

3.12 Li, Lin, Lu & Shum (2003)

This is the most recent technique [5] and probably also one of the most powerful techniques discussed in this paper. Instead of using a single cue for determining the light source locations, multiple cues (shading, shadow and specular reflections) are combined to assist in tackling the problem. The shading based component of the analysis used the critical point detection method proposed by Wang & Samaras.

Unlike any of the previous methods, this method is designed to be applied to textured objects.

4 Comparison

None of the techniques is flexible enough to be used in any situation. Each method has unique benefits that may make it more applicable to different situations.

One major attribute that differentiates the various techniques is the ability to operate on an unknown geometry. Interestingly, most of the earlier techniques attempted to accomplish this. Since Hougen & Ahuja's approach in 1993 the proposed methods have focused on known geometries. Another feature that only the more recent techniques show is the ability to detect multiple sources.

The following table compares the features of selected techniques in chronological order.

	Arbitrary Given Geometry	**Unknown Geometry**	**Multiple Sources**	**Textured Surfaces**
Pentland	✓	✓		
Weinshall	✓	✓		
Yang & Yuille	✓	✓	✓	
Hougen & Ahuja			✓	
Zhang & Yang			✓	
Wang & Samaras	✓		✓	
Li & Lin et. al.	✓		✓	✓

5 Conclusions

Within over 25 years since the problem of light source detection was first approached, numerous different solutions have been proposed. Most of them vary in terms of how the problem is tackled. This provides a good basis for future research.

There are still a number of challenges that remain for this field of research. Some of them are:

- **Dealing with different types of light sources.** Most current techniques assume directional lighting.
- **Handling textured objects.** Only one technique has been proposed so far.
- **Integration with shape from shading techniques.** Not many techniques are suited to be combined with shape from shading methods in order to deal with arbitrary unknown geometries.

References

[1] Brooks, M.J., Horn, B.K.P.: *Shape and Source from Shading*, Proceedings of the 9th Int. Joint Conference on Artificial Intelligence: pp. 932-936 (1985)

[2] Debevec, P.: *Rendering Synthetic Objects into Real Scenes: Bridging Traditional and Image-based Graphics with Global Illumination and High Dynamic Range Photography*, Proc. SIGGRAPH 98, pp. 189-198 (1998)

[3] Hougen, D.R., Ahuja, N.: *Estimation of the Light Source Distribution and its use in Integrated Shape Recovery from Stereo and Shading*, IEEE 4th Int. Conf. On Computer Vision, Berlin, Germany: pp. 148-155 (1993)

[4] Lee, C.H., Rosenfeld, A.: *Improved Methods of Estimating Shape from Shading Using the Light Source Coordinate System*, Artificial Intelligence, No. 26: pp. 125-143 (1985)

[5] Li Y., Lin S., Lu H. Shum H.: *Multiple-cue Illumination Estimation in Textured Scenes*, Proceedings of the Ninth International Conference on Computer Vision: pp. 1366-1373 (2003)

[6] Nillius, P., Eklundh, J.: *Automatic Estimation of the Light Source Direction*, Proceedings of the Conference on Computer Vision and Patter Recognition: pp. 1079-1083 (2001)

[7] Okatani K., Deguchi K.: *Estimation of Illumination Distribution Using a Specular Sphere* International Conference on Pattern Recognition: pp. 3596-3599 (2000)

[8] Pentland, A.P.: *Finding the Illuminant Direction*, J. Opt. Soc. Am. 72: pp. 448-455 (1982)

[9] Ullman, S.: *On Visual Detection of Light Sources*, Artificial Intelligence Memo 333, MIT (1975)

[10] Wang, Y., Samaras, D.: *Estimation of Multiple Illuminants from a Single Image of Arbitrary Known Geometry*, Computer Vision - ECCV 2002 Pt III Lecture Notes in Computer Science 2352: pp. 272-288 (2002)

[11] Wei, J.: *Robust recovery of multiple light source based on local light source constant constraint*, Pattern Recognition Letters 24(1-3): pp. 159-172 (2003)

[12] Weinshall, D.: *The Shape and the Direction of Illumination from Shading on Occluding Contours*, Artificial Intelligence Memo 1264, MIT (1990)

[13] Yang, Y., Yuille, A.: *Source from Shading*, Proceedings of the Conference on Computer Vision and Pattern Recognition: pp. 534-539 (1991)

[14] Zhang, Y.: *Illumination Determination and its Applications*, University of Saskatchewan (2000)

[15] Zhang, Y., Yang, Y.H.: *Multiple Illuminant Direction Detection with Application to Image Synthesis*, IEEE Transaction on Pattern Analysis and Machine Intelligence Vol. 23, No. 8: pp. 915-920 (2001)