

Novel Depth Cues from Uncalibrated Near-field Lighting

Sanjeev J. Koppal and Srinivasa G. Narasimhan

Robotics Institute, Carnegie Mellon University, Pittsburgh, USA

Email: (koppal,srinivas)@ri.cmu.edu

Abstract

We present the first method to compute depth cues from images taken solely under uncalibrated near point lighting. A stationary scene is illuminated by a point source that is moved approximately along a line or in a plane. We observe the brightness profile at each pixel and demonstrate how to obtain three novel cues: plane-scene intersections, depth ordering and mirror symmetries. These cues are defined with respect to the line/plane in which the light source moves, and not the camera viewpoint. Plane-Scene Intersections are detected by finding those scene points that are closest to the light source path at some time instance. Depth Ordering for scenes with homogeneous BRDFs is obtained by sorting pixels according to their shortest distances from a plane containing the light source. Mirror Symmetry pairs for scenes with homogeneous BRDFs are detected by reflecting scene points across a plane in which the light source moves. We show analytic results for Lambertian objects and demonstrate empirical evidence for a variety of other BRDFs.

1. Distant vs. Near Lighting

Many vision algorithms assume that the scene is illuminated by light sources that are far away, such as the sun or the sky. In this setting, the incident illumination angles at each scene point are identical, allowing easier estimation of scene properties. Approaches that exploit this fact to recover scene normals and albedos include the classical photometric stereo technique ([28]) and several extensions for non-lambertian low parameter BRDFs (Dichromatic ([11]), Diffuse + Specular ([4], [23],[15],[24]), Microfacet ([17],[27])).

Although these algorithms are popular, there are many scenarios where the distant lighting assumption is invalid. These include indoor, underground and underwater scenes, as well as outdoor scenes under night lighting. Scene analysis using near-field lighting is difficult since the incident angles at each scene point are different. However, the advantage is that the intensity fall-off in distance from the near light source is encoded in the images. Therefore it becomes possible to extract both depths and normals of the scene, and many methods have shown this (Shape from shading ([20], [19], [9]), Photometric stereo ([8], [3], [2], [10]), Calibrated camera-light source pairs ([15], [7]), Helmholtz stereopsis ([29])). Exploiting the intensity fall-off requires the 3D position of the source and therefore most of these techniques assume calibrated lighting. Instead, if the near light source is *uncalibrated*, can we still extract useful information about scene depth?

Let a near point light source move along a line or in a plane. Our key idea is to use this light source path as a 'baseline' with respect to which scene depth cues can be extracted. Analysis of the intensities at each pixel yields the following depth cues:

Plane-Scene Intersections: Let a lambertian scene be illuminated by a point source moving along a line. By simply detecting intensity maxima at a scene point, we produce plane-scene intersections similar to those obtained by sweeping a sheet of light over the scene. Intersections with planes of other orientations are produced by changing the direction of the light source path.

Depth Ordering: A point light source moving in a plane illuminates a homogeneous lambertian scene under orthographic viewing. Under certain conditions, it is possible to obtain an ordering of the scene in terms of perpendicular distances to this 'base plane'. This is done simply by integrating the measured intensities at scene points over time.

Mirror Symmetries: An orthographic camera views a lambertian scene illuminated by a light source moving along a line. We detect mirror symmetries by reflecting scene points across a plane containing the light source path. These symmetric pairs will observe identical incident angles and light source distances and are detected by simply matching scene points with the same brightnesses over time.

In practice, we relax many of the restricting assumptions described above. First, we do not require the source to move precisely along a line or in a plane. Instead, a user can simply move the source with his/her hand, making image acquisition easy. Second, simple heuristics can be used to extend all of our results to non-Lambertian BRDF's containing diffuse and specular components. We show strong supporting evidence using both rendered and real scenes with complex geometry and material properties.

Our approach computes depth information by moving the light source, complementing traditional stereo methods that exploit changes in viewpoint. An interesting similarity to camera-based methods is that increasing the length of the source path is similar to increasing stereo baseline and results in more reliable depth cues. One difference, however, is that our cues are defined with respect to the light source path. This allows us to obtain depth cues from different 'viewpoints', without moving the camera. These novel scene visualizations provide strong constraints for any depth estimation technique, while avoiding the correspondence problem. Our approach is simple, requiring neither geometric/radiometric calibration nor a complex setup. Therefore we believe our depth cues have broad application in computer vision.

2. Detecting Plane-Scene Intersections

Consider a scene with unknown geometry. Let it be illuminated by an isotropic point light source moving in a straight line (see Figure 1) at constant speed. At a line position $d = i$, the source S is nearest to scene points lying on a plane perpendicular to the light source path. In this section, we show that maxima will occur in the brightness profiles of all of these scene points, owing to a minimum in the inverse-square fall-off from the point light source. By simply detecting brightness maxima occurring at every light source position (or frame), we obtain plane-scene intersections. Our results appear to be created by intersecting the scene with a plane that is translating along the source path. This is similar in spirit to structured light striping where a sheet of light is swept across the scene. In the case of structured light striping, brightness maxima are detected by simply thresholding each image. The difference here is that we obtain the same effect using an uncalibrated isotropic near point light source.

2.1. Maxima in Brightness Profiles

Let us assume, without loss of generality, that the light source has unit intensity. If the BRDF is given by B , the foreshortening by f , the incident light angles by θ_s and ϕ_s , the viewing angles given by θ_v and ϕ_v , and the distance between the light source and the scene point at line position d by $R(d)$, then the brightness profile $E(d)$ at a scene point is:

$$E(d) = \frac{B(\theta_s(d), \phi_s(d), \theta_v, \phi_v) \cdot f(d)}{(R(d))^2} = \frac{F(d)}{(R(d))^2} \quad (1)$$

where the F term contains both BRDF and foreshortening. Taking the derivative with respect to position, d , gives us an expression for when the maxima of $E(d)$ should occur,

$$E'(d) = \frac{(R(d))^2 F'(d) - 2R'(d)R(d)F(d)}{(R(d))^4} = 0 \quad (2)$$

In the next sub-section we will investigate the maxima of $E(d)$ that occur when the light source is closest to a scene point. We call these *iso-planar* maxima since they indicate scene points that are located on the same plane (Figure 1). We assume that the scene is Lambertian and that the distance between the light source path and the scene point is small.

2.2. Plane-Scene Intersections for a Lambertian Scene

Without loss of generality, consider a scene point located at the origin, $P = (0, 0, 0)$, and let the light source move along a line parallel to the z -axis, and let its 3D position be $\vec{S}(d) = (D, 0, d)$ where D is its closest distance to the origin. The




Figure 1. **Detecting a Plane-Scene Intersection using Brightness Maxima:** Consider a light source moving in a line in front of a scene. At any line position $d = i$, the scene points that are closest to the light source path will display a maxima in their profiles. We use this maxima to create plane-scene intersections, similar to structured light results created by intersecting a sheet of light with the scene.




Figure 2. **Plane-scene intersections on a Lambertian Pot:** In this figure we show results of our method applied to a Lambertian clay pot. We obtain both horizontal and vertical plane intersections by moving a near point light source as shown by the marked paths. These results are similar to those obtained from structured light. The discontinuity in the result for horizontal planes is due to merging the intersections for the right and left halves of the pot. These were done separately to avoid blocking the camera view. Please see [13] for better visualization.

distance from the scene point to the light source is $R(d) = (D^2 + d^2)^{0.5}$. For a scene point with albedo ρ and surface normal $\vec{n} = (n_x, n_y, n_z)$ we write:

$$F(d) = \frac{\rho \vec{n} \cdot (\vec{S}(d) - \vec{P})}{R(d)} \quad (3)$$

An iso-planar maxima occurs if $F'(d)$ is zero when the light source reaches the closest distance to the scene point. Consider the expressions for $F'(d)$ and $R'(d)$,

$$F'(d) = \frac{\rho R(d) \cdot (n_z) - \rho R'(d) (Dn_x + n_z d)}{(R(d))^2} \quad (4)$$

$$R'(d) = \frac{d}{R(d)} \quad (5)$$

Putting these into $E'(d)$,

$$E'(d) = \frac{-\rho(2n_z d^2 + 3Dn_x d - D^2 n_z)}{(R(d))^5} \quad (6)$$

Setting $E'(d) = 0$ gives us a quadratic equation in d with two possible solutions. By checking the sign of the second derivative, $E''(d)$, for all possible normals, we found that one of these solutions is always a maxima and is given by,

$$d_{maxima} = D \left(\frac{-\frac{3}{4} + \sqrt{(\frac{3}{4})^2 + \frac{1}{2}(\frac{n_z}{n_x})^2}}{\frac{n_z}{n_x}} \right) \quad (7)$$

Ideally, iso-planar maxima should occur when $d = 0$, since that is when the light source is closest to the scene point and $R(0) = D$. Therefore, Equation 7 represents the error in the iso-planar maxima location. However, simulations show that this error remains bounded as $\frac{n_z}{n_x}$ is varied. For instance, d_{maxima} is zero when $\frac{n_z}{n_x} \rightarrow 0$ and becomes $\frac{D}{\sqrt{2}}$ when $\frac{n_z}{n_x} \rightarrow \infty$. Although we will investigate this bound more thoroughly in future work, here we assume the error is negligible if the light source path is close to the scene point (D is small).

In Figure 2 we show both horizontal and vertical plane-scene intersections for a lambertian pot, created by moving the light source first sideways and then upwards. We code the planes from blue to red, obtaining a continuum of color coded plane-scene intersections. For the second result, we merged two experiments for the left and right halves of the pot, to avoid blocking camera’s view of the scene. Although these results appear similar to structured light images obtained by sweeping a plane over the scene, they were obtained by a user hand-waving a near point light.

Shadows and Specularities: Non-isoplanar brightness maxima are usually rare in lambertian scenes illuminated by a near light source. In fact, for a maxima to occur when $R'(d) \neq 0$, a fortuitous occurrence of values for $R'(d)$, $R(d)$, $F(d)$ and $F'(d)$ would be required in Equation 2, which is less likely. However, scenes with sharp specularities still show non-isoplanar maxima. Fortunately, since specularities are characterized by a rapid rise in brightness, these are detected by thresholding the second derivative of the measured intensities. Glossy highlights are similarly removed using a method proposed by [16]. Shadows can eliminate iso-planar maxima from a profile, but are handled by repeating the experiment with different parallel light source paths. Since a scene point shadowed in one experiment may be illuminated in another, we usually detect an iso-planar maxima. Finally, many scenes may still exhibit non-isoplanar maxima due to material properties. We remove these by simply enforcing neighboring scene points to have iso-planar maxima that occur closely in time.

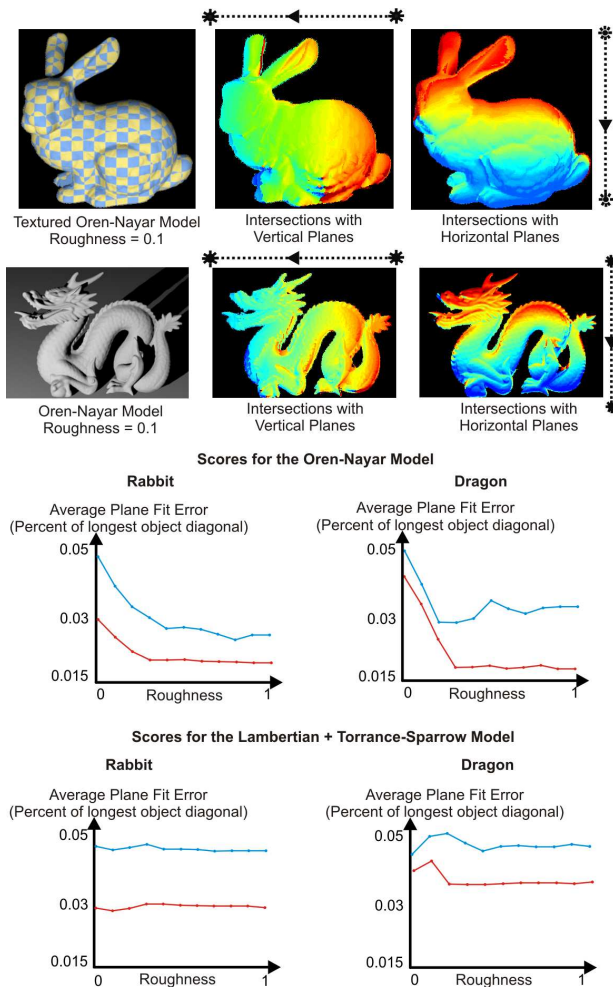


Figure 3. **Plane-Scene Intersections for Rendered Scenes:** We used a ray tracing software ([18]) to render two objects (dragon and bunny) using both the Lambertian + Torrance-Sparrow and the Oren-Nayar models. For each set of images, we created plane-scene intersections, and we show two examples of these at the top of the figure. We fit a plane to each intersected region and measure the plane fit error. The average of these errors is plotted, as a percent of the longest distance contained in the object. These empirical results support the idea that the plane-intersection algorithm can be used with a variety of non-lambertian scenes.

2.3. Experimental Results

It is possible to create iso-planar maxima in profiles if the derivative of the BRDF, $F'(d)$, is zero around a small interval where $R'(d) = 0$. In other words, the change in BRDF must be negligible for a short interval around the maxima location. This seems a reasonable assumption for diffuse BRDFs. In Figure 3, we show empirical evidence supporting this idea for scenes created using a ray-tracing tool ([18]). We rendered the bunny and dragon models by varying the parameters of the Oren-Nayar model (facet angle from 0 to $\frac{\pi}{2}$) and of a Lambertian + Torrance-Sparrow model (σ of facet distribution from 0 to 1). The light source was moved along the




Figure 4. **Plane-Scene Intersections for Real Scenes:** At the top of the figure we show horizontal and vertical plane-scene intersections for a painted house, even though this scene demonstrates glossy specularities. We also show plane-scene intersections for an office desk made of metal and plastic. At the bottom we show the sub-surface scattering effects of a wax candle, by shining a point laser. Our method is able to create horizontal plane-scene intersections for this object with complex BSSRDF ([6]) appearance.

x - and then y - axes to create horizontal and vertical plane-scene intersections. Using ground-truth, we fit a plane to the 3D locations of scene points in the plane-scene intersections and plotted the sum-of-squared errors. The low errors indicate the robustness of our technique for non-lambertian BRDFs.

At the top of Figure 4, we show horizontal and vertical plane-scene intersections for a painted house model, which are detected despite glossy specularities in the scene. Our algorithm also produces good results for scenes with sharp specularities, such as the metal office desk in Figure 4 and for rough objects with cracks, such as an earthen pot in Figure 4 (see [13] for the input sequence and a better visualization of the plane-scene intersections). We are also able to create plane-scene intersections for objects with sub-surface scattering (since BSSRDF ([6]) is smooth) such as a wax candle shown in Figure 4. The scattering effects are demonstrated using a laser pointer, and horizontal plane-scene intersections are shown. Such a result would be hard to obtain using traditional structured light methods.




Figure 5. **Overlap of incident angles for a simple 2D scene:** A light source moving in a line illuminates a 2D scene with homogeneous BRDF. Let scene point P be at a greater depth than Q and let both have the same surface normal. Some of the incident angles at P are repeated in Q , and we show these *overlapped angles* in gray. Because of inverse-square fall-off, the measured intensities in the overlapped region will be higher in Q than in P . We extend this idea to create scene depth ordering for 3D scenes as well.

3. Depth Ordering for Homogeneous Scenes

Consider a scene with homogeneous BRDF viewed by an orthographic camera. This scene is illuminated by a point source moving in a plane. Intuitively, scene points closer to the light source tend to be brighter than scene points further away. If we somehow manage to remove the effect of BRDF for any two scene points, then their appearance would depend only on their distances to the light source. Although this is not possible generally, we describe certain conditions under which we obtain scene depth ordering with respect to the “base plane” containing the light source path. We will provide two heuristics to achieve good results for depth ordering: **a)** Move the light source over a large area of the base plane and **b)** Bring the base plane as close as possible to the scene. We support our method with strong empirical evidence using both simulations and real scenes.

3.1. Integrating the Brightness Profile

Consider a scene point P illuminated by a light source moving in a plane. As before, we will assume the light source has unit intensity and let F contain both the BRDF and foreshortening. If $R(d)$ is the distance between the source and this scene point at position d on the light source path, then the measured intensity profile of P is given as:

$$E_p(d) = \frac{F_p(d)}{(R_p(d))^2}. \quad (8)$$

Let S_p be the sum of the P 's intensities along the light source path from positions d_1 to d_n ,

$$S_p = \sum_{d=d_1}^{d_n} \frac{F_p(d)}{(R_p(d))^2}. \quad (9)$$

Let there be another point Q whose perpendicular distance to the plane containing the light source is less than that of P .

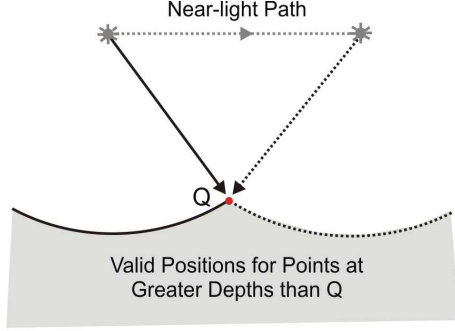


Figure 6. **Light Source Moves in a Path with Finite Length:** A light source moving in a line illuminates a scene with homogeneous BRDF. Our depth ordering method works for scenes where any point at a greater depth than Q must be further away from the light source at every time instance. This restricts us to a non-planar region around Q as shown.

At different light source positions, P and Q may observe identical incident angles. We term these as *overlapped* incident angles, and they are shown in gray in Figure 5. In this simple 2D scene, the longer the length of the light source path, the greater the overlapped region. For an infinite path, all the incident angles at P and Q will overlap. Since the inverse-square fall-off is smaller at Q than at P , $S_p < S_q$.

This result is harder to demonstrate for real 3D scenes where P and Q can have different surface normals. To make the analysis easier, we assume that the light source is *always* further from P than Q . This assumption of $R_p(d) > R_q(d)$ for all positions along the light source path is illustrated in Figure 6. We will show that this drawback is not severe since, in practice, we get good depth ordering results for scenes with a variety of geometries.

S_p and S_q consist of two components, one of which (denoted by O) comes from overlapped incident angles, as in the gray region in Figure 5. The other comes from non-overlapped or different incident angles (denoted by N). Since the order of the summation does not matter, we let the overlapped angles occur from d_1 to d_i , and the non-overlapped from d_{i+1} to d_n . We separate the summation accordingly as:

$$S_p = \sum_{d=d_1}^{d_i} \frac{F_p(d)}{(R_p(d))^2} + \sum_{d=d_{i+1}}^{d_n} \frac{F_p(d)}{(R_p(d))^2} \quad (10)$$

which we write concisely as:

$$S_p = O_p + N_p \quad (11)$$

Similarly for scene point Q we have $S_q = O_q + N_q$. The overlapped terms, O_p and O_q , have the same incident angles. Since P is further away than Q , the inverse-square fall-off from the light source ensures that $O_p < O_q$. Therefore, $S_p < S_q$ if $N_p < N_q$. Consider a pair of the summation terms from N_p and N_q . If the inequality holds for each of these pairs, then,

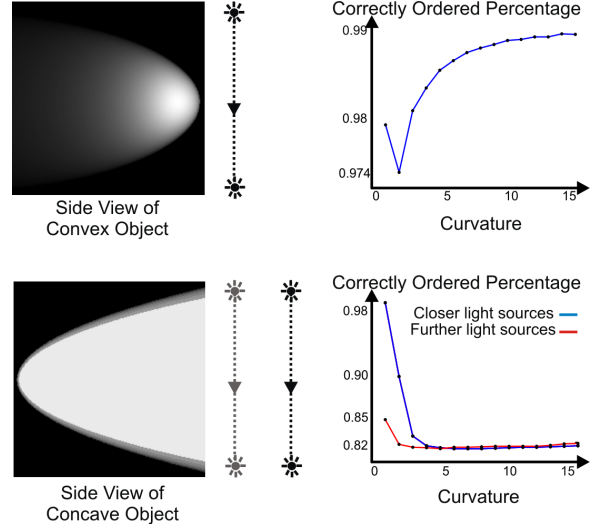


Figure 7. **Simulations of convex and concave Lambertian objects:** Our depth ordering method always works for convex lambertian objects. We show results with very low errors on renderings of a convex shape whose curvature we gradually increase. We also render concave objects at different curvatures. Although the depth orderings are much worse in this case, they stabilize and do not degrade with time. In addition, bringing the light source closer creates better orderings for shallow concave objects.

$$\frac{F_p(d)}{(R_p(d))^2} \leq \frac{F_q(d)}{(R_q(d))^2} \quad (12)$$

3.2. Depth Ordering for Lambertian Scenes

When does Equation 12 hold? We answer this question by first investigating depth ordering for lambertian scenes. We divide the problem into two cases, based on whether the local geometry between the scene points P and Q is convex or concave.

Convex neighborhood: If the normals at P and Q are given by \mathbf{n}_p and \mathbf{n}_q and if the 3D position of the light source is $\mathbf{s}(d)$ then we can rewrite Equation 12 as,

$$\frac{\mathbf{n}_p \cdot (\mathbf{s}(d) - \mathbf{P})}{\mathbf{n}_q \cdot (\mathbf{s}(d) - \mathbf{Q})} \leq \left(\frac{R_p(d)}{R_q(d)} \right)^3 \quad (13)$$

For a convex object, a scene point's foreshortening to the light source is higher than one farther away. Therefore the LHS of Equation 13 is always less than one, and our method always works for pairs in a convex lambertian neighborhood. In Figure 7 we show simulations of depth orderings for convex objects of different curvatures demonstrating above 97% accuracy. Small errors occur due to violations of the assumption in Figure 6. We also present results on non-lambertian convex objects, such as corduroy and dull plastic (Figure 9).





Figure 8. **Depth Ordering for Rendered Scenes:** We used the PBRT software ([18]) to render three scenes using both the Lambertian + Torrance-Sparrow and Oren-Nayar models. We covered the space of roughness parameters for each of these models. Using ground truth depths for these scenes, we measured the percentage of correctly ordered scene points and we plot the results above. These results provide empirical evidence supporting the use of our method with diffuse non-lambertian BRDFs.

Concave neighborhood (shadows and interreflections):

The foreshortening of P and Q are now opposite to the convex case, and the LHS of Equation 13 is always greater than 1. In fact, the further the two points are, the larger the left hand side of the equation will be. Fortunately, our method still performs well thanks to the effect of shadows. We rewrite Equation 13 by adding a visibility term, $V(d)$,

$$V(d) \frac{\mathbf{n}_P \cdot (\mathbf{s}(d) - \mathbf{P})}{\mathbf{n}_Q \cdot (\mathbf{s}(d) - \mathbf{Q})} \leq \left(\frac{R_P(d)}{R_Q(d)} \right)^3 \quad (14)$$

Since P is further away from the light source than Q , it is more often in shadow. Whenever this happens, $V(d) = 0$ and Equation 14 is true. Concave depth ordering has more errors




Figure 9. **Depth Ordering Results on Real Convex Scenes:** We acquired images by waving a near point light in a plane. For each scenes we display the depth ordering obtained as a "pseudo depth", by plotting the ordering in 3D. The objects presented are made up of a variety of materials such as plastic and corduroy.

than the convex case (Figure 7), the results are still reasonable (above 82%). In addition, the closer we bring the light source to the scene, the better the results get since the RHS of Equation 14 increases. Finally, we note that smooth interreflections do not severely degrade performance, since the errors for concave depth ordering remain roughly constant with increase in curvature.

3.3. Experimental Results

In Figure 8, we have rendered three scenes (buddha, bunny and dragon), using ray tracing ([18]). As before, we varied the roughness parameters for the Oren-Nayar model (facet angle from 0 to $\frac{\pi}{2}$) and the Lambertian + Torrance-Sparrow model (σ of facet distribution from 0 to 50), getting 10 sets of images, each with 170 different light source locations. We compare the depth ordering with ground truth depth and show accuracy above 85% for the Oren-Nayar model and above 75% for the Lambertian + Torrance-Sparrow model. In Figure 10 we display the ordering for a polyresin statue containing shadows and interreflections as a "pseudo depth" in Maya. We have shown different views as well as its appearance under novel lighting. Finally, an advantage of our method is that these depth orderings are with respect to the light source and not the camera. We obtain orderings from different planes, without camera motion and avoiding the difficult problem of scene point correspondence. In Figure 11 we show ordering results for a wooden chair, for three different planes.

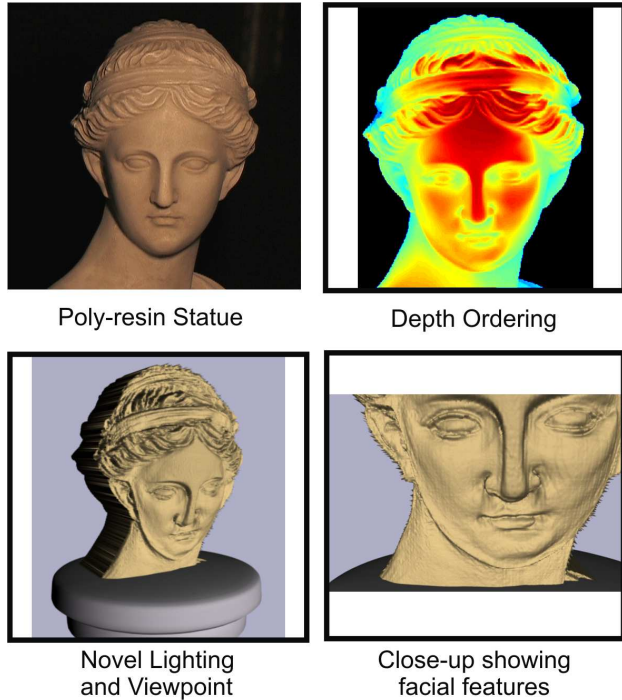


Figure 10. **Depth Ordering Results for a Statue:** In the top row we show a polyresin statue and a 2D depth ordering image obtained from an uncalibrated near light source moving in a plane. We visualize the ordering as a "pseudo depth" in the bottom row, applying novel lighting, viewpoint and BRDF. Although this object has complex effects such as interreflections and self-shadowing, our depth ordering is accurate, as we show in the close-up. Please see [13] for many other view points.

4. Mirror Symmetries in Homogeneous Scenes

In this section, we will investigate a way of finding mirror symmetry pairs for scenes with homogenous BRDFs. Although our method is simple, it is useful for objects that have the mirror symmetry in their shapes and produces very sparse pairs otherwise.

Mirror Symmetry Pairs in Lambertian Scenes Consider a point P in a homogeneous Lambertian scene containing a mirror symmetry, viewed by an orthographic camera. Let the scene be illuminated by a point light moving in a plane parallel to the viewing direction. Reflecting P and its surface normal across such a plane does not change its incident elevation angles and light source distances (Figure 12). Therefore, all the reflected pairs will have identical appearance over time and can be easily located. We show the result of matching identical profiles for a homogeneous Lambertian pot at the top of Figure 13. The pot has a reflection symmetry across a vertical plane passing through its center and perpendicular to the image plane. We wave a point light source in this plane and mark the matched scene points with the same color. Note that if the light source was moved in a different plane, different symmetry pairs would be obtained. However, these would most likely




Figure 11. **Depth Ordering from different viewpoints:** Here we show depth ordering of a wooden chair from the viewpoints of three different planes. Brighter scene points are closer than farther ones. The second depth ordering is from the camera point of view, and is similar to conventional depth results obtained from stereo algorithms. We also get the orderings from the left and right planes without moving the camera. This allows us to obtain novel visualizations of the scene, while avoiding the problem of scene point correspondence.

be sparse, since the reflection plane would not coincide with the mirror symmetry of the object.

Shadows and Inter-reflections: A scene with mirror symmetry has the same geometry across the plane of reflection. Therefore, geometric appearance effects such as shadows and inter-reflections will be identical between symmetry pairs, creating identical intensity profiles.

Isotropic BRDFs: In homogeneous non-lambertian scenes the effect of the incident azimuth angles must be considered. In general these will be different for P and its symmetry pair. However, the absolute difference between the azimuth angles are the same, and therefore our method works for isotropic BRDFs. As before, the symmetry pair of P is found by reflection across the plane containing the light source path and the orthographic viewing direction. Using the laws of symmetry we can easily show that P and its symmetric pair have the same light source distances, the same elevation angles, and identical absolute differences between the azimuth angles of viewing and illumination. We find symmetry pairs by matching identical intensities for an office desk at the bottom of Figure 13, even though this scene contains non-lambertian materials such as metal and plastic.




Figure 12. **Matching Identical Profiles to give Symmetry Pairs:** We show a light source moving in a line illuminating a scene point P . We find a symmetry pair by reflecting P across a plane containing the light source path and the orthographic viewing vector. The elevation angles (θ) and the absolute differences in azimuth angles (ϕ) are the same for the pair.




Figure 13. **Symmetry Pairs in Real Scenes:** We show two real scenes, a lambertian pot and a office desk. A near point light source was hand-waved approximately along the axis of reflection. We match identical profiles in the scene and mark them with the same color, giving us symmetric pairs in the scene.

5. Conclusions

Algorithms for obtaining scene geometric information from uncalibrated, varying lighting are well known in the case of distant light sources. For example, many photometric stereo methods have been introduced that handle unknown, distant illumination ([5], [12], [1], [26]). In addition, variation in shadows has been used extensively ([22], [14], [21]) to obtain cues such as depth edges. Recently, Takahiro and Sato ([25]) combined near and distant-lit images to obtain scene shape. However, none of these methods have been extended beyond the distant illumination scenario. This is surprising, since, intuitively, the more complex the set of rays that illuminate the scene, the richer should be the information obtained from that scene. In this paper, we have made a first step towards understanding this concept for near point light sources.

6. Acknowledgements

This research was partly supported by NSF awards #IIS-0643628 and #CCF-0541307, and ONR contracts #N00014-05-1-0188 and #N00014-06-1-0762. The authors thank Mohit Gupta and Shuntaro Yamazaki for useful technical discussions.

References

- [1] R. Basri and D. Jacobs. Photometric stereo with general, unknown lighting. *ICCV*, 2001.
- [2] J. J. Clark. Active photometric stereo. *CVPR*, 1992.
- [3] J. J. Clark. Photometric stereo with nearby planar distributed illuminants. *CRV*, 2006.
- [4] E. Coleman and R. Jain. Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry. *Intl Conf. Color in Graphics and Image Processing*, 1982.
- [5] A. Georghiadis. Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. *ICCV*, 2003.
- [6] M. Goesele, H. P. A. Lensch, J. Lang, C. Fuchs, and H. Seidel. Acquisition of translucent objects. *Siggraph*, 2004.
- [7] K. Hara, K. Nishino, and K. Ikeuchi. Determining reflectance and light position from a single image without distant illumination assumption. *ICCV*, 2003.
- [8] Y. Iwahori, R. Woodham, H. Tanaka, and N. Ishii. Moving point light source photometric stereo. *IEICE*, 1994.
- [9] S. Kao and C. Fuh. Shape from shading using near point light sources. *Image Analysis Applications and Computer Graphics*, 1995.
- [10] B. Kim and P. Burger. Depth and shape from shading using the photometric stereo method. *CVGIP*, 1999.
- [11] G. Klinker, S. Shafer, and T. Kanade. A physical approach to color image understanding. *IJCV*, 1990.
- [12] S. J. Koppal and S. G. Narasimhan. Clustering appearance for scene analysis. *CVPR*, 2006.
- [13] S. J. Koppal and S. G. Narasimhan. Novel depth cues video. <http://www.cs.cmu.edu/~koppal/depth.cues.mov>, 2007.
- [14] M. S. Langer, G. Dudek, and S. W. Zucker. Space occupancy using multiple shadowimages. *IROS*, 2005.
- [15] S. Magda, T. Zickler, D. Kriegman, and P. Belhumeur. Beyond Lambert: reconstructing surfaces with arbitrary BRDFs. *ICCV*, 2001.
- [16] S. K. Nayar, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of lambertian, specular, and hybrid surfaces using extended sources. *MIV*, 1989.
- [17] M. Oren and S. K. Nayar. Generalization of the lambertian model and implications for machine vision. *IJCV*, 1995.
- [18] M. Pharr and G. Humphreys. Physically based rendering: From theory to implementation. *Elsevier*, 2004.
- [19] E. Prados and O. Faugeras. A mathematical and algorithmic study of the lambertian SFS problem for orthographic and pinhole cameras. *Technical Report RR-5005, INRIA*, 2003.
- [20] E. Prados and O. Faugeras. Shape from shading: A well posed problem? *CVPR*, 2005.
- [21] R. Raskar, K. Tan, R. Feris, J. Yu, and M. Turk. Non-photorealistic camera: Depth edge detection and stylized rendering using multi-flash imaging. *Siggraph*, 2004.
- [22] D. Raviv, Y. Pao, and K. Loparo. Reconstruction of three-dimensional surfaces from two-dimensional binary images. *Transactions on Robotics and Automation*, 1989.
- [23] A. Shashua. On photometric issues in 3D visual recognition from a single 2D image. *IJCV*, 1997.
- [24] H. Tagare and R. deFigueiredo. A theory of photometric stereo for a class of diffuse non-lambertian surfaces. *PAMI*, 1991.
- [25] O. Takahiro and S. Yoichi. Does a nearby point light source resolve the ambiguity of shape recovery in uncalibrated photometric stereo? *IPSI*, 2007.
- [26] P. Tan, S. P. Mallick, L. Quan, D. Kriegman, and T. Zickler. Isotropy, reciprocity and the generalized bas-relief ambiguity. *CVPR*, 2007.
- [27] K. Torrance and E. Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 1967.
- [28] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 1980.
- [29] T. Zickler, P. Belhumeur, and D. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *ECCV*, 2002.