A Review of Selected Single Light Source Detection Methods

Nathan Funk
University of Alberta
Assignment 1 for CMPUT 605

September 27, 2004

1 Pentland

Pentland [3] was the first to attack the problem of estimating a light source that is not directly visible in an image. Before him, Ullman [5] published the paper "On Visual Detection of Light Sources" which however only focuses on recovering the locations of light sources that are visible within an image.

One of the most significant aspects of this paper is that it attempts to recoverillumination conditions without assuming that an object with known geometry is visibile in the image. Instead it uses a statistical approach to recover lighting for arbitrary scenes.

Three main assumptions are made:

1. All surfaces reflect according to the Lambertian reflectance model
2. All illumination originates from a single distant light source
3. The surface normals of the objects in the scene are isotropically distributed

The first and second assumptions are made in many other early papers as well. Although it appears limiting, the study of images lit by a single light source proves to be difficult enough.

The third assumption is also somewhat limiting but necessary for this type of statistical analysis. Nearly spherical objects and complex scenes with many randomly distributed objects will almost satisfy this condition. But images of flat or cylindrical surfaces are not as appropriate. In fact any scene where the surface normal orientation is strongly biased to certain directions are not very well suited for this approach.

Pentland begins his derivations with explaining the process of image formation under the Lambertian model. After following the path of the light to the surface and then being reflected into the camera he arrives at

$$I = \rho f(N \cdot L), \tag{1}$$

where I is the image intensity, ρ is the local albedo, f is the incoming flux per unit area, N is the surface normal, and L is the light source vector. It can be noted that the case where $N\dot{L}$ produces a negative value is not considered. This has the effect that in the application of this method, self-shadowed regions must be excluded from the domain of analysis.

Further assuming that the albedo and illumination are constant within a small region of the surface the following can be shown:

$$\mathrm{d}I = \rho f(\mathrm{d}N \cdot L). \tag{2}$$

This equation serves as the basis for determining the x and y components of the light source direction. Under the assumption that $E(dN) = 0$, a regression model is constructed. The regression model relates the average change in image intensity along specific directions to the x and y components of the light source as follows:

$$\begin{bmatrix} \mathrm{d}\bar{I}_0 \\ \vdots \\ \mathrm{d}\bar{I}_n \end{bmatrix} = \begin{bmatrix} \cos\theta_0 & \sin\theta_0 \\ \vdots & \vdots \\ \cos\theta_n & \sin\theta_n \end{bmatrix} \begin{bmatrix} \hat{x}_L \\ \hat{y}_L \end{bmatrix}. \tag{3}$$

After solving for $\hat{x}_L$ and $\hat{y}_L$ these two values can be used to determine the tilt with:

$$\tau_L = \arctan\left(\frac{\hat{y}_L}{\hat{x}_L}\right). \tag{4}$$

The derivation of the slant equation involves calculation of the variance of dI. The final equation for the slant can be written as

$$\sigma_L = \arccos\left[\left(1 - \frac{\hat{x}_L^2 + \hat{y}_L^2}{k^2}\right)^{1/2}\right], \tag{5}$$

where

$$k = [E(\mathrm{d}I^2) - E(\mathrm{d}I)^2]^{1/2}. \tag{6}$$

Pentland goes on to compare the performance of his estimator to human estimates but does not does not provide a performance analysis under ideal conditions. As later papers show, the proposed slant equation does not give exact estimates under ideal conditions. If all the assumptions of the algorithm are reflected in an ideal test case (such as the image of a sphere), the slant should in fact be estimated with high accuracy. Lee and Rosenfeld [2] as well as Chojnacki et al. point out that Pentland's slant estimation does not achieve the expected result.

2 Lee and Rosenfeld

Lee and Rosenfeld's 1985 paper [2] concentrates on a shape from shading technique, but they also develop a new light source detection technique. It is in many ways similar to Pentland's, but the derivation is clearer and experiments on ideal cases are included.

They implement Pentland's method as a basis for comparison for their new technique. In their discussion, they point out that the regression model using derivatives in many directions is not necessary. Since the first derivative in any direction is a linear combination of the derivatives I_x and I_y, the regression model offers little benefit other than perhaps reducing the effects of noise. Their Pentland implementation was tested on an image of a sphere for which it should give accurate results. They note that this is however not the case:

> We see from these results that Pentland's method is a very good tilt estimator, as well as a very good slant estimator if the slant is below 40 degrees. However, the slant estimation gives 90 degrees whenever the slant is above 50 degrees, due to the small variance of $\mathrm{d}I$, i.e., $E\{\mathrm{d}^2 I\} - E^2\{\mathrm{d}I\}$.

The slant errors they report for a 0 tilt case are as follows:

Actual slant	0	11	22	33	44	56	67	78	89
Slant error	0	1	1	3	23	34	23	12	1

Lee and Rosenfeld do not provide a discussion on where the flaw in Pentland's derivation is.

The new tilt estimation technique is nearly the same as Pentland's approach, however the derivation is greatly simplified. They start with expressing the image derivatives I_x and I_y in terms of the tilt τ and slant σ of points on the sphere:

$$I_x = \frac{\lambda\rho}{R}(-l_1 + l_3 \cdot \tan\sigma \cdot cos\tau), \tag{7}$$

$$I_y = \frac{\lambda\rho}{R}(-l_2 + l_3 \cdot \tan\sigma \cdot sin\tau), \tag{8}$$

where (l_1, l_2, l_3) is the light source direction. Then the expected values of I_x and I_y can be determined by integrating over the visible hemisphere. For this integration, the sampling distribution needs to be considered. This plays an important role in their analysis. The final results of the integration are:

$$E\{I_x\} = -\frac{\lambda\rho}{R}l_1, \tag{9}$$

$$E\{I_y\} = -\frac{\lambda\rho}{R}l_2. \tag{10}$$

It is simple to now formulate the light source tilt as $\tau_l = \arctan(E\{I_y\}/E\{I_x\})$.

Lee and Rosenfeld propose two different methods of estimating the slant. The first method only uses the intensity information along the line with the same tilt as the light source (which is assumed to have been already determined). The intensity profile along this line is $I = \lambda\rho\cos(\sigma - \sigma_l)$. Using a similar approach to the tilt estimation, the expected value of the intensity $E\{I\}$ is calculated by integrating over the lit region. The result is

$$E\{I\} = \lambda\rho\frac{\pi\cos\sigma_l - \sigma_l\cos\sigma_l + \sin\sigma_l}{2(1+\cos\sigma_l)}. \tag{11}$$

Instead of using the intesity values directly, the same approach can be applied to I^2 resulting in the simpler equation

$$E\{I^2\} = \frac{\lambda^2\rho^2}{3}(1+\cos\sigma_l). \tag{12}$$

Either one of these equations can be used to determine the slant. It is pointed out that when using equation 11 the brightest spot in the image needs to be found in order to determine $\lambda\rho$. According to the paper this is not necessary for eqation 12. In their experiments, this slant estimation technique performed well for a tilt of 0, but not as well for other tilt values.

The second slant estimation technique is slightly different. Instead of using only the intensity information along a single line, the entire image is considered. The derivation is similar again and they finally arrive at:

$$E\{I\} = 4\lambda\rho\frac{\pi\cos\sigma_l - \sigma_l\cos\sigma_l + \sin\sigma_l}{3(1+\cos\sigma_l)}. \tag{13}$$

Note that surprisingly this is the same as equation 11 except for the multiplication by a factor of 2/3.

Their implementation of the second slant estimator produced good results with slant errors below 7 degrees, with the majority of estimates having only 2 degree or less error. Experiments were also performed on an image of an ellipsoid. As expected the errors are much larger for this case. The maximum tilt error was 16 degrees, and the maximum slant error was 22 degrees.

3 Chojnacki et al.

This 1994 paper reviews Pentland's estimator, pointing out it's main flaw in the slant estimation, and proposes a new slant estimation technique.

Instead of using Pentland's original equations, Chojnacki brings them from differential form into integral form. This adds some clarity to the equations since Pentland's original paper makes little note of the direction in which certain differentials are taken. In the integral form the direction of differentiation is clearly stated.

Like Lee and Rosenfeld's article, this paper points out that the partial derivatives of the image intensity in arbitrary directions is always a linear combination of the derivatives in the horizontal and vertical directions. This allows determining the tilt τ_l without the regression model proposed by Pentland. Instead the tilt is simply calculated from the image derivatives E_x and E_y:

$$\tau_l = \arctan\left(\frac{\mathbb{E}^{\Omega}\{E_y\}}{\mathbb{E}^{\Omega}\{E_x\}}\right). \tag{14}$$

The slant also does not rely on the regression model:

$$\sigma_l = \arccos\left[1 - \frac{(\mathbb{E}^{\Omega}\{E_x\})^2 + (\mathbb{E}^{\Omega}\{E_y\})^2}{(\mathrm{Var}^{\Omega}\{E_s\})^2}\right]^{1/2}. \tag{15}$$

The slant formula is assumed to be independent of the choice for the direction $\mathbf{s}$. It was noted in the implementation that this choice affects the results. Neither Pentland or Chojnacki go into details about how this direction might be chosen or how the choice affects the results.

The paper includes a derivation of the tilt estimator, verifying Pentland's result although in a slightly different form. The invalidity of the slant formula is approached by showing that the denominator in the equation $\mathrm{Var}^{\Omega}\{E_s\} = +\infty$. Although this only holds for continuous images, Chojnacki shows through experimental results that the slant estimate accuracy decreases with increasing sampling frequency. A theoretical discussion of how slant results are affected in the discrete domain is however not provided. His experimental analysis measures the slant error for actual slant values ranging from 0 to 90 deg. The slant error increases from 0 with increasing actual slant, then decreases again after peaking at about 45 deg. The experiment analyses 6 different resolutions ranging from 64 x 64 to 768 x 768. The maximum error is consistantly larger for higher resolutions.

Finally, the Disk Method is proposed in order to fix the resolution dependency problem. In this method the image domain over which the derivatives are calculated is reduced. For the image of a sphere, the edge regions are excluded by introducing a parameter α. Given R as the radius of the sphere, and letting $0 < \alpha < 1$, a closed disk $D_{\alpha R}$ with a radius of αR is defined. The center of the disk coincides with with projected center of the sphere. Since $\alpha < 1$ there is a small region along the edge of the sphere which is excluded from this domain. The paper continues with the derivation of the tilt and slant based on this new domain. The final tilt and slant formulas are:

$$\tau_l = \arctan\left(\frac{\mathbb{E}^{D_{\alpha R}}\{E_y\}}{\mathbb{E}^{D_{\alpha R}}\{E_x\}}\right), \tag{16}$$

$$\sigma_l = \arccos\left\{1 + \frac{\theta(\alpha)[(\mathbb{E}^{D_{\alpha R}}\{E_x\})^2 + (\mathbb{E}^{D_{\alpha R}}\{E_y\})^2]}{(\mathrm{Var}^{D_{\alpha R}}\{E_s\})^2}\right\}^{1/2}. \tag{17}$$

where

$$\theta(\alpha) = -\frac{1}{2} - \frac{1}{2\alpha^2}\ln(1-\alpha^2) \tag{18}$$

The same experiment is performed on the new slant estimator, and the results show that the slant estimation is accurate independent of the resolution. The authors however do not provide details on how the

value of α is to be selected or calculated. They also do not report whether the value of α affects the results. Daniel Gibbins' Ph.D. thesis [1] also does not provide more detail on the value of α.

4 Sinha and Adelson

The vast majority of illumination detection techniques assume that objects in the scene have a smoothly curved surface. Sinha and Adelson [4] instead decided to focus on recovering lighting and shape from images of polyhedral objects with exclusively flat surfaces.

Approaches like the previous three reviewed papers do not perform well on scenes where the distribution of normals is not isotropic. A completely different approach is taken in this paper. The one assumption that is the same is that the scene is illuminated by a distant point light source. The first half of the paper concentrates on the preliminary steps that are necessary prior to recovering the lighting.

The preprocessing of the image is broken down into:

1. Segmenting areas of constant brightness.
2. Finding edges between areas of different brightness. This assumes that the edges are straight lines specified through a start and end point.
3. Classification of junctions into Arrow, Y, and T.
4. Classification of edges into *illumination edges* or *reflectance edges.*

The last step is a crucial part of this technique. The edge classification is based on the classification of the associated junction points. If both junctions at the ends of the edge are T type, then the edge is classified as reflectance edge. This is due to the fact that in most cases this edge is not cause by lighting but instead a change in the reflectance properties of the surface. For other cases where an Arrow or Y type junction is found, the edge is more likely to be caused by a different orientation of the adjacent areas and thus differing intensity of reflection towards the camera.

Recovering the depth of the junction points is approached by using a metric of angle variance, planarity of each polygon, and compactness of the entire structure. The point out that there is also necessity to verify of the consitancy of the recovered structure and the image. In a sense this means that one should be able to roughly reproduce the image by illuminating the recovered model with a single light source.

Sinha and Adelson do not choose to make specific assumptions about the BRDF for checking the consistancy. Their main reason is that when the BRDF is in fact slightly non-Lambertian it will cause a major disortion of the recovered model. Similarily, they point out that minor alterations in the image grey-level values also radically alter the computed solution. They try to make their technique "gracefully tolerant" of this possible variances.

This is achieved by only using binary relations between areas of differing brightness values. While recovering the lighting the only information besides the direction of the normals is whether one surface is brighter or darker than another. For every edge, the binary relation of the adjacent surfaces restricts the direction of the incoming light. As Figure 4 shows, given the two normals of surfaces S1 and S2 and the binary relation, half of the gaussian sphere can be excluded from the possible directions of the light source. The figure also assumes that the light source is constrained to lie above the horizontal plane.

The mathematical constraints can be defined as follows. First compute the vector $\bar{s}$ for which $\bar{s} \cdot ((\bar{n}_1 + \bar{n}_2)/2) = 0$ and $\bar{s} \times (\bar{n}_1 \times \bar{n}_2) = 0$. If S1 is brighter than S2, then $\bar{s} \cdot \bar{n}_1 > 0$ and $\bar{s} \cdot \bar{n}_2 < 0$. Otherwise $\bar{s} \cdot \bar{n}_1 < 0$ and $\bar{s} \cdot \bar{n}_2 > 0$. Then the direction of the light source t can be any vector that satisfies $\bar{s} \cdot t > 0$.

After constraining the light source direction through the analysis of all edges, the final result might look something like the sphere on the right of Figure 4.

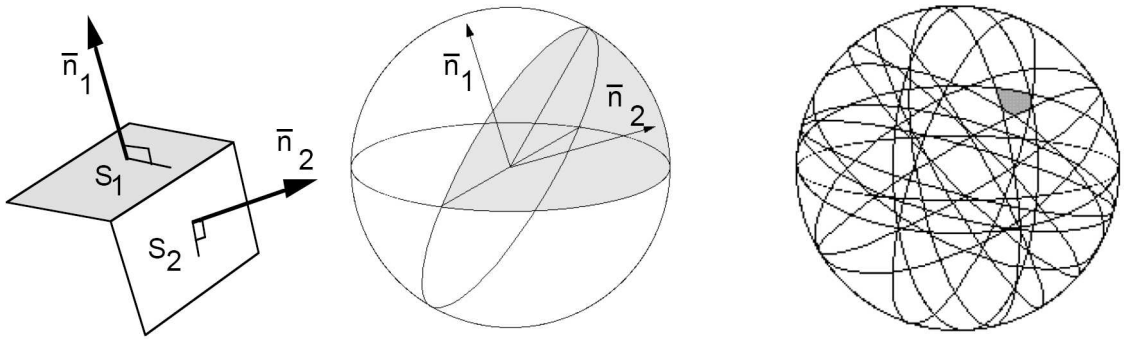

Figure 1: Left and centre: Illustration of two adjacent surfaces and their relative brightness restricting the light source direction to the shaded area on the Gaussian sphere. Right: Example of the constrained region of the possible light source directions for a more complex scene.

Depending on the number of surfaces visible in the image and the distribution of their orienatations the accuracy of the method will vary. It is interesting to observe that the method will likely perform well on most non-Lambertian surfaces. Also, since illumination edges and reflectance edges are distinguished, certain surfaces with varying albedo can also be permitted.

5 Implementation

The implementation component of this assingment focuses on the methods proposed by Pentland and Chojnacki et al. The following techniques were implemented in a common framework:

1. Pentland's original estimator
2. Chojnacki's differential formulation of Pentland's estimator
3. The Disk Method proposed by Chojnacki et al.

The experiments confirmed that the tilt estimation is good in all of the three methods. Typically the tilt estimate was within 2 degrees of the actual value.

The slant error however was more significant in all three techniques. As Figure 5 shows, the slant error is acceptable for all methods as long as the actual slant is below 50 degrees. The errors in methods 1 and 2 are roughly the same, confirming that Chojnacki's reformulation is correct. For these techniques however, for slants above 55 degrees the estimator returns 90 degrees. This is due to the expression under the square root becoming negative. The resulting number is complex with a real component of 0, making the slant estimate 90 degrees. It is possible that this is the same effect that Lee and Rosenfeld noticed. Their slant error values for slants under 50 degrees do however not agree exactly with the results from this experiment. The difference between methods 1 and 2 at 90 degrees slant was not studied in detail.

The results of the Disk Method are also not as good as those presented by Chojnacki et al. One possible reason is a lack of information on how the value of the parameter α is to be used. Since the paper does not offer any details on α, although it influences the results, it is difficult to reproduce the original results.

Finally, the effect of increasing the resolution was studied in attempt to confirm the results of Chojnacki et al. Figure 5 confirms that as the resolution increases, so does the error in the slant.

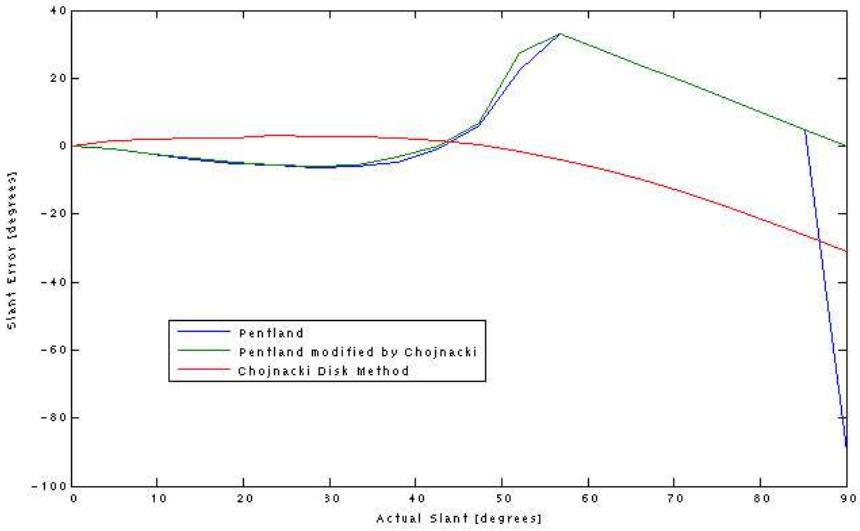

Figure 2: A comparison of the slant error for the three implemented techniques.

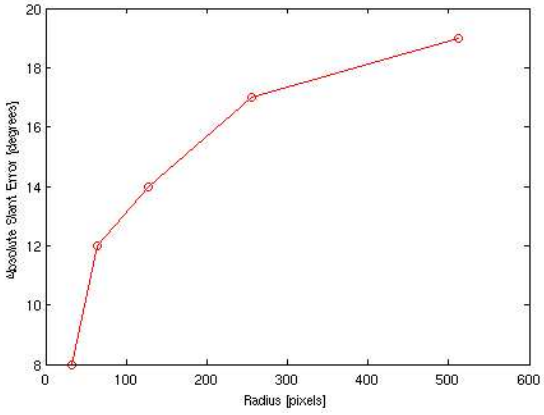

Figure 3: The effect of increasing the resolution of the image on Pentland's slant error. The resolution is measured here as the radius of the sphere in pixels.

References

[1] D. Gibbins. *Estimating Illumination Conditions for Shape from Shading*. PhD thesis, Flinders Univeristy of South Australia, May 1994.

[2] C. H. Lee and A. Rosenfeld. Improved methods of estimating shape from shading using the light source coordinate system. *Artificial Intelligence*, (26):125–143, 1985.

[3] A. P. Pentland. Finding the illuminant direction. *Journal of the Optical Society of America*, (72):448–455, 1982.

[4] P. Sinha and E. Adelson. Recovering reflectance and illumination in a world of painted polyhedra. In *IEEE 4th Int. Conf. On Computer Vision*, pages 156–163, 1993.

[5] S. Ullman. On visual detection of light sources. *Artificial Intelligence Memo*, (333), 1975.