

Diffuse Structured Light

Shree K. Nayar
Columbia University
120th St, New York, NY 10027
nayar@cs.columbia.edu

Mohit Gupta
Columbia University
120th St, New York, NY 10027
mohitg@cs.columbia.edu

Abstract

Today, structured light systems are widely used in applications such as robotic assembly, visual inspection, surgery, entertainment, games and digitization of cultural heritage. Current structured light methods are faced with two serious limitations. First, they are unable to cope with scene regions that produce strong highlights due to specular reflection. Second, they cannot recover useful information for regions that lie within shadows. We observe that many structured light methods use illumination patterns that have translational symmetry, i.e., two-dimensional patterns that vary only along one of the two dimensions. We show that, for this class of patterns, diffusion of the patterns along the axis of translation can mitigate the adverse effects of specularities and shadows. We show results for two applications - 3D scanning using phase shifting of sinusoidal patterns and separation of direct and global components of light transport using high-frequency binary stripes.

1. Introduction

Whenever illumination can be controlled, a practitioner will almost always use a structured light based vision system over a passive one, for the simple reason that it recovers information more reliably. Today, structured light methods are widely used for visual inspection, object recognition, estimating material properties and recovering three-dimensional structure. In particular, they are the workhorse of the multi-billion dollar industry that is focused on machine vision for factory automation.

In virtually any real-world application, a vision system is required to handle a wide range of shapes and materials. For instance, a single image of a circuit board may include diffuse bodies of electronic components, glossy strips of copper connectors and mirror-like solder joints. When a scene that includes diffuse and specular objects is lit by a single point source, the specular reflections received by the sensor can be many orders of magnitude stronger than the diffuse reflec-

tions [15]. Such a wide range of radiance values remains a major challenge for any vision system. A common technique for capturing scenes with wide brightness ranges is exposure bracketing [2]. However, this requires acquiring multiple images at different exposures, thereby increasing the overall measurement time.

A second limitation with active vision systems pertains to shadows. Most structured light methods require an appreciable distance (baseline) between the sensor (e.g., camera) and the source (e.g., projector). This separation between the sensor and the source is necessary, for instance, to be able to compute depth via triangulation. A major downside of this requirement is that it is almost guaranteed to produce strong shadows in the captured images. In the case of depth estimation, this results in “holes” in the depth map. One way to reduce the holes is by using multiple sources so that a hole created by one may be filled by another. However, the multiple sources need to be activated sequentially, thereby increasing the number of acquired images, and hence the measurement time. Figure 1 (a) illustrates these two challenges - specularities and shadows - using a simple scene with two objects. To be able to recover dense and reliable information, a structured light vision system must be able to address the problems of bright highlights and dark shadows.

2. Diffuse Structured Light

We describe a simple but effective way to mitigate both specularities and shadows in structured light methods. Our approach uses a diffuser placed between the light source and the scene. Diffusers with isotropic scattering functions have been extensively used in machine vision to create extended gradient light sources for photometric stereo [5, 12, 10, 8, 7]. While isotropic diffusers reduce specularities, they also blur the incident illumination. Consequently, they are not directly applicable to structured light methods that rely on the projection of coded illumination patterns.

We make two observations that allow us to exploit the advantages of diffusers while maintaining the coding information present in the projected patterns. First, several structured light methods use patterns




Figure 1. **Conventional versus diffuse structured light:** The scene consists of a specular sphere and a diffuse cube. (a) In conventional structured light, point P lies on a specular highlight and is saturated in the image. Point Q is self-shadowed and point R is in a cast shadow. Depths cannot be computed for these points. The same problem arises with the right face of the cube, which is fully visible to the camera but is not illuminated. (b) In diffuse structured light, a single light ray incident upon a linear diffuser results in a fan of rays. Consequently, a one-dimensional set of incident rays is converted to a two-dimensional light field. Point P is illuminated by a fan of rays, emulating an extended (line) source. This reduces the problem due to specularities. (c) Linearly diffuse illumination can illuminate points in the scene (such as, Q and R) that would have been shadowed in the case of conventional structured light. Information can also be recovered at points (such as, S) that receive light from multiple disjoint segments of the light strip on the diffuser.

that have translational symmetry (see [13] for a survey). Coding is done along only one of the two spatial dimensions of the projected patterns (see Figure 2 for examples). For such patterns, the projected light field is a pencil of *illumination planes*, where each plane consists of a one-dimensional set of rays with the same intensity. Second, it is possible to construct diffusers that are *linear*, in that, they diffuse incident light only along one of the two spatial dimensions. For instance, diffusers based on lenticular arrays can have wide scattering along the diffusion axis, and no scattering in any other direction.

Based on these two observations, we present a technique called *diffuse structured light*. The idea is to place a linear diffuser between the light source and the scene such that every illumination plane intersects the diffuser along the diffusion axis. This setup achieves two conditions. First, each point in the scene is illuminated by an extended source rather than a point source. This results in reduction of both specularities and shadows, as illustrated in Figures 1 (b,c). Second, it ensures that the light rays scatter only *within* each illumination plane. As a result, each point receives light from a single iso-brightness line in the projected illumination pattern (Figure 2). Since there is no cross-talk between the illumination planes corresponding to different iso-brightness lines of the pattern, the coding information in the pattern is preserved.

As shown in Figures 1 (b,c), a one-dimensional set of rays is converted by the diffuser into a two-dimensional light field. We refer to this light field as *linearly diffuse illumination*. In the next section, we will derive the radiance of a scene point due to such illumination and show how it reduces specularities and shadows.




Figure 2. **Examples of commonly used illumination patterns with translational symmetry:** (a) Uniform illumination for defect detection. (b) Single light stripe, (c) binary pattern, (d) De Bruijn pattern and (e) N-ary coded pattern for depth estimation [13]. (f) Ramp brightness function, (g) triangular function and (h) sinusoid used for depth estimation [3]. (i) High-frequency stripe pattern used for separation of direct and indirect illumination [11].

3. Scene Radiance due to Linearly Diffuse Illumination

Consider the case of conventional structured light (SL) shown in Figure 3(a), where the point P is illuminated by a thin sheet of collimated light rays that are aligned with the vertical axis. Since we are considering

illumination patterns with translational symmetry, all the light rays in the sheet may be assumed to have the same radiance E_o . If the normal at the point P is \hat{n} , its irradiance is:

$$E_{conv} = E_o (\hat{n} \cdot \hat{z}) V, \quad (1)$$

where V is the binary visibility term for point P. If P is occluded from the light source, then $V = 0$, and if P directly sees the light source, then $V = 1$. The subscript *conv* denotes conventional SL. Let the bi-directional reflectance distribution function (BRDF) of P be $f(\hat{n}, \hat{v}, \hat{s})$, where \hat{v} is the viewing direction and \hat{s} is the illumination direction. In our setting, $\hat{s} = \hat{z}$. The radiance at point P is:

$$L_{conv} = E_o K_{conv}, \quad \text{where,} \quad (2)$$

$$K_{conv} = f(\hat{n}, \hat{v}, \hat{z}) (\hat{n} \cdot \hat{z}) V. \quad (3)$$

Now consider the case shown in Figure 3(b), where a linear diffuser with scattering function $D(\theta)$ is placed between the source and the scene. The diffuser converts a one-dimensional set of rays into a two-dimensional light field. The light field can be viewed as a collection of an infinite number of collimated sheets with different angles of incidence. For P, the collimated sheets range in incidence angle from α_1 to α_2 (see Figure 3(b)). The irradiance of P due to the sheets with incidence angles between α and $\alpha + d\alpha$ can be written as:

$$dE_{diff} = E_o D(\alpha) (\hat{n} \cdot \hat{s}(\alpha)) V(\alpha) d\alpha, \quad (4)$$

where $V(\alpha)$ is the binary visibility term for point P corresponding to the incidence angle α . The subscript *diff* denotes diffuse SL. Then, the radiance of P measured by the camera due to the above infinitesimal range of collimated sheets is:

$$dL_{diff} = E_o f(\hat{n}, \hat{v}, \hat{s}(\alpha)) D(\alpha) (\hat{n} \cdot \hat{s}(\alpha)) V(\alpha) d\alpha. \quad (5)$$

The total radiance of P due to the entire illuminating strip on the diffuser is found by integrating over the complete range of incidence angles:

$$L_{diff} = E_o K_{diff}, \quad \text{where,} \quad (6)$$

$$K_{diff} = \int_{\alpha_1}^{\alpha_2} f(\hat{n}, \hat{v}, \hat{s}(\alpha)) D(\alpha) (\hat{n} \cdot \hat{s}(\alpha)) V(\alpha) d\alpha. \quad (7)$$

It is worth noting that the term K_{diff} is *constant* with respect to the illumination intensity. When a method uses multiple illumination patterns, we get a set of radiance measurements at each scene point. The brightness of P due to the i^{th} pattern is:

$$L_{diff}^{(i)} = E_o^{(i)} K_{diff}. \quad (8)$$




Figure 3. Irradiance due to a collimated sheet of rays in (a) conventional structured light and (b) diffuse structured light.

The constant K_{diff} depends on the BRDF, location and normal of the point P, as well as the scattering function of the diffuser. As we will show in the next section, K_{diff} serves to reduce the brightness of specular points while maintaining the brightness of diffuse points. Moreover, if a point is in shadow in conventional SL ($K_{conv} = 0$), but sees a finite illuminated strip of the diffuser, then $K_{diff} > 0$. As a result, information may be recovered at points such as Q and R in Figure 1 (c) that are in shadow in the case of conventional SL. We now discuss these benefits in detail.

4. Reduction of Specularities and Shadows

Specularities: Isotropic diffusers have been used in photometric stereo techniques to reduce specularities and increase the area of the measured region (coverage) [5, 12, 10, 8, 7]. However, these diffusers blur the incident illumination equally in all directions and are not suitable for coded structured light methods. In contrast, the linear diffusers that we use not only reduce specularities and increase surface coverage, but also fully preserve the coding information for illumination patterns with translational symmetry.

Consider a diffuse scene point P. Its radiance is influenced by light received from the entire hemisphere aligned with its normal, as shown in Figure 4(a). In conventional SL, P would receive a single collimated beam of light. In diffuse SL, it receives the aggregate of weaker illumination from several directions instead of strong illumination from a single direction. If the width of the scattering function of the diffuser is chosen appropriately, all diffuse points will produce approximately the same brightness under collimated illumination and diffuse illumination, i.e., $K_{diff} \approx K_{conv}$.

Now, let us assume that the point P is mirror-like, as shown in Figure 4 (b). In conventional SL, if P is ap-




Figure 4. The illumination acceptance cone and the section of the strip on the diffuser that contributes to scene radiance in the case of (a) a diffuse point, (b) a highly specular (mirror) point, and (c) a rough specular (glossy) point. (d) The range of scene brightness values in the conventional case can be larger than the sensor’s dynamic range. This results in saturation and dark pixels. A diffuser compresses the scene’s brightness range to lie within the sensor’s dynamic range.

appropriately oriented (its illumination acceptance cone is pointed in the direction of the illumination), it would have very high radiance in the direction of the camera. In the case of diffuse SL, however, P receives weaker illumination from an infinitesimal patch on the diffuser. As a result, its irradiance (and hence radiance) would be much lower than in the case of conventional SL ($K_{diff} \ll K_{conv}$). As surface roughness increases (see Figure 4 (c)), the reflectance in the camera direction decreases, but the area of the diffuser contributing to its radiance in the camera direction increases. In short, the variation in the brightness of a specular point with roughness is much lesser in the case of diffuse SL compared to conventional SL.

Intuitively, the diffuser reduces the range of scene brightness values, as illustrated in Figure 4(d). The range of brightness values in the conventional case can be significantly larger than the sensor’s dynamic range, resulting in saturation and dark pixels. By using the diffuser, the range can be made to fall within the sensor’s dynamic range.

Shadows: Diffuse illumination also alleviates the problem of shadows by “looking into” surfaces that are tilted away, or occluded from, the original direction of light projection, as illustrated in Figure 1 (c). While the points Q and R are certain to be shadowed in the conventional system ($E_{conv} = L_{conv} = 0$), they receive significant irradiance in the diffuse case. Both Q and R are able to “see” segments of the light strip on the diffuser, resulting in $K_{diff} > 0$, and hence, $L_{diff} > 0$. In fact, the section of the light strip visible to a scene point does not have to be contiguous. It can even be a set of disjoint segments, as is the case of point S¹.

Consider a scene point that lies on an illumination plane but does not receive illumination in a conventional system as it is obstructed by other points. With a diffuser, this point will receive irradiance from the

¹Shadow reduction due to diffuse illumination often occurs in outdoor scenes as well. Consider the appearance of buildings on a cloudy day. An overcast sky is a nearly uniform hemispherical source (also referred to as a Ganzfeld source [9]). Photographs taken on such days do not have strong shadows.

same illumination plane. *This is an important property of linearly diffuse illumination. In the case of 3D scanning, it allows us to estimate depth by using exactly the same triangulation algorithm as in a conventional system (without a diffuser).*

5. Quantitative Analysis using Simulations

We now quantify the reduction in specularities and shadows as a function of the diffuser’s scattering angle (width of the scattering function).

Specularities: We computed the radiance distribution on a specular hemisphere for different scattering angles, using the Torrance-Sparrow reflectance model [14] (see Appendix A for the derivation of radiance). The hemisphere is illuminated and viewed orthographically from the same direction (top).

For illustration, points $p(\theta, \phi)$ on the hemisphere (spherical coordinates) are mapped to points $q(r, \omega)$ on a unit disk (polar coordinates), where $r = \frac{2\theta}{\pi}$ and $\omega = \phi$. Figure 5 shows the radiance distributions for different roughness parameters: $\rho = 0.05$ (very smooth material), $\rho = 0.45$ (moderately rough) and $\rho = 0.9$ (very rough). The portion of the hemisphere shown in grey has radiance greater than 10 grey levels (for an 8-bit camera), and is considered covered. The uncovered portion, shown in red, can not be measured reliably. Figure 6 shows the plots for coverage versus the diffuser scattering angle. Increasing the scattering angle increases the coverage. The diffuser used in our experiments has a scattering angle of approximately 100° .

Shadows: For shadows (self, not cast), we computed the radiance distribution on a Lambertian sphere (the derivation of the radiance is given in Appendix A). The sphere is mapped to a unit disk ($r = \frac{\theta}{\pi}$). Figure 7 shows the radiance distributions for different diffuser scattering angles. The intensity threshold is again 10 grey levels. In the absence of diffuser ($\delta = 0^\circ$), only the top hemisphere receives illumination. With the diffuser, it is possible to cover a large fraction of the




Figure 5. **Coverage of metallic hemisphere (specularities):** The radiance distribution on a specular hemisphere is simulated for different diffuser scattering angles and material roughness parameters. The roughness parameter (Torrance-Sparrow reflectance model) is $\rho = 0.05$ for the top row (very smooth material), $\rho = 0.45$ for the middle row (moderately rough) and $\rho = 0.9$ for the bottom row (very rough). The hemisphere is mapped to a unit disk for illustration. The portion of the hemisphere with radiance greater than 10 grey levels (for an 8-bit camera) is considered covered (shown in grey; brighter shade corresponds to higher radiance). The uncovered portion, shown in red, cannot be measured reliably.




Figure 6. **Coverage of metallic hemisphere:** The fraction of the covered area increases with the diffuser scattering angle. The diffuser used in our experiments has a scattering angle of approximately 100° .

bottom hemisphere as well. Figure 8 shows a plot of coverage fraction versus the scattering angle.

6. Experiments

Our implementation of diffuse SL uses a projector as the source and a linear diffuser between the projector

and the scene², as shown in Figure 9. The projector used in our experiments is an AAXA M2 Projector with

²Linearly diffuse illumination can also be generated by placing a lenticular array, or a micro-louvre filter with deep and narrow walls, next to a widely scattering computer display (CRT, LED, LCD, etc.). The advantage of such an apparatus is that it is more compact than using a projector and diffuser.




Figure 7. **Coverage of Lambertian sphere (shadows):** The radiance distribution of a Lambertian sphere for different diffuser scattering angles. The sphere is mapped to a unit disk for illustration. Without the diffuser, only the top hemisphere is illuminated (shown in grey). With the diffuser, it is possible to cover a fraction of the bottom hemisphere as well.

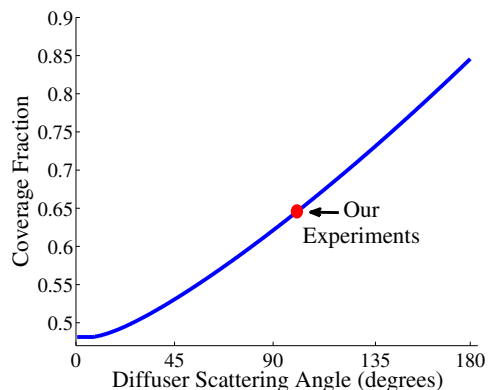



Figure 8. **Coverage of Lambertian sphere:** Coverage increases with the diffuser’s scattering angle.

XGA (1024 × 768) resolution. The camera is a Canon EOS XSi SLR with a resolution of 4290 × 2856 pixels. The diffuser is a lenticular array (Edmund Optics part number NT43-029) that is 12” × 12” in size. The scattering is approximately ±50 degrees along the diffusion axis, and is negligible perpendicular to the axis. To handle larger objects/scenes, it is possible to construct larger diffusers by either tiling multiple smaller diffusers or custom fabricating a single large diffuser.

6.1. Results: Structured Light 3D Scanning

We have used diffuse SL for 3D scanning using sinusoidal phase shifting. We show results for objects made of various polished metals with strong specularities, as well as scenes with shadows (attached and cast). The projected patterns are vertical sinusoids of different frequencies and phase shifts. For each scene, 6 – 9 projected patterns (images) were used.

Specularities: Our first example, shown in Figure 10 (a), is a glossy metal hemisphere on a diffuse plane. For both conventional and diffuse SL, we acquired images at 12 different camera exposures ranging from 0.1s to 4.0s. For both the cases, the best exposure value was



Experimental Setup for diffuse SL

Figure 9. The setup for diffuse SL consists of a linear diffuser placed between the object and the projector. The diffusion axis is horizontal in the image. Diffuse SL uses exactly the same algorithms as used by conventional SL.

chosen so that the percentage of recovered area is maximized. In the best case of diffuse SL (Figure 10 (b)), the specularities are suppressed considerably, resulting in a highly accurate reconstruction. In the best case of conventional SL (Figures 10 (c)), shape within specular highlights is not reconstructed due to saturation. When lower exposures are used to avoid saturation (Figures 10 (d,e)), some of the non-highlight regions of the object are too dark to be reconstructed reliably.

Figures 13, 14 and 15 show results for other glossy metallic objects. For conventional SL, reconstructions at different camera exposures are shown, with the best case corresponding to the highest shown exposure (see Figures 13 (f), 14 (c) and 15 (c)). The best case reconstructions achieved with diffuse SL consistently outperform conventional SL. The table in Figure 12 shows a quantitative comparison between the best cases of diffuse SL and conventional SL for the hemisphere scene in Figure 10. The mean square error (MSE) is higher

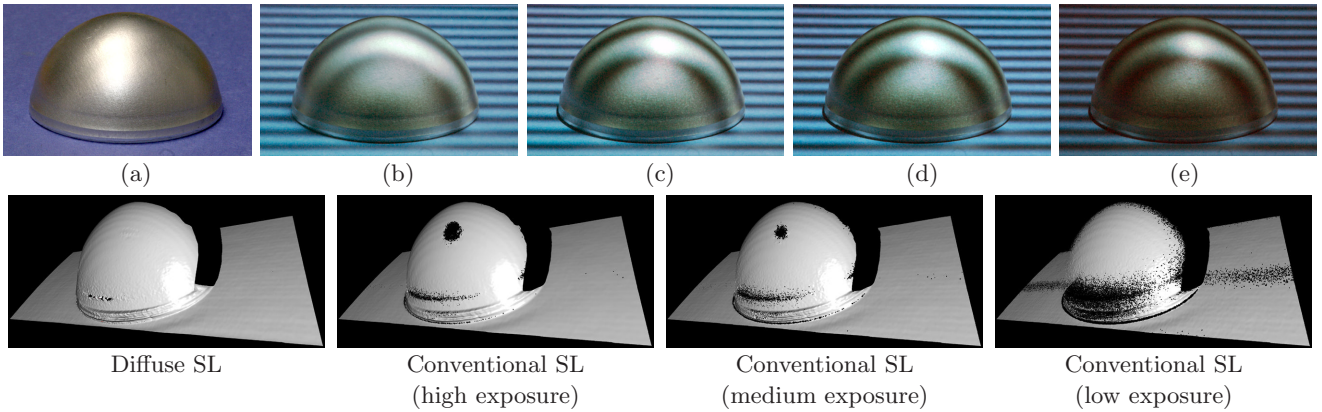


Figure 10. **Structured light 3D scanning of specular objects:** (a) A specular metal hemisphere on a diffuse plane. (b) One (out of six) input image used for diffuse SL. (c-e) Corresponding input images for conventional SL under high (1.6s), medium (1.0s) and low (0.3s) exposure settings, respectively. Note that each of these images includes strong (saturated) specularities, dark (noisy) regions, or both. (f) Reconstructed shape for diffuse SL. (g-i) Reconstructed shapes for conventional SL with high, medium and low exposures, respectively. At high exposures, regions with specular highlights are not reconstructed due to saturation. At low exposures, some scene regions are too dark to be reconstructed reliably.

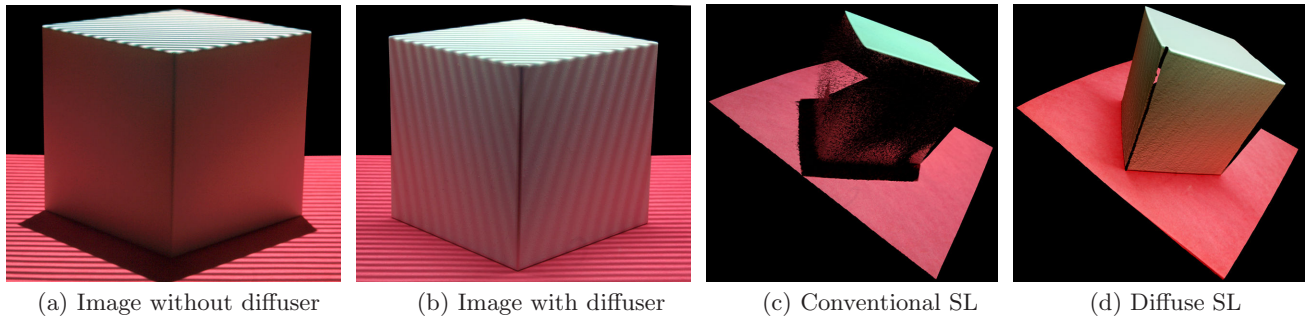


Figure 11. **3D scanning of shadowed regions:** A cube on a flat plane. The projector is above the cube - the direction of projection is normal to the top face of the cube. (a) For conventional SL, the two vertical faces of the cube are within the attached shadow. Additionally, the cube casts a shadow on the plane it sits on. (b) With a diffuser, the entire scene receives illumination. (c,d) Reconstructions using conventional and diffuse SL, respectively.

for conventional SL due to specularities.

Shadows: Figure 11 shows a cube on a flat plane. In conventional SL, the two vertical faces of the cube are in shadow. The cube also casts a shadow on the plane. Depth cannot be computed for these shadow regions. However, with a diffuser, these regions of the scene also receive illumination. Notice the projected stripes on the vertical walls and near the base of the cube in Figure 11 (b). As a result, the coverage of the scene with diffuse SL is significantly higher compared to conventional SL, as shown in the table in Figure 12. In Figure 16, the goal is to recover the 3D structure of the fingerprint. Without the diffuser, a large portion of the finger tip is in shadow, resulting in noisy reconstruction. With the diffuser, most of the finger receives illumination, resulting in an accurate reconstruction.

6.2. Results: Direct-Global Separation

The second application that we show is the separation of direct and global components of light transport

	Diffuse SL (Best Exposure)		Conventional SL (Best Exposure)	
	MSE	Coverage	MSE	Coverage
Hemisph.	0.82mm	92.8%	5.02mm	88.7%
Cube	0.76mm	98.2%	0.77mm	29.3%

Figure 12. The reconstruction mean square error (MSE) and the percentage of reconstructed area (coverage) for conventional and diffuse SL. The ground truth dimensions of the hemisphere and the cube were obtained by physical measurement.

using high-frequency illumination [11]. In our experiments, we used 18 shifted high frequency *binary stripe* patterns. Since the linear diffuser does not blur these patterns perpendicular to the stripes, the light field incident on the scene after diffusion is also binary and retains the high frequencies. As a result, the direct and global components can be separated faithfully. Figure 17 shows an example of direct-global separation in the presence of subsurface scattering. Note that the direct and global components are somewhat different for

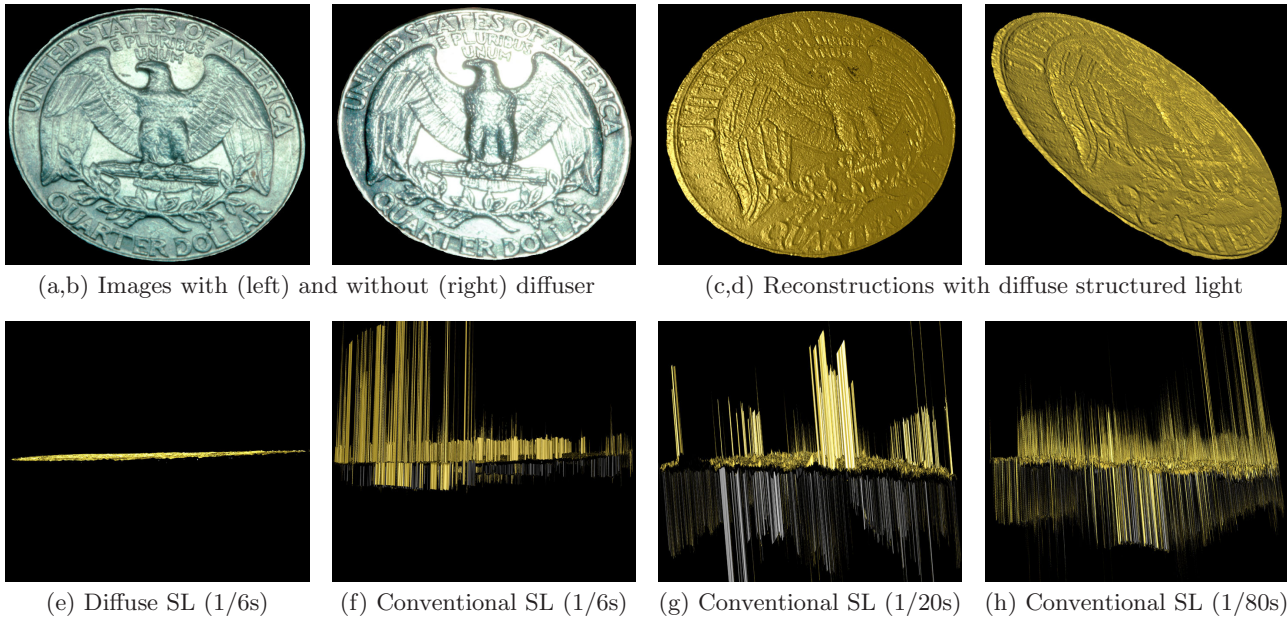


Figure 13. **Reconstruction of a quarter coin:** (a,b) Images with and without diffuser, respectively. (c-e) Reconstructions with diffuse SL. (f-h) Reconstructions obtained with conventional SL for different camera exposures. Since the coin is specular, conventional SL results in depth errors due to specular highlights at higher exposures and low signal-to-noise ratio at lower exposures.




Figure 14. **Reconstruction of an earring.** (a) The earring is made of polished (specular) copper. (b) Reconstruction with diffuse SL. (c,d) Reconstructions with conventional SL at high and low exposure, respectively.

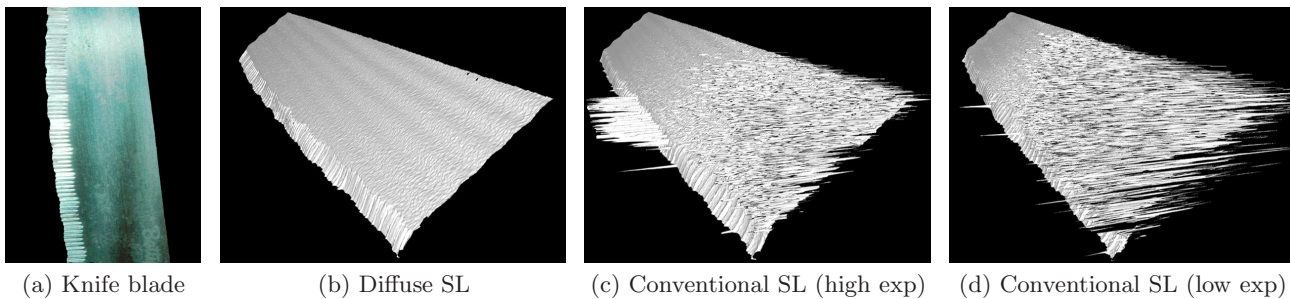


Figure 15. **Reconstruction of a knife blade with serrated edge:** (a) The knife blade is made of polished stainless steel. (b) Reconstruction with diffuse SL. Conventional SL results in erroneous reconstructions due to (c) specular highlights with high exposure, and (d) low signal-to-noise ratio with low exposure.

the cases with and without the diffuser, as they correspond to different light sources - a line source (with diffuser) and a point source (without diffuser). The direct component without the diffuser has strong high-

lights on the grapes. With the diffuser, the highlights are spread out as expected. Figure 18 shows an example of direct-global separation for a scene with strong interreflections. The direct component without the dif-




Figure 16. **Reconstruction of a fingerprint:** (a) Without the diffuser, a large portion of the finger tip is in shadow, resulting in (c) noisy reconstruction. (b) With the diffuser, most of the scene receives illumination, resulting in (d) an accurate reconstruction.




Figure 17. **Direct-global separation in the presence of subsurface scattering:** Diffuse structured light can also be used for direct-global separation by projecting high-frequency *binary stripe* patterns. Since linear diffusion does not blur these patterns perpendicular to the stripes, the light field after diffusion is also binary and high-frequency. As a result, the direct and global components can be separated faithfully.




Figure 18. **Direct-global separation in the presence of interreflections:** Notice the strong shadows in (b) the direct component without the diffuser. (d) In comparison, the shadows are softer in the direct component with the diffuser.

fuser has sharp shadows. As expected, the shadows are softer and reduced in the direct component measured with the diffuser.

7. Discussion

Effectiveness and Scope: Diffuse SL mitigates the problems due to specularities and shadows, but may not completely remove them. As shown in the simulations (Section 5) and the experimental results (Section 6), the amount of reduction in errors is a function of the scene geometry, material properties and the diffuser scattering angle. For example, for the cube shown in Figure 11, the amount of coverage is increased significantly. If the cube is rotated so that the wall of the cube is perpendicular to the diffusion axis, the wall will remain in shadow even with the diffuser. In general, if the normal of a self-shadowed scene point is

perpendicular to the diffusion direction, it will remain in shadow with the diffuser. If, however, the normals in the scene are evenly distributed, the fraction of points that remain in shadow with the diffuser will be small.

Diffuser Placement: An important design choice for diffuse structured light is the distance between the diffuser and the target scene. Smaller the distance, wider the range of incident angles for the diffuse illumination, and hence greater the reduction in specularities and shadows. On the other hand, if the diffuser is too close to the scene, it could obstruct the camera’s field of view. Ultimately, the appropriate distance would depend on the application at hand. It is also possible to work around the obstruction issue by having the camera view the scene through a small hole in the diffuser. Note that the effect of such a hole is less severe in the case of diffuse structured light (Figure 1 (c)).

Polarization and Diffusion: It is a common practice in machine vision to reduce specular reflections by using polarized illumination and sensing [6, 16]. This is achieved by placing two polarizers with orthogonal polarization directions, one in front of the source and the other in front of the camera. The same approach can be used in the case of diffuser structured light - a polarization filter can be incorporated within the illumination optics to further mitigate specularities.

Defocus and Diffusion: Thus far, we have assumed that the rays of light incident upon the diffuser are collimated. However, in practice, any illumination source has a limited depth of field. If the projected illumination pattern is $p(x, y)$, the illumination received by the scene is $p(x, y) * h(x, y)$, where $h(x, y)$ is the defocus kernel, which is depth dependent. When a diffuser is incorporated, the received illumination is blurred by an additional diffusion blur kernel $d(x, y)$ [18]:

$$p'(x, y) = p(x, y) * h(x, y) * d(x, y). \quad (9)$$

Note that, in our case, the diffusion blur kernel is only one-dimensional (along the diffusion axis). As a result, while diffusion changes the blur size along the diffusion axis, the blur kernel remains the same in the orthogonal direction. For this reason, linearly diffuse illumination can also be incorporated into structured light methods that measure the amount of illumination defocus to recover scene depth [17, 4].

Illumination Patterns Without Translational Symmetry: Linear diffusers cannot be used with all two-dimensional patterns that can be expressed with a one-dimensional function. For example, while a pattern with concentric circles can be described with a one-dimensional brightness function, a tangential diffuser will not be able to preserve the coding information - rays corresponding to different circles will overlap. However, diffusers can be implemented for any pattern whose iso-intensity contours are straight lines. For example, radial diffusers [1] can be used with illumination patterns that are radially invariant.

Effect of Partially Incoherent Light and Diffraction: The physical features of the diffuser are typically large compared to the wavelength of light. Hence, the effects of diffraction are small, and moreover, they only impact the exact shape of the 1D scattering function, which is not critical to our approach. More importantly, diffraction and partial coherence do not introduce scattering in the orthogonal direction. As a result, the angular spread in the orthogonal direction remains zero for all practical purposes.

A. Derivation of Radiance

Points on Specular Hemisphere: Without loss of generality, let the hemisphere be centered at the origin with the zenith at $(0, 0, 1)$. The hemisphere is viewed orthographically from the top, so that the viewing vector $\hat{v} = (0, 0, 1)$.

The diffuser is placed so that the diffusion axis is along x-axis. The diffuser scattering function is constant within the diffuser scattering angle (δ). The size of the diffuser is assumed to be sufficiently large so that all the scene points receive illumination from all possible scattering angles.

Since the four quadrants of the hemisphere are symmetric, we derive the radiance distribution for one quadrant. Consider the point $p(\theta, \phi)$ with spherical coordinates (θ, ϕ) , $0 \leq \theta \leq \frac{\pi}{2}$ and $0 \leq \phi \leq \frac{\pi}{2}$. Point p receives illumination from a single diffuser strip parallel to the x-axis. According to the Torrance-Sparrow reflectance model [14], the radiance due to a single infinitesimal diffuser element corresponding to the incidence angle α is given as:

$$dL(\theta, \phi) = C \left(\frac{F \cdot G \cdot e^{-\left(\frac{\tan \beta}{\rho}\right)^2}}{\cos \theta_l \cos \theta_v} \right) d\alpha, \quad (10)$$

where C is the proportionality constant which includes albedo and light source intensity, θ_l is the angle between the surface normal and the illumination direction, θ_v is the angle between the normal and the viewing direction, and ρ is the material roughness parameter. $F(\theta, \phi, \alpha)$ is the Fresnel term, for which we used the Schlick approximation: $F = F_o + (1 - F_o)(1 - \cos \theta_l)^5$. In our simulations, we used $F_o = 0.9$, which corresponds to the metal silver. $G(\theta, \phi, \alpha)$ is the geometric attenuation factor and β is the angle between the normal and the half-vector.

Due to self-shadow, p does not get illumination from diffuser elements corresponding to $\alpha < \alpha_o$, where:

$$\alpha_o = \min \left(\max \left(\text{atan} \left(\frac{-\cos \theta \cos \phi}{\sin \theta} \right), \frac{-\delta}{2} \right), \frac{\delta}{2} \right). \quad (11)$$

The total radiance is found by integrating over all lighting directions α , $\alpha_o \leq \alpha \leq \frac{\delta}{2}$:

$$L(\theta, \phi) = C \int_{\alpha_o}^{\frac{\delta}{2}} \left(\frac{F \cdot G \cdot e^{-\left(\frac{\tan \beta}{\rho}\right)^2}}{\cos \theta_l \cos \theta_v} \right) d\alpha. \quad (12)$$

Points on Lambertian Sphere: Next, we derive the radiance distribution over a Lambertian sphere. Consider the point $p(\theta, \phi)$, $0 \leq \theta \leq \pi$ and $0 \leq \phi \leq \frac{\pi}{2}$. The normal at p is $\hat{n}_p = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)$. The radiance at p due to a single infinitesimal diffuser element e_α is:

$$dL(\theta, \phi) = C (\hat{n}_p \cdot \hat{s}(\alpha)) d\alpha, \quad (13)$$

where C is the proportionality constant which includes albedo and light source intensity. The total radiance is found by integrating over all lighting directions α , $\alpha_o \leq \alpha \leq \frac{\delta}{2}$ (α_o is the same as before):

$$L(\theta, \phi) = C \int_{\alpha_o}^{\frac{\delta}{2}} (\hat{n}_p \cdot \hat{s}(\alpha)) d\alpha, \quad (14)$$

After simplification, we get the following closed-form expression:

$$L(\theta, \phi) = C \left(\cos \theta \left(\sin \frac{\delta}{2} - \sin \alpha_o \right) - \right. \quad (15)$$

$$\left. \sin \theta \cos \phi \left(\cos \frac{\delta}{2} - \cos \alpha_o \right) \right). \quad (16)$$

Acknowledgments: This research was supported in parts by NSF (grant number IIS 09-64429) and ONR (grant number N00014-11-1-0285).

References

- [1] O. Cossairt, C. Zhou, and S. K. Nayar. Diffusion coding photography for extended depth of field. *Proc. of ACM Trans. on Graphics*, 29(3), 2010. [10](#)
- [2] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. *ACM SIGGRAPH*, 1997. [1](#)
- [3] Y. Gong and S. Zhang. Ultrafast 3-D shape measurement with an off-the-shelf DLP projector. *Opt. Express*, 18(19), 2010. [2](#)
- [4] M. Gupta, Y. Tian, S. G. Narasimhan, and L. Zhang. (De) Focusing on global light transport for active scene recovery. *Proc. IEEE CVPR*, 2009. [10](#)
- [5] K. Ikeuchi. Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE PAMI*, 3(6), 1981. [1, 3](#)
- [6] K. Koshikawa and Y. Shirai. A model-based recognition of glossy objects using their polarimetric properties. *Advanced Robotics*, 2(2), 1987. [10](#)
- [7] W.-C. Ma, T. Hawkins, P. Peers, C.-F. Chabert, M. Weiss, and P. Debevec. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proc. EGSR*, 2007. [1, 3](#)
- [8] G. McGunnigle. Photometric stereo with graduated extended sources for recovery of specular surfaces. *J. Opt. Soc. Am. A*, 27(5), 2010. [1, 3](#)
- [9] W. Metzger. Optische untersuchungen am ganzfeld. *Psychologische Forschung*, 13, 1930. [4](#)
- [10] S. K. Nayar, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of lambertian, specular, and hybrid surfaces using extended sources. In *Proc. of IEEE International Workshop on Industrial Applications of Machine Intelligence and Vision*, 1989. [1, 3](#)
- [11] S. K. Nayar, G. Krishnan, M. D. Grossberg, and R. Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *Proc. of ACM Trans. on Graphics*, 25(3), 2006. [2, 7](#)
- [12] S. K. Nayar and A. C. Sanderson. Determining surface orientations of specular surfaces by intensity encoded illumination. In *Proc. SPIE*, volume 850, 1987. [1, 3](#)
- [13] J. Salvi, S. Fernandez, T. Pribanic, and X. Llado. A state of the art in structured light patterns for surface profilometry. *Pattern Recognition*, 43(8), 2010. [2](#)
- [14] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *J. Opt. Soc. Am. A*, 57(9), 1967. [4, 10](#)
- [15] L. B. Wolff. On the relative brightness of specular and diffuse reflection. In *Proc. IEEE CVPR*, 1994. [1](#)
- [16] L. B. Wolff and T. Boult. Constraining object features using a polarization reflectance model. *IEEE PAMI*, 13(7), 1991. [10](#)
- [17] L. Zhang and S. K. Nayar. Projection defocus analysis for scene capture and image display. *Proc. of ACM Trans. on Graphics*, 25(3), 2006. [10](#)
- [18] C. Zhou, O. Cossairt, and S. K. Nayar. Depth from diffusion. In *Proc. IEEE CVPR*, 2010. [10](#)