Implementation of an Illuminant Detection Algorithm

Nathan Funk
University of Alberta
Class Project for CMPUT 615

December 16, 2003

Abstract

This paper discusses the implementation of the light source detection technique proposed by Zhang and Yang. The approach attempts to recover multiple light sources from an image of a sphere by locating cutoff curves. Three other approaches have extended this technique and generalize it to allow the analysis of arbitrary known geometries.

The implementation is performed in MATLAB and is tested using synthetic image generation code which allows arbitrary placement of light sources. This allows the analysis of the performance under a variety of conditions. Experiments are performed on two synthetic images. Three light sources are present in the first experiment. The recovered directions are accurate within 0.3 degrees, and the intensities are accurate within 1%. The second experiment considers 12 light sources, 8 of which are recovered successfully. The reasons for not detecting the four other sources are determined to be an inadequate method of searching the sphere for cutoff curves, and difficulties in locating of peaks corresponding to light sources in the Hough transform space. It is concluded that other methods of cutoff curve detection should be considered to improve robustness. Also, a more robust method of locating peaks in the Hough transform space is recommended.

1 Introduction

This project focuses on the implementation of a light source detection method as described by Zhang and Yang [11] [12]. This method is relatively recent (developed in 2000) and has shown outstanding results in comparison to previous techniques. Most previous methods could only recover a single light source from an image. The algorithm described by Zhang and Yang operates on the image of a sphere and can recover multiple directional light sources with high accuracy of both their direction and intensity.

At least three papers have been published that extend the original technique proposed by Zhang and Yang. Some of these new approaches are more general in the sense that they make less assumptions about the scene. This is evidence for the significance of the underlying concepts of the original technique.

This report provides some background information including a discussion of the problem of light source detection, and a survey of approaches found in literature. A brief comparison of the surveyed methods is also included.

The implementation was performed in MATLAB. In addition to the light source detection functions, a function was implemented to generate synthetic images for the analysis. This provides the advantage of total control over the light placement and intensities, and a ground truth comparison to evaluate the performance of the method. The theory involved in the method is discussed in the Approach section, together with some specific notes on the implementation. Finally, some experiments and their results are presented.

2 Background

2.1 General Background

Describing the problem of light source detection is quite simple: Given an image of a scene, the goal is to find the locations and intensities of the light sources illuminating that scene. For example, from an image of an outdoor scene, find the position of the sun.

Although *we* can solve this problem without great difficulty when seeing an image, it is much more difficult to implement in software. The main reason for this is that humans can typically recognize objects independent of the lighting. With this as a first step, we can then go on making conclusions about where the light is coming from. When solving this problem with a computer however, the knowledge of object shape is typically not available as a starting condition. This makes the process of determining the location of lights a whole lot more complicated.

In addition, the lighting conditions might not always be as ideal as in an outdoor scene with a clear sky and only the sun as the main light source. If the sky is overcast, the problem already becomes much more difficult. Similar problems can also result in scenes with multiple light sources (as is the case for most indoor lighting).

After considering the various types of problems that might be encountered, it can be concluded that substantial assumptions will need to be made to make progress in solving the problem. For the currently common techniques, these assumptions often include:

1. *The geometry of the scene is known.* Typically, a single object with known geometry is assumed to be visible in the image.

2. *The number of light sources is known.* For many of the early techniques a single light source is assumed.

3. *The light sources are directional.* Therefore, each light source location can be specified by two angles.

4. *The surface has a uniform albedo and Lambertian reflectance.*

5. *The surfaces are smooth.*

These assumptions can simplify the problem substantially. Especially knowing the geometry of the object in the scene is a major benefit. The assumption that light sources are directional also allows for point light sources at a large distance from the scene with respect to the size of the scene. The method described in this paper will use all these assumptions except the second one.

2.2 Literature Survey

There have been many different approaches proposed over the past 20 years. They differ greatly in the assumptions they make about the scene and also in the methods they employ to solve the problem. For example, some techniques assume that the geometry of parts of the scene is known. Others attempt to recover light sources in any scene without prior information about the geometry. Some techniques only analyze intensity information along the edges of objects, whereas others consider all intensity information available in the image.

First we will discuss the techniques up to and including Zhang and Yang's approach in chronological order:

Pentland (1982): Pentland's approach [6] 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. The average intensity changes in the x and y directions are used to determine the tilt and slant of a single directional light source. Due to the assumptions made, this approach is not applicable for all object shapes.

Lee & Rosenfeld's approach [4] (1985) is similar to Pentland's, however is targeted specifically at light source recovery from an image of a sphere.

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.

Weinshall (1990): Weinshall [9] 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. So the tilt can be estimated by searching for the extreme points of the intensity along the boundary. Finally the slant can be estimated by locating the brightest spot on the object.

Yang & Yuille (1991): This method [10] also employs the analysis of intensity information along the occluding boundaries. But it instead uses points at the silhouette to determine the light source direction. This approach allows the detection of multiple sources, however the number of sources is limited.

Hougen & Ahuja (1993): Hougen & Ahuja [3] use predefined directions as light source candidates, and assume a known shape (surface normal information is available). A drawback of this method is that the solution does not always converge.

Debevec (1998): Debevec's approach [2] was not originally intended for use in computer vision, but rather computer graphics. It is included in this survey due to its similarities with Zhang's technique. Debevec proposes the use of a mirrored sphere to find the light field illuminating the scene. A high dynamic range image is captured, and the intensities measured on the sphere are used to determine the intensities of the incoming light. The disadvantages associated with this approach with respect to computer vision, are that the scene may not be able to be manipulated and the use of a mirrored surface may also interfere with the lighting of the environment.

Zhang & Yang (2000): Similar to Debevec's method, this approach [11] [12] assumes that a sphere is visible in the scene. However, instead of a mirrored sphere, the surface is assumed to be lambertian. This makes the determination of the light source directions more complicated, but it has the benefit of not interfering with the environment to the same extent. More importantly, the method can be extended or generalized to operate on arbitrary known geometries.

A full discussion of this technique is provided in the Approach section.

2.3 Methods extending Zhang & Yang's approach

The following three documents cite Zhang and Yang's article [12] and develop techniques similar to, and extending upon Zhang's approach.

Wei (2001): This method [8] recovers multiple illuminants from the image of a sphere. The main difference from Zhang's approach is in the determination of the critical points (see Theory section). Wei claims that his approach is less computationally expensive and works with a more dense data set increasing the robustness of the technique.

Wang & Samaras (2002): Similarly, this method [7] also claims to outperform Zhang'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.

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. Unlike any of the previous methods, this method can also operate on textured objects.

2.4 Comparison of Techniques

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.

	Allows arbitrary geometry	Geometry can be unknown	Allows multiple sources	Allows textured surfaces
Pentland	✓	✓		
Weinshall	✓	✓		
Yang & Yuille	✓	✓	✓	
Hougen & Ahuja			✓	
Zhang & Yang			✓	
Wang & Samaras	✓		✓	
Li, Lin, Lu & Shum	✓		✓	✓

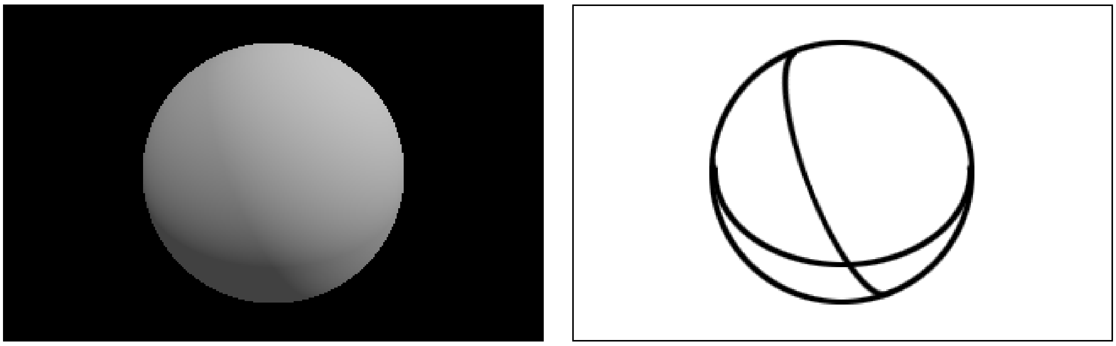

(a) Example synthetic image with two light sources (b) Corresponding cutoff curves

Figure 1: Illustration of a synthetic image and the corresponding cutoff curves.

3 Approach

As mentioned earlier, the method of Zhang & Yang operates on an image of a sphere. The surface of the sphere is assumed to be lambertian, meaning that light is reflected equally in all directions from the surface. Also, the albedo of the sphere is uniform. The light sources are assumed to be directional.

In contrast to the earlier techniques, this method determines lighting information from the *cutoff curves* caused by illumination. Figure 1(a) displays a typical scene with a sphere lit by two directional light sources. *Cutoff curves* are actually great circles around the sphere where the intensity dropoff from a single light source is a maximum. They can also be described as the edge of the silhouette generated by each individual light source. The cutoff curves corresponding to the image are depicted in Figure 1(b). Under orthographic projection they appear as segments of ellipses in the image.

By detecting the location and orientation of the cutoff curves on the sphere, it is a relatively simple task to then infer the *direction* of the light sources. Each cutoff curve lies on a plane perpendicular to the direction of the corresponding light source. Or mathematically, letting the light sources be L_i for $i = 1, 2, \ldots, m$, we can express each cutoff curve Σ_i with

$$\Sigma_i = \{P | L_i \cdot N_P = 0\}, \quad (1)$$

where P are points on the surface, and N_P is the surface normal at P.

It would be difficult to locate cutoff curves directly from the image without any intermediate steps. Zhang accomplishes the task of finding the cutoff curves by searching for points that lie on these curves, then grouping these points together to find individual curves. Zhang uses the term *critical points* to name all the points along these cutoff curves (as expressed in equation 1). He establishes that they are easily distinguishable from other points on the sphere under usual circumstances.

In particular he shows that the critical points divide the image of the sphere into a finite number of regions in which the intensity is determined by a simple function of position. In Figure 1(b) it can be seen that four such regions exist. This simple expression defining the intensities within each region becomes a crucial tool for locating critical points.

Figure 2 highlights a single great circle (projected as a line) on a sphere illuminated

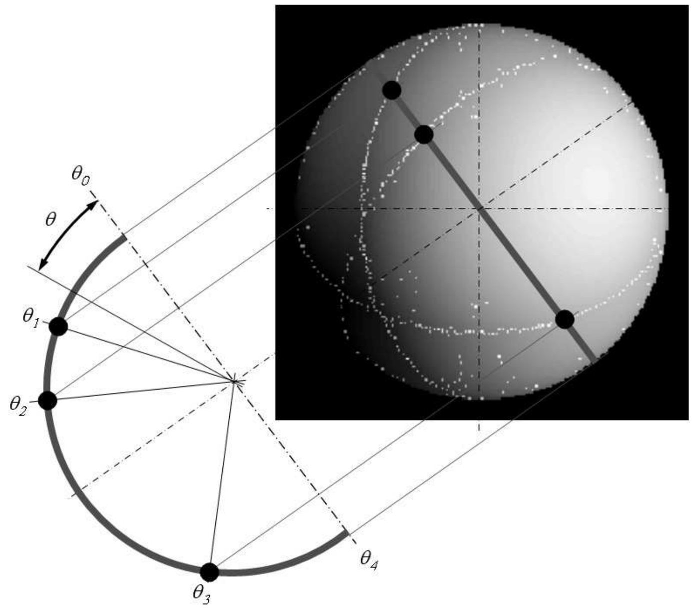

Figure 2: Illustration of critical points (at θ_1, θ_2, θ_3) on the visual section of a great circle.

by three light sources. The diagrammatic cross section illustrates how the position of any point on this great circle can be described with a single angle measurement θ. The visual part of the circle is divided up into four sections separated by θ_0, θ_1, θ_2, θ_3, and θ_4. Zhang proves that for each section $[\theta_{i-1}, \theta_i]$, the intensity $\mathrm{E}(\theta)$ is determined by the equation

$$\mathrm{E}(\theta) = \mathrm{b_i} \sin\theta + \mathrm{c_i} \cos\theta, \tag{2}$$

where the coefficients $\mathrm{b_i}$ and $\mathrm{c_i}$ are constants for each section i. These coefficients play a major role in determining the location of critical points as section 3.1 will show.

It is important to point out that there are cases in which light sources are extremely difficult to detect, and possibly even undetectable. One example is a light source which is located along the line between the camera and the center of the sphere. The cutoff curve falls onto the boundary of the sphere for orthographic projection, and might not even be visible in perspective projection. In the case of perspective projection, any light source within a small angle of this line will be undetectable since its cutoff curve is occluded.

Another difficult situation occurs, when two light sources are pointing along the same line, but in opposite directions. Their cutoff curves will overlap. Since these cases are rare for randomly located light sources and camera positions they are not considered in the implementation. However, they are discussed and handled in greater detail in [11] and [12].

Before discussing the details of the approach, we will outline the necessary steps for clarity:

1. **Finding Critical Points:** The visible area of the sphere is searched for points lying on cutoff curves - critical points.

2. **Grouping Critical Points:** The critical points found in the first step are grouped

into cutoff curves from which the *pre-direction* of the light sources can be determined. The side on which the light source is located is still unknown.

3. **Determining Light Source Intensities:** Using the pre-direction of each light source, the final direction and intensities are determined.

3.1 Finding Critical Points

The process of finding critical points involves searching the image along great circles of the sphere, which project to arcs in the image. Great circles are chosen because the formulation of the intensities along them is simple, and has already been derived. The measured intensities from the image can be compared to the formulated intensities, and from this the location of critical points derived.

The great circles chosen in this implementation are configured in a star pattern. Each great circle projects to a line passing through the center of the sphere in the image. The search along each circle is performed in small angular steps.

Each point on a great circle can be identified with a single angle measurement. We will use θ to represent a point on a given circle. At a given point θ, two small segments, "subarcs", to either side of the point can be defined with $[\theta - H, \theta]$ and $[\theta, \theta + H]$. According to the constraints identified earlier, the intensities along each subarc are defined by equation (2), so it is possible to find the coefficients b_i and c_i for each subarc. These can be found with a least squares solution to a set of irradiance equations (each like equation (2)). We denote the coefficients as b_-, c_- for the first subarc, and $\mathrm{b}_+, \mathrm{c}_+$ for the second. If there is no critical point between $\theta - H$ and $\theta + H$, then ideally b_- would equal b_+ and c_- would equal c_+. So we now can formulate a method of identifying whether there is a critical point in this region. To make it robust against slight inaccuracies of the coefficient values, we can compare the geometric distance to a threshold value T with

$$\sqrt{(\mathrm{b}_- - \mathrm{b}_+)^2 + (\mathrm{c}_- - \mathrm{c}_+)^2} < T. \tag{3}$$

Hence, this equation serves as a test for finding critical points. Zhang shows that the geometric distance expression is equal to the sum of all intensities of light sources involved in only one of the two adjacent sections. This is however not necessarily important for the implementation and is therefore not discussed in greater detail. Only the fact that light sources with intensities below the threshold values can therefore not be detected, should be noted.

Simply knowing that the critical point is somewhere between $\theta - H$ and $\theta + H$ does not give us an exact position. From the coefficient values, an estimate of the critical point position can be made.

3.1.1 Initial estimate

The coefficients define two intensity curves. So an estimate can be made by finding the angle at which these intensity curves intersect. This can be accomplished by solving

$$\mathrm{E} = \mathrm{b}_- \sin\theta + \mathrm{c}_- \cos\theta = \mathrm{b}_+ \sin\theta + \mathrm{c}_+ \cos\theta, \tag{4}$$

with respect to θ. It follows that

$$\theta = \arctan\left(\frac{\mathrm{c}_+ - \mathrm{c}_-}{\mathrm{b}_+ - \mathrm{b}_-}\right). \tag{5}$$

This will in fact give the exact position of the critical point if it is located precisely between the two subarcs. This is however quite unlikely. When a critical point lies on one of the subarcs, its calculated coefficient values will not accurately represent the associated intensity curve and the calculated position θ will be inaccurate. So we need a way to refine the initial guess further in order to obtain a more exact value.

3.1.2 Refinement of Critical Points

The refinement process implemented uses a simple idea. Shift the subarcs so they are centered around the estimated position of the critical point and repeat the original process. In addition, the subarcs are pushed away from each other slightly by an angle ϵ so they are located at $[\theta' - H - \epsilon, \theta' - \epsilon]$ and $[\theta' + \epsilon, \theta' + H + \epsilon]$ where θ' is the estimated position. After finding the coefficient values for the shifted subarcs, a new estimate of the critical point location can be made. The refinement process can be repeated until the estimated position does not change significantly.

3.2 Grouping Critical Points

Now that a set of critical points have been found, they need to be grouped into cutoff curves. The main challenges involved are that the number of cutoff curves is unknown, and there can be a significant amount of error in the measured critical point locations. Although Zhang mentions that a Hough transform can be used to approach this problem, it is not explained in what exact manner this is accomplished.

The method employed in this project defines two spaces:

1. **Critical Point Space (CPS):** This is a discrete 2D space of binary values. Spherical coordinates are used to represent the location of critical points.
2. **Light Direction Space (LDS):** This is the Hough transform space or *accumulator*. It is also a discrete 2D space with spherical coordinates describing the direction of light sources.

When many critical points are found with little error in their location, they form curves in the CPS. Essentially, we are trying to find best fitting curves for the critical points. So to fill the accumulator, we need to count the number of critical points that lie on each curve corresponding to a point in the LDS. To do this we need to formulate an equation for the curve in CPS corresponding to a point in LDS. Or more precisely, we need a function of the form $\beta = f(\alpha, \theta, \phi)$ where α and β are the spherical coordinates of points in CPS, and θ and ϕ are the spherical coordinates of light directions in LDS.

Through a derivation based on simple geometric constraints it can be shown that

$$\beta = \arctan\left(\frac{1}{\tan\phi(\cos\theta\cos\alpha + \sin\theta\sin\alpha)}\right). \tag{6}$$

Given this relationship it is now easy to count the number of critical points contributing to each considered cutoff curve.

The final task for this step is to try to determine the light source positions given the values in the LDS. Simply thresholding the values was found to be insufficient since for some cutoff curves much less critical points are found than for others. And in addition, the peaks are not always clearly located at one point in the LDS, but spread out over a region.

To approach these problems, the peak location process is broken down into three steps. (1) Find local maxima and plateaus, (2) select only the maxima above a specified threshold and (3) remove points that are within *close range* of a "higher" point. *Close range* is formulated in the implementation as an angular limit between two potential light source directions.

3.3 Determining Light Source Intensities

With the cutoff curves identified, the only pieces of information left to recover are the light source intensities. These can be found by solving an equation system of intensity equations. By selecting a set of surface points P_j with corresponding surface normals N_j, an equation system can be set up to relate the intensity values at each of these points in the image $\mathrm{E}(P_j)$, to the irradiance equation:

$$\mathrm{E}(P_j) = \sum_i \max(e_i L_i \cdot N_j, 0), \tag{7}$$

where e_i is the intensity of the light L_i. All variables in this system are known except the intensities e_i. By applying a least-squares technique, each intensity can be determined.

It should be pointed out that the effects of undetected light sources are ignored in this discussion and implementation. They are included in Zhang and Yang's work [11] [12].

3.4 Implementation

All the steps discussed in this report were implemented as MATLAB functions. Zhang's discussion does not focus on implementation aspects, so some decisions needed to be made regarding the implementation. The most important decisions are listed here:

Finding Critical Points: As mentioned earlier, the great circles to search along are arranged in a star pattern with each of the circles projecting to a line in the image. This pattern was chosen for simplicity of implementation, while still covering the sphere reasonably.

It was decided to extend the refinement step a little to avoid false detection of critical points. So in the implementation, if the refinement does not converge to under 1% change within 5 iterations, the critical point is not included in the final set.

Grouping Critical Points: Zhang does not discuss how the Hough transform is implemented, so section 3.2 is entirely based on decisions made for this specific implementation.

Determining Light Source Intensities: A simple approach was chosen to solve the equation system of this step. Since it is non-linear, the built-in MATLAB least squares could not be used. So instead the `fminsearch` minimization function is employed.

4 Experiments

4.1 Experimental design

The goal of these experiments is in part to demonstrate how well the technique works under good conditions, and also to point out the limits of this implementation and the general approach.

Experiment 1: A test is performed on a three light source setup in which the lights are located to produce easily detectable cutoff curves. This experiment should provide a measure of how well the method operates under "good conditions".

Experiment 2: This experiment compares the results for a case with 12 light sources to those obtained by Zhang [11]. The light sources are randomly placed however none of them are undetectable, so the assumption of not handling undetectable light sources made in this implementation is avoided.

All experiments use the following parameters unless noted otherwise:

- Synthetic image of a sphere with a diameter of 180 pixels. Pixel intensity values are recorded with `double` precision.
- Subarc length of $H = 5$ degrees with 10 points on each subarc. Threshold of $T = 0.1$.
- Refinement parameters: $\epsilon = 0.2 * H$, maximum of 5 iterations
- Hough transform parameters: CPS resolution 180x180 (1 degree accuracy of critical points), LDS resolution 720x720 ($\frac{1}{4}$ degree accuracy for light directions), threshold for peak finding: 40% of maximum peak value, merging angle: 5 degrees ("close range")
- Least squares solution: Intensity at 100 points on the sphere are evaluated in the minimized function.

4.2 Results

Experiment 1: The results of the experiment are summarized in the following table:

	Original			Recovered		
Source	1	2	3	1	2	3
Tilt (degrees)	90.0	0.0	-40.0	90.0	0.3	-40.0
Slant (degrees)	130.0	130.0	120.0	130.0	130.0	120.0
Intensity	0.400	0.500	0.300	0.400	0.498	0.303

As expected the results are very accurate for this case. All angle measurements are accurate within 0.05 degrees except the recovered tilt value of source 2, which shows an error 0.3 degrees. The intensity values are also accurate with a maximum error of 1% in the intensity of source 3.

It is essential to note though, that the high accuracy of the angular results can be partially attributed to the fact that the Hough transform space was divided up in quarter degree intervals. By having specified light directions with integer angles, the accuracy appears higher than it actually is. The intensity values are not subject to the same issue.

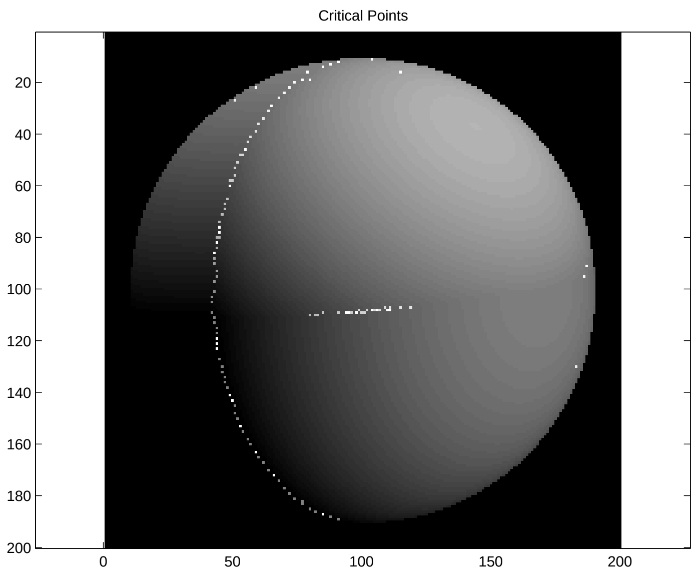

Figure 3: Illustration of problems in finding critical points when using the star pattern. Many critical points are found along the vertical cutoff curve, but few are found on the horizontal cutoff curve due to the small angle between the cutoff curve and the search direction.

But it can also be noted that their accuracy depends on the number of iterations performed in the minimization algorithm.

Experiment 2: Only 8 of the 12 light sources are detected with this implementation. However, all 12 were recovered with Zhang's implementation. The light source directions are not provided here since they are not considered relevant in comparison to the following implications of these results.

One reason of light sources not being detected is that too few critical points are found on their cutoff curves. This was also noticed in other experiments not documented here. Specifically, the star pattern arrangement of the searched circles was discovered to be inadequate. If a cutoff curve is tangential or at a shallow angle to a searched circle, critical points along it will not be detected reliably. This effect is illustrated by a comparison of two light sources in Figure 3. The critical point detection algorithm operates best when the circle is orthogonal to the cutoff curve. Although it is not clearly stated in [11] what pattern the search was performed in, the paper [12] mentions searching two orthogonal arcs locally around each pixel in the image. This approach has the benefit that the cutoff curves will always be at an angle of at least 45 degrees to one of the searched arcs, thus obtaining better critical point measurements.

Since for some of the cutoff curves less critical points are detected, their associated peak in the Hough transform is not as high as the others. This causes difficulties in finding all the peaks. If the threshold is set too high, some peaks are not detected (as was the case for this experiment). If the threshold is set too low, false peaks are found.

5 Conclusions & Recommendations

The following conclusions were drawn from the literature survey and the implementation of this technique:

- The assumption of a sphere being visible might be considered too restrictive in some cases. Depending on the situation, it may not be possible to manipulate the scene. And obviously a sphere with lambertian surface can not be expected for any arbitrary scene. Therefore applications of this technique are limited to cases where the scene can be manipulated.

- The method performs with great accuracy in cases where the critical points are easily detected.

- The algorithm for searching for critical points does not perform well when a cutoff curve is at a small angle to the search direction. The method mentioned in [12] where orthogonal arcs are employed is expected to give better results. Other critical point detection techniques such as the one proposed by Wei [8] are also expected to give better results. This method was tested and showed promising results, however a discussion could not be included due to the page limit.

- If the peak finding threshold is set too low, false peaks are detected in the Hough transform space. If it is set too high, some light sources might not be detected.

Recommendations for future extensions of this implementation, and further experiments are as follows:

- Instead of using the star pattern search as implemented in this project, examine other options for critical point detection and analyze their effectiveness in different situations (noisy images, real images...)

- Implement a more robust peak finding algorithm for the Hough transform.

- Perform tests on the maximum number of light sources detectable.

- Analyze robustness to different sources of noise and distortion.

Finally, I would like to thank Professor Herb Yang for his helpful comments and recommendations during the implementation of this project.

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

[7] 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)

[8] 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)

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

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

[11] Zhang, Y.: *Illumination Determination and its Applications*, Master's Thesis, University of Saskatchewan (2000)

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