

A Practical Method for Estimation of Point Light-Sources

Martin Weber* and Roberto Cipolla
Department of Engineering, University of Cambridge
Cambridge, CB2 1PZ, UK
mw232@cam.ac.uk

Abstract

We introduce a general model for point light-sources and show how the parameters of a source at finite distance can be estimated by shading on an object with known geometry and Lambertian reflectance. The parameters we estimate include not only the direction but also the location of the source. Furthermore, we argue that the system introduced here can be used in arbitrarily complex background illumination, provided one can switch on and off the light-source that is to be estimated.

1 Introduction

The effect of a light-source can be seen as a map that supplies each point in space with a radiance value. The correct estimation of illumination plays an important role in many branches of Computer Vision. The importance of light estimation is explicitly apparent in all applications which use photometric measurements (intensities) which obviously depend on the illumination. However, the estimation of light can also be crucial for techniques such as texture based vision which are not explicitly intensity dependent. Concrete applications include all reconstruction problems that use photometric measurements (such as shape from shading [4, 3], photometric stereo [11, 13]), but also for instance Augmented Reality applications [5] where virtual objects have to be rendered into a scene with correct lighting and shadow. We focus our attention on reconstruction problems while keeping the set-up general enough not to exclude other applications.

2 Previous Work

In most of the early work done on reconstruction [4, 3], a point light-source at infinite distance is assumed [8], in some cases in special alignment with the camera. Comparably much effort was made in some cases [12], to calibrate the position and intensity of light-sources manually to the required degree of accuracy. Some authors also address the related question of reflectance estimation in their frameworks [10, 14]. More recent work discusses methods that were developed to estimate complex *illumination distributions* in infinite distance approximation. In the literature, one finds two different approaches to model the illumination distribution over the hemisphere of incident directions: either as a *functional expansion* [6] or as a *discretisation by sub-dividing* the hemisphere [9].

* Supported by the EPSRC, the Cambridge European Trust and the DAAD (Germany).

3 Motivation and Applications

As we will demonstrate, the standard assumptions of infinitely far and/or sources aligned with the camera, are not necessary. Furthermore, we show how to estimate the light-source reliably, solely based on the shading observed in one image.

Our approach to the estimation of light-sources is motivated by the reconstruction problem (i.e. the image based recovery of shape and other properties of viewed objects). The reconstruction process can be viewed as the *inverse* of the physical process that generates the image. The physical process, that is simulated by Computer Graphics systems, consists of an interplay of the following ingredients which we have to model.

- **Light-source (Emitter):** We introduce the general model of a point light-source in section 4.
- **Surface (Reflector):** We assume a known geometry (objects such as cubes, cylinders or spheres) with homogenous Lambertian reflectance and convex shape to exclude inter-reflections.
- **Camera (Receiver):** The model of a perspective camera is employed and pixels are converted into radiance values using the inverse response function determined in radiometric calibration.

Assume that we have a camera that is geometrically and photometrically calibrated¹. For the process of geometric calibration, we refer the reader to [2]. Now, assume further, that we have a single image of a known object which has a homogeneous surface of Lambertian reflectance. Then it is conceivable that the set of all possible light-sources that can generate the viewed image is severely restricted. Intuitively, it should be possible to determine the direction of the incident light, provided one observes the reflection on at least three points with linearly independent surface-orientations. As will be demonstrated, this is indeed the case. Beyond this intuition we will demonstrate that apart from the direction of the source, other parameters (e.g. its distance) can be estimated too.

In the following, we will assume that the illumination can be viewed as originating from one point-like source. We give the following reasons to justify this assumption:

- Many sources can be approximated as point-like if their distance to the viewed object is much greater than their extensions.
- A reflector can be approximated as a point light-source with a virtual origin lying behind the actual reflector.
- The effect of arbitrary complex background illumination can be eliminated by subtracting images taken with the source to be estimated switched on and off.²

In order to check the result of the estimation with the reality, we simulated the image forming process and compared images rendered using the virtual source, object and camera with the real images.

¹As was pointed out in [7], any camera can be calibrated photometrically provided the pixel-exposure can be varied.

²Note that in order to subtract, we first need to convert pixel values into radiance values using the inverse response function of the camera.

4 A General Point Light-Source Model

4.1 The assumptions

- Point light-source: Light originates from one point.
- Transparent media: No significant attenuation of light by the media (air).
- Direct illumination: The change of the light field from inter-reflections is neglected.
- S^1 -rotational symmetry: We assume that there exists at least one axis of rotation which leaves the distribution of radiated light invariant.

4.2 The model

Given these assumptions we can state the most general model for the light-source as a map $I : \mathbb{R}^3 \rightarrow \mathbb{R}, p \mapsto I(p)$ that assigns a radiance to each point p (in an empty space):

$$I(p) = \frac{1}{(l-p)^2} \rho(s(p), z) \quad (1)$$

With $s(p) = \langle \frac{p-l}{|p-l|}, d \rangle \in \mathbb{R}$. Where $l \in \mathbb{R}^3$ represents the position of the source, $d \in \mathbb{R}^3$ the direction and $z \in \mathbb{R}^n$ parameters of the angular dependent strength ρ . (For a multi-spectral source, one would have to replace I and ρ by functions depending also on the wavelength.)

For later reference, we summarise the (first) derivatives of I with respect to the model parameters l, z and d :

$$\frac{1}{I} \frac{\partial I}{\partial l_i} = 2 \frac{(p-l)_i}{|p-l|^2} + \frac{1}{\rho} \frac{\partial \rho}{\partial s} \frac{\langle d, p-l \rangle (p-l)_i - (p-l)^2 d_i}{|p-l|^3} \quad (2)$$

$$\frac{1}{I} \frac{\partial I}{\partial z_j} = \frac{1}{\rho} \frac{\partial \rho}{\partial z_j} \quad (3)$$

$$\frac{1}{I} \frac{\partial I}{\partial d_i} = \frac{1}{\rho} \frac{\partial \rho}{\partial s} \frac{(p-l)_i}{|p-l|} \quad (4)$$

4.3 Proof of the claims

We have to convince ourselves that using the above assumptions any model is of the stated form. For this purpose, imagine a sphere of radius 1 containing the light-source as its centre (say the origin). Since we assumed a S^1 -symmetry, we can find a direction $d \in S^2$ with the property $I|_{S^2}(p) = f(\angle(d, p))$ (with some function f) and since $\langle d, p \rangle = |d| \cos(\angle(d, p))$ we have proved that $I|_{S^2}$ is of the form claimed.

Now consider a surface patch $d\Omega$ on S^2 . Assuming that the media is transparent, we can deduce that the power radiated in the direction is independent of the distance while the size of the patch increases quadratically³. This proves the generality of the stated form, and the derivatives follow by standard differentiation.

³The surface area of a sphere is $\int_{rS^2} d\Omega = r^2 \int_{S^2} d\Omega = 4\pi r^2$

4.4 Examples

In this paper we restrict ourselves to model sources with $n = 1$ and an angular dependency of the form

$$\rho(s, z) = \begin{cases} z + s & \text{if } |s| < s_0 \\ 0 & \text{else} \end{cases}$$

(where $s_0 \in \mathbb{R}$ such that $\rho(s, z) \geq 0$ is ensured)⁴ which includes two important cases:

- Isotropic source: $d = 0$, ρ is independent of the direction.
- A simple spot light-source: $d \neq 0$, points in the direction of maximal intensity and encodes the extent of the spot-nature in its absolute length. s_0 limits the angle in which the spot radiates (or can be approximated by this simple angular dependence)⁵.

5 Parameter Estimation from Shading

In the following two sections we gather the information needed to estimate the parameters of the point light-source solely from shading.

5.1 Model evaluation

Given any guess of the parameters, we need to be able to evaluate the model and predict the measurements. As stated earlier, we assumed the camera to be fully calibrated and the object to be known. Assuming Lambertian reflectance we can then predict the irradiance m observed at a given point $w \in \mathbb{R}^2$ in the image to be

$$m = I(p) \left\langle \frac{l-p}{|l-p|}, n \right\rangle \quad (5)$$

where p is projected by the camera to w and n is the (outward) normal of the object at point p (where it is differentiable).

5.2 Dependency on changes in the parameters

We can use the derivatives derived in 4.2 to determine the overall change of the prediction:

$$\begin{aligned} dm &= d \left(I(p) \left\langle \frac{l-p}{|l-p|}, n \right\rangle \right) \\ &= I(p) \left(\left\langle \frac{l-p}{|l-p|}, n \right\rangle \frac{dI}{I}(p) + \sum_i \frac{(l-p)^2 n_i - \langle l-p, n \rangle (l-p)_i}{|l-p|^3} dl_i \right) \end{aligned} \quad (6)$$

Here m is viewed as function of the model parameters l, d and z and hence $dm = \sum_i (\frac{\partial m}{\partial l_i} dl_i + \frac{\partial m}{\partial d_i} d(d_i)) + \sum_j \frac{\partial m}{\partial z_j} dz_j$ and $\frac{dI}{I}$ has been calculated in section 4.2.

⁴Note that ρ can have discontinuities for $s = \pm s_0$ and that we have to ensure that the estimation is restricted to values of non-vanishing ρ .

⁵Note, that s has its maximum value in the direction of the spot given by d . For other directions s declines as cosine of the angle multiplied by the constant $|d|$

5.3 Gauss-Newton method

If one assumes normally distributed errors in the measurements, m the optimal parameters are estimated as solution of a non-linear least square problem. We solve this problem using the Gauss-Newton method which is comparably fast and converges rapidly (as our experiments confirm). The merit of this method to our problem was also tested independently by data simulated with Gaussian noise.

Besides the estimated optimum parameter combination, the method also delivers χ^2 which we use in the standard way to estimate the agreement of the data to our model.

6 The Estimation Process

In this section we describe the full estimation process in some detail. The input for our method are two images of a known object taken of the identical scene, one with the point light-source switched on and one with it turned off. Figure 1 demonstrates the process we outline in the following.




Figure 1: Estimation process: Shown (left to right) are the effective one source image, an image indicating the geometric calibration as well as the area where data is acquired for the estimation, and finally the synthetic image rendered using the estimated parameters.

- **Effective one source image:** In order to eliminate effects of possibly complex background illumination, we use the *inverse response function* [7] to transform the input images into radiance images which we can then subtract from each other to give an effective one source image.
- **Geometric calibration:** To proceed further, we need to find the pose of the known convex Lambertian object viewed. In this paper we use the model of a cube which naturally defines a coordinate system: we define the origin to be the inner corner viewed and use the three neighbouring corners to introduce the unit base vectors: The first unit vector as leading from the origin along the edge to the top-right corner, the second unit vector as leading to the top-left corner and finally the third unit vector as leading to the bottom corner. Note that here we introduced a length scale. The calibration is performed by manually clicking at the seven visible corners and the initial calibration can be refined using *snakes* [1]. The perspective projection matrix can be calculated from the corresponding points in the real world and in the image. At this stage, we know how to project our model into the image geometrically but do not yet know about the correct illumination.

Object	Light-model	direction	distance	$\chi^2_{\text{normalised}}$
cube (view 1)	infinite distance	(0.923, 0.224, 0.314)		8.9
	finite distance	(0.907, 0.241, 0.345)	6.84	1.4

Table 1: Comparison of the infinite distance estimation with the finite distance estimation: $\chi^2_{\text{normalised}}$ is expected to be approximately one. Clearly the finite distance model is much more consistent with the data.

- **Sampling of radiance values:** Using the known pose of the cube we can record the radiance value at each pixel and associate it with the corresponding surface-point and -normal. As indicated in Figure 1, we also use the geometric calibration to segment neighbourhoods of edges in the image. We exclude these regions from the light estimation to avoid errors due to imperfections of the real cube near the edges⁶.

Note that each non-excluded pixel on the cube delivers data which can be used for the light estimation and we have a large data set.

- **Light estimation:** Our aim is to estimate the parameters of an isotropic point light-source. To initialise the non-linear finite distance model we first solve the *linear* least square problem for the infinite-distance light model. Finally the Gauss-Newton method is used to estimate the parameters. Table 1 demonstrates the improvements achieved using the finite distance model. Note that the direction estimated in the finite distance model is different from the one in the infinite model.
- **Comparison with the original image:** As is illustrated in Figure 1, we can now proceed to render a synthetic version (using ray-tracing) and compare it with the actual input image⁷. Figure 2 highlights the differences between original and synthetic images and the histogram quantifies the quality of the estimation.

7 Experimental Results

7.1 Experimental set-up

We use an inexpensive desktop halogen lamp (20W) as the light-source. As object of known geometry, we use two cubes which were covered with paper of the same type to approximate homogenous Lambertian reflectance. The approximate dimensions of the cubes are 9.5cm (referred to as *cube*) and 13cm (referred to as *CUBE*). The cubes are in turn placed on a turn-table with a dark surface to avoid inter-reflections. (We used the turn-table later (section 7.3) for a consistency check of positions estimated at different rotations.) As recording system, we used the video-camera (Sony DC-77RR-CE and amplifier XC-77RR-CE) which can produce a gamma-corrected signal so that we did not need to calibrate the camera radiometrically, although this could be easily done as described in [7]. We estimated the uncertainty of the pixel-values, which is assumed in our modeling) empirically to approximately 8 (of 256).

⁶Recall that the normal field has discontinuities at the edges.

⁷More precisely with the effective one-source radiance image.




Figure 2: Comparison of original and synthetic images: The first image shows the absolute value of radiance differences multiplied by a factor of 30. In the ideal case the cube should be black. The second image shows the radiance differences multiplied by a factor of 20 and offset by the radiance value corresponding to the displayed square. In the ideal case the cube should have the intensity of the square. The histogram shows the distribution of the differences. The x-axis corresponds to the radiance difference and is normalised such that a value of 1 corresponds to the maximum radiance observed.

Object	View	$\chi^2_{\text{normalised}}$	l	$\frac{l}{ l }$	$ l $	ρ
<i>cube</i>	1	1.40	(-6.2, -1.7, -2.4)	(0.91, 0.24, 0.35)	6.8	$2.7e + 04$
	2	1.02	(-6.0, -3.4, -2.6)	(0.81, 0.47, 0.36)	7.4	$3.0e + 04$
	3	0.95	(-4.0, -4.3, -2.3)	(0.63, 0.68, 0.37)	6.3	$2.2e + 04$
	4	0.97	(-2.2, -4.3, -1.9)	(0.42, 0.84, 0.36)	5.2	$1.6e + 04$
<i>CUBE</i>	1	3.38	(-4.0, -0.9, -1.5)	(0.92, 0.21, 0.34)	4.4	$1.3e + 04$
	2	3.11	(-3.3, -1.8, -1.4)	(0.82, 0.46, 0.35)	4.0	$1.1e + 04$
	3	3.55	(-2.1, -2.4, -1.2)	(0.63, 0.70, 0.35)	3.4	$0.9e + 04$
	4	3.62	(-1.3, -2.7, -1.1)	(0.39, 0.85, 0.34)	3.2	$0.8e + 04$

Table 2: Results of the estimation for the various images. $\chi^2_{\text{normalised}}$ is expected to be approximately one. Note that the variations in $|l|$ are much larger than variations in $\frac{l}{|l|}$ and hence the strength ρ , which is correlated to $|l|^2$ cannot be estimated very accurately.

7.2 Estimation of parameters and comparison with original images

We performed the analysis outlined in section 6 for four different views of both cubes. Each time the turn-table was rotated approximately 15° . The full estimation process (coded in C++) takes about 10 seconds per image on a Pentium-III 900MHz machine without any particular optimisation. The results are displayed in Table 2⁸ and Figures 3, 4. Note that the histograms confirm our assumption on nearly Gaussian errors and the small value of around 2% for the standard deviation means that the model predicts the intensities successfully.

⁸The larger value of $\chi^2_{\text{normalised}}$ and originates possibly from errors in the non-perfectly flat surfaces of *CUBE*.




Figure 3: Estimation Results (*cube*): The first row shows the original images of the cube from the four different views. In the second row, the differences of synthetic and real radiance values are displayed (multiplied by a factor of 20 and offset by the radiance of the square area). The histograms show the distribution of radiance differences along with the total number of pixels that were evaluated (N), the mean (μ) and the standard deviation (σ).




Figure 4: Estimation Results (*CUBE*): (Images as in Figure 3 but for *CUBE*). The structure that can be seen in the subtraction images (second row) and the presence of *tails* in some of the histograms (third row) indicate that the surfaces are not perfectly flat. This interpretation is in agreement with the larger values for $\chi^2_{\text{normalised}}$ in Table 2.

<i>cube</i>			<i>CUBE</i>		
$-l'$	$\frac{-l'}{ l' }$	$ l' $	$-l'$	$\frac{-l'}{ l' }$	$ l' $
(4.0, 0.9, 1.5)	(0.92, 0.21, 0.34)	4.4	(6.2, 1.7, 2.4)	(0.91, 0.24, 0.35)	6.8
(4.2, 1.0, 1.6)	(0.92, 0.21, 0.34)	4.5	(6.8, 1.8, 2.8)	(0.90, 0.24, 0.38)	7.5
(3.2, 0.6, 1.2)	(0.92, 0.18, 0.34)	3.5	(6.1, 2.0, 2.5)	(0.89, 0.29, 0.36)	6.9
(2.8, 0.5, 1.0)	(0.93, 0.18, 0.33)	3.1	(4.8, 1.5, 2.0)	(0.89, 0.28, 0.36)	5.5
μ	(0.92, 0.20, 0.34)	3.9	μ	(0.90, 0.26, 0.36)	6.7
σ	(0.004, 0.02, 0.01)	0.6	σ	(0.007, 0.02, 0.01)	0.8

Table 3: Agreement of the estimations in different views: Light positions l estimated for different views are transformed into the same coordinate system and denoted by l' . Values in each column should agree. In the lower half, mean and standard deviation of the above quantities are displayed.

7.3 Rotational consistency check and accuracy

In this section we investigate if the estimations performed independently for each view agree. Moving the object in front of the fixed light-source is equivalent to moving the light-source with the object fixed. To that extent we can use the projection matrices to extract the relative motion and this should enable us to transform all estimates into a common coordinate system in which the position of the light-source should be invariant. First, from each projection matrix, we can extract a rigid-body transformation⁹:

$$\phi : \mathbb{R}^3 \rightarrow \mathbb{R}^3, x \mapsto Rx + T \text{ with } R \in O(3), T \in \mathbb{R}^3$$

Next, we can express the relative motion as $\phi_{ij} = \phi_i^{-1} \circ \phi_j$ which transforms coordinates of view j into coordinates of view i . Table 3 shows the result of transforming all light-estimations into the coordinates of the first image.

Furthermore, we can combine the estimations to get an improved estimate on the light-source and to estimate the uncertainty (table 3). Note that as anticipated, the uncertainty in distance is much larger than in orientation.

8 Future Work and Conclusion

We aim to extend the analysis done using cubes to other objects like cylinders and spheres and also include non-homogenous and non-Lambertian objects. In a further step, we plan to implement the estimation to non-isotropic sources as introduced in section 4.2. Finally, it might be possible to extend the estimation to situations where more than one but a small number of sources are to be estimated.

We have presented a model for point light-sources which allowed us to successfully recover the position of a source and its parameters. As we have demonstrated, it is possible and worthwhile for many applications to work with a finite distance model. The method is practical in that it does not demand any special imaging system or environment and just one image with a known object is sufficient for the estimation. The direct comparison of real images with images rendered underline the power of the method presented.

⁹We use the standard method which corresponds essentially to a QR-decomposition.

References

- [1] A. Blake and M. Isard. *Active Contours*. Springer, New York, 1998.
- [2] R. Cipolla and P. Giblin. *Visual Motion of Curves and Surfaces*. Cambridge University Press, Cambridge, 2000.
- [3] B.K.P. Horn and M.J. Brooks. The variational approach to shape from shading. *Computer Vision, Graphics, and Image Processing*, 33(2):174–208, 1986.
- [4] K. Ikeuchi and B.K.P. Horn. Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17:141–184, 1981.
- [5] C. Loscos, M. Frasson, G. Drettakis, B. Walter, X. Granier, and P. Poulin. Interactive virtual relighting and remodeling of real scenes. In *Rendering Techniques '99 (10th EG workshop on Rendering)*. Springer, Wien, June 1999.
- [6] S.R. Marschner and D.P. Greenberg. Inverse lighting for photography. In *Proceedings of IS&T/SID Fifth Color Imaging Conference*, pages 262–265, November 1997.
- [7] S.K. Nayar and T. Mitsunaga. Radiometric self calibration. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, I:374–380, June 1999.
- [8] A.P. Pentland. Finding the illumination direction. *J. Opt. So. Am.*, 72(4):448–455, 1982.
- [9] I. Sato, Y. Sato, and K. Ikeuchi. Illumination distribution from brightness in shadows: Adaptive estimation of illumination distribution with unknown reflectance properties in shadow regions. In *Proceedings of IEEE ICCV'99*, volume 2, pages 875–883, 1999.
- [10] P. Sinha and E.H. Adelson. Recovering reflectance and illumination in a world of painted polyhedra. *Proc 4th Int. Conf. on Comp. Vis.*, pages 156–163, 1993.
- [11] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.
- [12] R.J. Woodham. Gradient and curvature from the photometric-stereo method, including local confidence estimation. *JOSA*, 11:3050–3068, 1994.
- [13] P.L. Worthington and E.R. Hancock. New constraints on data-closeness and needle map consistency for shape-from-shading. *IEEE-PAMI*, 21:1250–1267, 1999.
- [14] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. *SIGGRAPH99 conference proceedings*, pages 215–224, 1999.