

Ambient Occlusion via Compressive Visibility Estimation

Wei Yang^{*1} Yu Ji^{*1} Haiting Lin¹ Yang Yang¹ Sing Bing Kang² Jingyi Yu¹
¹University of Delaware ²Microsoft Research
 {wyangcs, yuji, haiting, yyangwin, jingyiyu}@udel.edu sbkang@microsoft.com

Abstract

There has been emerging interest on recovering traditionally challenging intrinsic scene properties. In this paper, we present a novel computational imaging solution for recovering the ambient occlusion (AO) map of an object. AO measures how much light from all different directions can reach a surface point without being blocked by self-occlusions. Previous approaches either require obtaining highly accurate surface geometry or acquiring a large number of images. We adopt a compressive sensing framework that captures the object under strategically coded lighting directions. We show that this incident illumination field exhibits some unique properties suitable for AO recovery: every ray's contribution to the visibility function is binary while their distribution for AO measurement is sparse. This enables a sparsity-prior based solution for iteratively recovering the surface normal, the surface albedo, and the visibility function from a small number of images. To physically implement the scheme, we construct an encodable directional light source using the light field probe. Experiments on synthetic and real scenes show that our approach is both reliable and accurate with significantly reduced size of input.

1. Introduction

The problem of recovering intrinsic properties of a scene/object from images has attracted much attention in the past decade. Tremendous efforts have been focused on intrinsic properties related to shading and reflectance [11, 2]. The problem is inherently challenging as it is under-constrained: we have multiple unknowns but a smaller set of constraints (images). Additional constraints such as priors are needed to make the problem trackable. For example, the earlier work by Weiss [35] used a maximum likelihood (ML) method to recover illumination invariant intrinsic properties from a sequence of images captured at a fixed viewpoint but under significantly different lightings.

More recent work has focused on recovering intrinsic properties related to surface reflectance and visibility [3].

In this paper, we explore a challenging type of intrinsic properties called the ambient occlusion map. Ambient Occlusion (AO) characterizes the visibility of a surface point due to local geometry occlusions. Given a scene point x , its AO measures the occlusion of ambient light caused by local surface geometry:

$$A(x) = \frac{1}{\pi} \int_{\Omega} v(x, \hat{w}) \langle \hat{w} \cdot \hat{n} \rangle d\hat{w} \quad (1)$$

where \hat{w} is the direction of incident light; \hat{n} is the normal of x ; and $\langle \cdot \rangle$ refers to the dot product. $v(x, \hat{w})$ is the local visibility function and is equal to 0 if the light ray from \hat{w} is occluded from x . Ω is the unit sphere over the point x .

AO reveals local visibility of illumination and hence affects the appearance of objects under shading. Applications are numerous in both computer vision and graphics, ranging from inverse rendering [1], image based relighting [25] to photometric stereo [20, 32, 6]. Recovering AO from images, however is highly ill-posed (See details in Sec. 3.1). A brute-force approach is to first capture the 3D geometry of object and then compute AO using ray casting [28]. This task, however, is challenging since accurate 3D reconstruction remains difficult. The seminal work by Hauage *et al.* [12] analyzes a stack of images acquired under a moving point light source. To make the problem trackable, they adopt a parametric visibility model, a cone-shaped function centered at the normal. A simple, per-pixel statistics was then used to cancel out surface albedo and to estimate visibility. Although highly effectively, their approach used a large number of images to obtain high accuracy and the parametric visibility model is not always valid (e.g., point A and D in Fig. 1).

In this paper, we present a novel computational imaging solution for recovering AO. We adopt a compressive sensing framework that captures the object under strategically coded directional lights. We show that the incident illumination field exhibits some unique properties suitable for AO recovery: every ray's contribution to the visibility function is binary and their distribution for AO measurement is

*These authors contributed to the work equally.

sparse. This enables a sparsity-prior based solution for iteratively recovering the surface normal, the surface albedo, and the sparse visibility function from a small number of images. To physically implement the scheme, we construct an encodable directional light source using the light field probe [36]. We validate our approach on both synthetic and real data and show that our scheme produces AO estimation at comparable accuracy to [12] but with a much smaller set of images. In addition, we can recover more general visibility functions beyond the normal-centered cone-shaped models.

2. Related Work

In this section, we review state-of-the-art intrinsic image recovery techniques and emerging computational imaging solutions.

Intrinsic Image Recovery. Traditional approaches have focused on intrinsic properties related to reflectance (albedo), normal and geometric occlusion. Many single images based methods [17, 33, 19] aim to perform the reflectance and illumination decomposition. Theoretically, these methods can potentially be used to recover AO when the scene is illuminated by a uniform and constant light source. However, they rely on smoothness prior and hence are not reliable in the presence of complex albedo and surface geometry. Recent studies show that image sequences based methods (e.g. [35]) cannot robustly handle soft and persistent shadows produced by AO. Our work is also related to visibility estimation in photometric stereo [20, 32, 6]. Sunkavalli *et al.* [32] approximate lighting visibility via dimensionality reduction in illumination spaces. Their method, however, is not suitable for handling a large number of varying lighting conditions. Aldrian and Smith [1] conduct inverse-render on simple geometry such as faces under uniform and constant illumination. Hauagge *et al.* [12] estimate AO from a stack of images captured under a moving point light source. They adopt a simplified visibility model and show impressive results on recovering AO. A major drawback of their technique is the requirement of a large number of input images, which we aim to reduce through computational imaging and compressive reconstruction.

Computational Imaging. The core of our technique is to create a controllable illumination field. Masselus *et al.* [23] position an object on a rotational table and illuminate it using a projector. By exploiting the rich shading and geometric information through spatial-angular analysis, they can relight the object with high realism. Later, Cossairt *et al.* [8] apply computational illumination to produce synthetic illumination between real and synthetic objects. Debevec [9] uses the Light Stage system to exploit the richness in

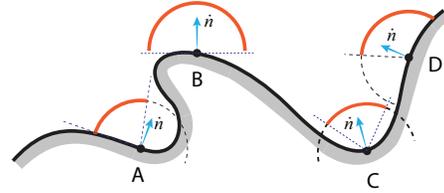


Figure 1: An 1-D illustration of ambient occlusions for different scene points with different occlusion levels.

angular and spatial variation of the light field, and the object can be relit with nearly realism. Ezra *et al.* [4] proposes a novel BRDF measurement device consisting exclusively of LEDs. Their device uses no cameras and is fast and simpler for measuring BRDF.

Our idea resembles the inverted light field camera. Light fields [18] are image based representations that gather rays sampled both spatially and angularly. Physically, they can be collected through a light field camera array or more recently a light field camera. The concept of light field illumination is relatively new. Wetzstein *et al.* [36] put coded image patterns behind a microlens array to encode rays spatially and angularly. The device, called the light field probe, can be used to determine light path variations and then to reconstruct transparent objects [37]. In a similar vein, Ji *et al.* [15] place a color coded pattern behind the light field probe to capture light paths through a 3D gas flow for volumetric heat reconstruction. In this paper, we build a controllable directional illumination field using the light field probe.

Compressive Sensing. Finally, our work is closely related to latest advances on Compressive Sensing (CS) [5] for signal reconstruction. The literature of CS is huge and we refer the readers to the comprehensive survey [27]. In photography, CS allows a reduction in image acquisition energy per image by as much as a factor of 15 at the cost of complex decompression algorithms [30]. There is also emerging interest on applying CS for multi-spectral imaging [21], light field imaging [22] and super-resolution [13], light transport [26], and depth sensing [7]. Our work employs the compressive sensing technique to reduce the number of input images for AO reconstruction.

3. Ambient Occlusion Recovery

In this section, we first show why AO recovery from images is an ill-posed problem and then discuss how to solve it using CS. Before proceeding, we clarify our assumptions. Similar to previous work [12], we assume Lambertian reflectance and fixed camera location. The AO map is measured w.r.t. the camera’s viewpoint.

3.1. Ill-Posedness

By Eqn. 1, we can discretize the lighting directions and approximate AO as:

$$\tilde{A} = \frac{1}{c} \sum_{i=1}^N v_i \langle \dot{w}_i \cdot \dot{n} \rangle \quad (2)$$

where c is the normalization factor that regularizes \tilde{A} into the range $[0, 1]$, \dot{w}_i is the discretized lighting direction and v_i is the local visibility function.

To measure \tilde{A} , we can illuminate the object using a dense set of uniform directional lights \dot{w}_i and measure its corresponding intensity as:

$$I_i = \rho v_i \langle \dot{w}_i \cdot \dot{n} \rangle \quad (3)$$

where ρ is the surface albedo.

Summing up images captured from all directions, we have

$$\sum_{i=1}^N I_i = \rho \sum_{i=1}^N v_i \langle \dot{w}_i \cdot \dot{n} \rangle = \rho \tilde{A} c \quad (4)$$

Notice that the AO term \tilde{A} in Eqn. 2 cannot be resolved based on the intensity of images since the albedo ρ is also unknown. To address this issue, Hauagge *et al.* [12] adopt a simplified parametric visibility model. Specifically, they assume that the visibility function follows cone-shaped distribution centered at the normal as $A = \pi \sin^2 \alpha$, where α is the cone's half angle. Computing the visibility function is to estimate α . Under uniformly distributed lighting directions, they have shown that computing $\kappa = E[I]^2 / E[I^2]$ ($E[\cdot]$ stands for expectation) directly cancels the albedo and can be used to directly solve for α . For their assumption to work, densely distributed light sources will be needed and an image needs to be captured for each lighting direction.

3.2. AO Estimation Via Compressive Sensing

Our approach is motivated by recent compressive sensing schemes. Instead of capturing one lighting direction at a time, we aim to enable multiple lighting directions in one shot. A downside though is that we can no longer use the κ statistics mentioned above for canceling out the albedo for visibility estimation. Instead, we build our solution on compressive signal reconstruction.

We use a binary vector $b = [l_1, \dots, l_N]$ to represent the status of N lighting directions, where $l_i = 1$ or 0 corresponds to if the lighting direction \dot{w}_i is enabled or disabled. Under this formation, we have:

$$I = \rho \sum_{i=1}^N l_i v_i \langle \dot{w}_i \cdot \dot{n} \rangle \quad (5)$$

We can now use a set of M strategically coded directional lighting patterns. For each pattern b^j , $j = 1 \dots M$, we

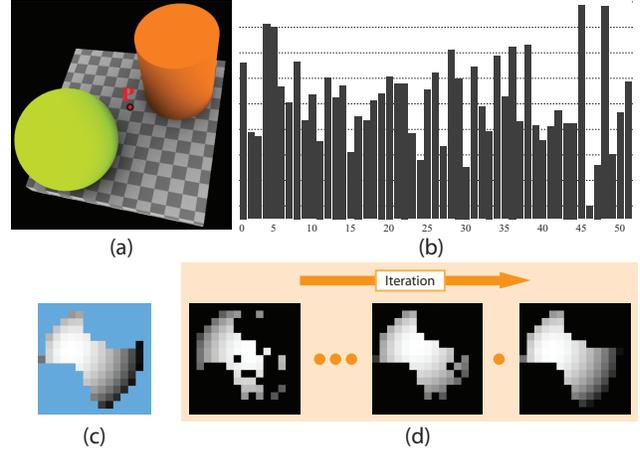


Figure 2: (a) A simple scene illuminated by one of our directionally encoded light sources. Point P's AO is affected by the sphere and the cylinder. We use 50 different patterns and (b) shows the P's intensity with respect to different patterns. (c) shows P's ground truth visibility distribution over the hemisphere (foreshortened by its normal). (d) shows intermediate visibility estimation results at different iterations by our algorithm.

capture an image I^j . This results in an $M \times N$ measurement matrix $B = [b^1, b^2, \dots, b^M]^T$ and we can rewrite Eqn. 5 as

$$I = \rho B [V * W(\dot{n})] \quad (6)$$

where $W(\dot{n}) = [\langle \dot{w}_1 \cdot \dot{n} \rangle, \langle \dot{w}_2 \cdot \dot{n} \rangle, \dots, \langle \dot{w}_N \cdot \dot{n} \rangle]$ and $[\ast]$ refers to the pairwise element-wise product.

Given the measurements, for each pixel, we aim to solve for ρ (a scalar), V (a $N \times 1$ vector) and \dot{n} (a unit vector with two degrees of freedoms). Hence the number of unknowns is $N + 3$. Again, we want to use fewer constraints (number of input images), i.e., $M < N + 3$. Our solution is to reduce the problem to two sub-problems and solve them using iterative optimization.

Visibility Recovery Sub-problem. The simplest initialization of the surface normal is to directly use the camera's viewing direction $\dot{n}_0 = [0, 0, 1]$ for all normals. This works well in our scheme but incurs slower convergence. A better scheme is to use the lighting pattern that incurs highest intensity and use their averaged direction. This results in the initial assignment of W_0 . We then set out to find the optimal ρ and V by optimizing the following objective function:

$$\begin{aligned} \rho, V \leftarrow \arg \min_{\rho, V} & \|\rho B (W_0 * V) - I\|_2 \\ \text{Subject to: } & V = [0, 1] \end{aligned} \quad (7)$$

Notice that V is a binary pattern and solving V in this optimization is NP-hard.

Recent works in compressive sensing [5] and signal representation [34, 16] show that this problem can be reduced to an ℓ_∞ regularized ℓ_1 minimization. Specifically, we transform the optimization objective function to:

$$\hat{\rho}, \hat{V} \leftarrow \arg \min_{\rho, V} \left\{ \|\rho B(W_0 * V) - I\|_2^2 + \lambda_1 \|V\|_1 + \lambda_1 \|V - 0.5\|_\infty + \lambda_2 \|\nabla V\|_1 \right\} \quad (8)$$

where λ_1, λ_2 and λ_3 are weighting factors. The new objective function consists of four terms: 1) $\|\rho B(W_0 * V) - I\|_2^2$ corresponds to the fidelity term where the estimated V should be consistent with the observed pixel intensities I ; 2) $\|V\|_1$ is the sparse prior term that forces the visibility of negligible light directions, i.e., those do not affect the observation should be zero. With this term, the solution would favor a sparse set of visible light directions; 3) $\|V - 0.5\|_\infty$ is the binary prior term. It is used to clamp the elements of V with high values to 1 and low values to 0. Combining $\|V\|_1$ and $\|V - 0.5\|_\infty$ with weighting factors allows us to obtain an *approximate* binary solution; and 4) $\|\nabla V\|_1$ is the total variation term, i.e., to bias towards a solution with compact visible areas. Under this formulation, Eqn. 8 can be modeled as a second order cone problem (SOCP), we use the CVX optimization toolbox [10] to obtain the solution.

We have validated our approach on both synthetic and real experiments.

Normal Recovery Sub-problem Recall that under our formulation, the per-direction visibility vector \hat{V} obtained by solving Eqn. 8 is not binary. We need to regularize \hat{V} to binary using a predefined threshold k , we set $k = 0.35$ in our experiment. This results in a binary visibility vector \tilde{V} . Now that we have both the visibility vector and albedo, we can refine the estimation of normal \hat{n} by solving for the following least square problem:

$$\begin{aligned} \bar{\rho}, \bar{n} &\leftarrow \arg \min_{\rho, \hat{n}} \|\hat{\rho} B[W(\hat{n}) * \tilde{V}] - I\|_2 \\ &\text{Subject to } \|\hat{n}\|_2 = 1 \end{aligned} \quad (9)$$

Specifically, we relax the constraint to $\|\hat{n}\| \leq 1$ and directly solve it via constrained least square minimization (e.g., the *lsqlin* function in Matlab). The results are normalized to have unit length \bar{n} . Next, we use \bar{n} to update W . We repeat the process to iteratively improve the visibility and normal estimation. In our implementation, we preset the maximum number of iterations (15 in most cases). In order to perform CS acquisition, it is critical that the sampling basis are incoherent. We choose to use the Hadamard transform (HT) [29, 9] to generate the sampling pattern.

4. Experiments

4.1. Synthetic Experiment

For synthetic experiments, we render the scene using the POV-Ray ray tracer (www.povray.org). We place an object at the origin of the coordinate system and the camera on z-axis viewing towards the negative z direction. We create two 4D light field sources, each covers nearly a hemisphere. Each light field source consists of 88 uniformly distributed lighting directions. We generate 50 Hadamard patterns and render the corresponding images.

We test our technique on two scenes: the turtle scene and the tentacle scene. We apply our optimization scheme Eqn. 8 and 9 to iteratively compute the visibility term and surface normal. For Eqn. 8, we use the gradient of V on the grid to compute the total variation term. We set the weighting factors as $\lambda_1 = 0.25$, $\lambda_2 = 10.5$ and $\lambda_3 = 0.25$ and the maximum number of iterations as 15.

Next, we compare our technique with [12]. [12] requires uniformly distributed lighting directions. We emulate the sampling by using uniformly partitioning a geo-sphere and map the position of the vertices to direction. We render 161 images, each with only one lighting direction on. The complete results with the ground truth ambient occlusion and albedo is shown in Fig. 3. Overall, our technique produces comparable results as [12]. Furthermore, for points lying in a valley, our method produces better estimations since [12] assumes cone-shaped visibility model which is no longer valid in this case. Also note that [12] requires all 3 color channels in their optimization algorithm. If one color has zero albedo, it will fail to produce accurate AOs, (e.g., the tentacle and tail of the turtle). Finally their estimated AO tends to be smoother since they use statistical behavior of visibility within a patch of pixels.

Further, our algorithm can significantly reduce the size of input. Fig. 4 plots the rooted means squared error of AO and albedo w.r.t the number of input images. Our results show that with a small number of inputs (e.g., 40), we can produce comparable results as [12].

4.2. Photometric Stereo

We also show that our technique can improve photometric stereo by exploring two public datasets SCHOLAR and FROG which provide the images captured under calibrated directional lightings. We emulate our coded lighting process by summing the multiple images that correspond to the code pattern. In our experiment, we synthesize 10 images from the Scholar dataset and 6 for the Frog dataset and test our algorithm for normal recovery. The first column in Fig. 5 shows a sample image.

Recall that both datasets use a sparse set of directional lights: 12 for the SCHOLAR and 8 for the FROG. To robustly apply total variation in our ℓ_∞ regularized ℓ_1

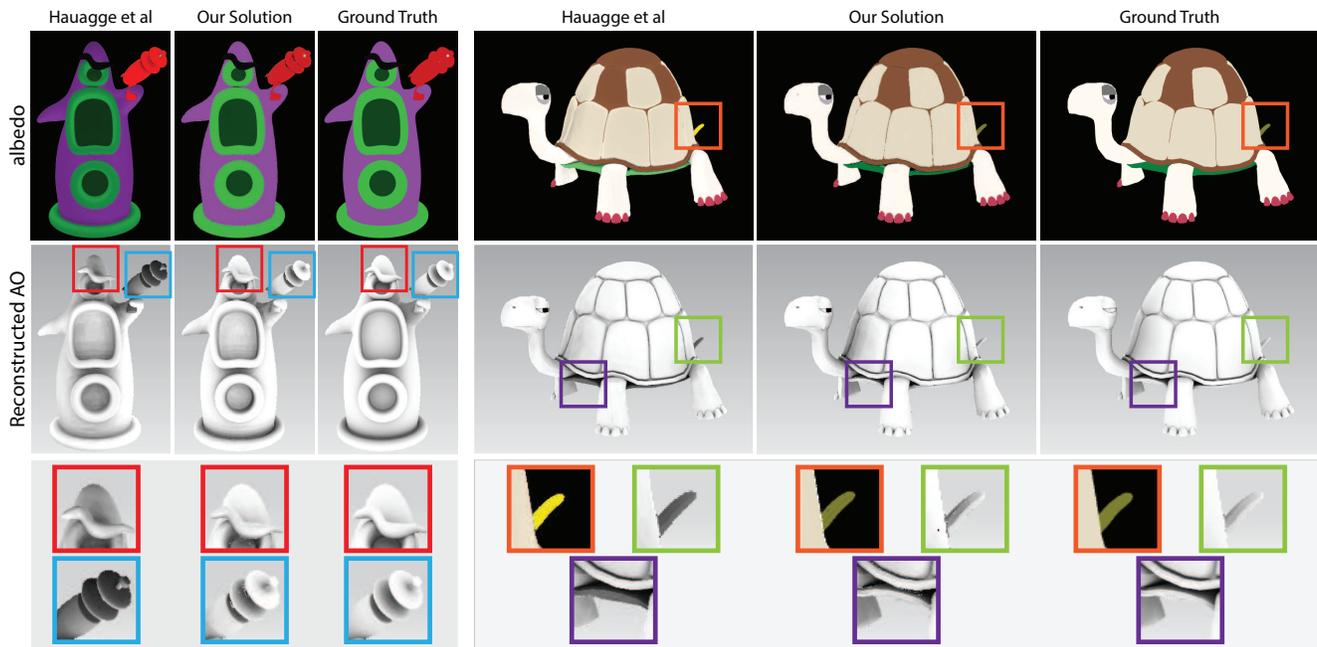


Figure 3: Comparison between our result and [12] on AO estimation on synthetic datasets. Left column shows the ambient occlusion results on the Tentacle scene. Right three columns shows the recovery albedo and ambient occlusion with closeup views.

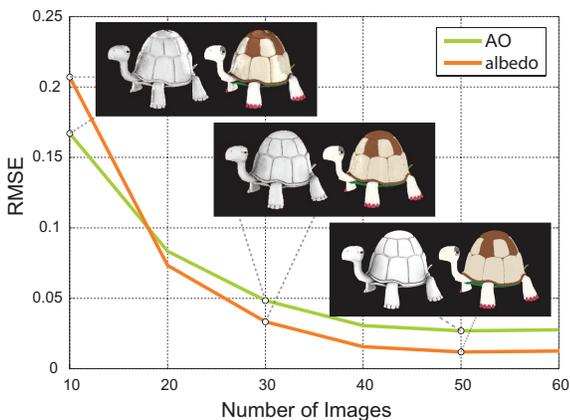


Figure 4: The Root Mean Square Error of AO and albedo estimations with different input sizes.

optimization, we find, for each direction, its nearest lighting direction and compute the corresponding difference in visibility term. For both datasets, we set the weighting factors as $\lambda_1 = 0.25$, $\lambda_2 = 8.5$ and $\lambda_3 = 0.05$ and the maximum number of iterations as 15. The second column of Fig. 5 shows our estimated lighting visibility and the middle two columns show our estimated normal and albedo maps. The last two columns shows the computed normal and its error map w.r.t the ground truth normal. Our results are comparable to state-of-the-art solutions. However, we

not only obtain the normal map but also the AO map, which is largely missing using classical photometric stereo techniques.

4.3. Real Data

Light Field Probe: To produce encodable directional light sources, we use the light field probe that consists of a microlenslet array. To cover both sides of the object, we use two Epson HC8350 projectors and two Fresnel 300 micro lenslet arrays to create the left and right light field source. The projected resolution is 1920 x 1080. We attach a diffuser onto the back of the micro lenslet array. To expand over the full hemisphere, the left micro lenslet arrays is set to have an included angle about 75° with the horizontal plane, and the right micro lenslet arrays is 45° . The object is placed on a rotation table. When we rotate the object one cycle, the illumination field will spread the whole hemisphere. A PointGray Flee 3 camera is on the rotation axis and viewing the object. The captured image has a resolution of 4096 x 2160. The camera is aligned with the rotation axis of the turning table. We also use the telephoto lenses to minimize view-dependent occlusions.

Calibration: To calibrate the emit directions of our light source, we use a mirror ball of 2 inch diameter. The center of the ball is aligned with the rotation axis. We encode the projection pattern for each microlens to enforce that all

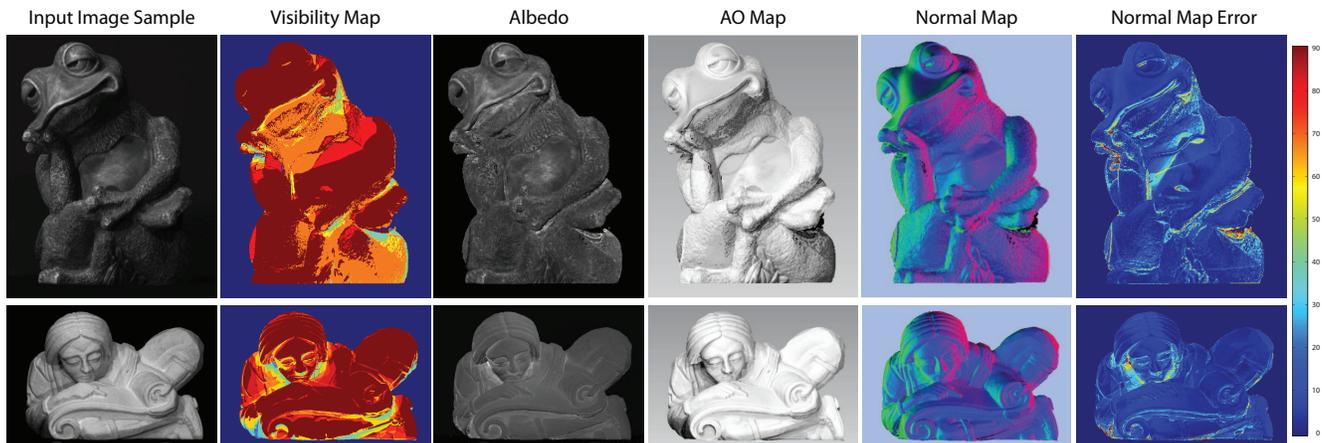


Figure 5: Ambient occlusion and albedo estimation using our approach on the scholar and frog dataset.

light rays through microlenslets have a uniform direction. We then detect the location of specular highlight on the ball and use it to compute the lighting direction. Fig. 7 shows our light field probe setup with angularly coded color wheel pattern. The emitted rays with the same directions have the same color.

Due to low angular resolution of the micro lenslets, we need to rotate the object about 60 degree each time and we capture 12 images at each position with different lighting patterns. At each position, we solve for the visibility and normal independently. The visibility measures for the same patch are stitched over the hemisphere. We choose the one that incurs least error (with respect to input) as the estimated normal and albedo. The weighting factors we use in our experiments are: $\lambda_1 = 0.1$, $\lambda_2 = 10.5$ and $\lambda_3 = 0.2$.

Fig. 9 shows the reconstruction results on a toy duck. In our experiment, a total of 72 images were captured at a resolution of 4096 x 2160. To reduce image noise and computational cost, we downsample the images and only process regions of interest (about 500x400 image resolution). To calibrate the camera response curves, we use the color calibration toolbox provided [14].

We use the Hadamard pattern with 75% of the lighting directions enabled at each capture. This is more advantageous than acquiring an image under a single light direction since it provides more lights to reduce noise. Fig. 9 shows our recovered AO, normal map and albedo map. Our results are comparable to the synthetic ones. The recovered albedo map, however, exhibits some color bleeding artifacts (e.g., the head and belly region of the duck), the might be caused by the error when we align the rotated images together. Also notice that our approach is per-pixel based and does not impose smoothness priors to adjacent pixels and therefore our normal map appears noisier than traditional photometric stereo. Additional results can be found in the supplementary materials.

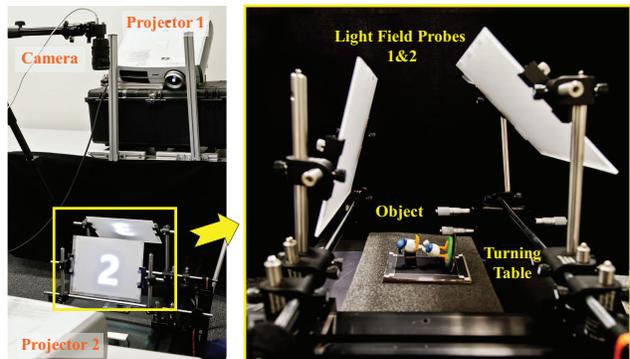


Figure 6: Left: our system uses two light field probes to emulate encodable directional light sources. Right: Closeup views of the light field probes.

5. Conclusions and Discussions

We have presented a novel computational imaging solution for recovering the ambient occlusion (AO) map of an object. Our technique borrows the compressive sensing framework by capturing the object under strategically encoded lighting directions. We have developed a sparsity-prior based solution for robustly recover the surface normal, the surface albedo, and the visibility function. To physically implement our scheme, we have constructed an encodable light source by exploiting the recent light field probe designs. Experiments on synthetic and real scenes have shown that our approach is reliable and accurate and can greatly reduce the input size.

Compared with [12], our technique uses much fewer images but requires known lighting directions, which is a major limitation. In the future, we will explore eliminating this requirement, e.g., by correlating the 3D geometry of the object with coding patterns. Our solution is also per-pixel

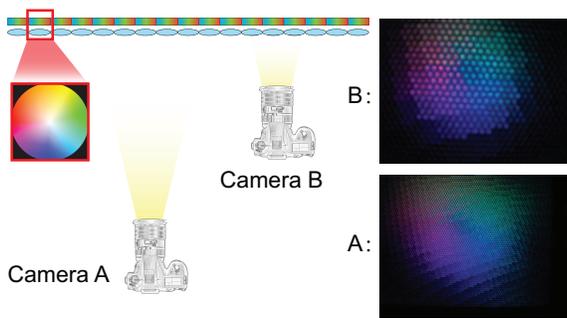


Figure 7: Directional invariance of light field probes. We project identical color wheel images for every microlenslet so that each lighting direction maps to a unique color.

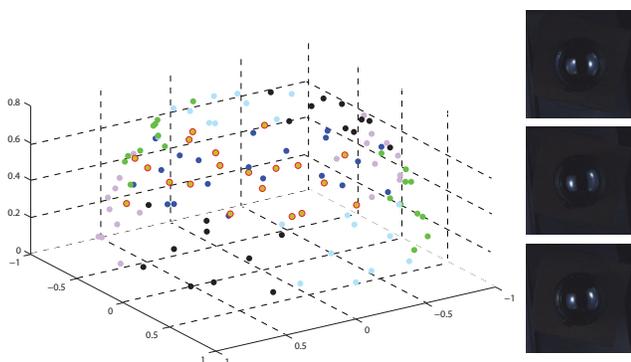


Figure 8: Calibrating the lighting direction of a mirror ball. Left: lighting directions represented by points on a uniform sphere. Identical directions are shown in the same color. Right: we detect specular highlights on the mirror ball for calibrating the lighting direction.

based and does not consider coherence between neighboring pixels. A straightforward extension is to incorporate the Markov Random Field work to enforce smoothness priors on normal and albedo estimations. Finally, same as most existing photometric stereo techniques, our method cannot handle subsurface scattering and inter-reflection. It is our important future direction to include light transport analysis [31] and acquisition [24] in our AO measurement.

Acknowledgements

This project was supported by the National Science Foundation under grant IIS-1422477.

References

- [1] O. Aldrian and W. A. Smith. Inverse rendering of faces on a cloudy day. In *Computer Vision—ECCV 2012*, pages 201–214. Springer, 2012. 1, 2
- [2] S. Bell, K. Bala, and N. Snavely. Intrinsic images in the wild. *ACM Trans. on Graphics (SIGGRAPH)*, 33(4), 2014. 1

- [3] S. Bell, P. Upchurch, N. Snavely, and K. Bala. Opensurfaces: A richly annotated catalog of surface appearance. *ACM Transactions on Graphics (TOG)*, 32(4):111, 2013. 1
- [4] M. Ben-Ezra, J. Wang, B. Wilburn, X. Li, and L. Ma. An led-only brdf measurement device. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8, June 2008. 2
- [5] E. J. Candès and M. B. Wakin. An introduction to compressive sampling. *Signal Processing Magazine, IEEE*, 25(2):21–30, 2008. 2, 4
- [6] M. Chandraker, S. Agarwal, and D. Kriegman. Shadowcuts: Photometric stereo with shadows. In *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*, pages 1–8. IEEE, 2007. 1, 2
- [7] A. Colaço, A. Kirmani, G. A. Howland, J. C. Howell, and V. K. Goyal. Compressive depth map acquisition using a single photon-counting detector: Parametric signal processing meets sparsity. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 96–102. IEEE, 2012. 2
- [8] O. Cossairt, S. Nayar, and R. Ramamoorthi. Light field transfer: global illumination between real and synthetic objects. In *ACM Transactions on Graphics (TOG)*, volume 27, page 57. ACM, 2008. 2
- [9] P. Debevec. The light stages and their applications to photo-real digital actors. In *SIGGRAPH Asia*, Nov. 2012. 2, 4
- [10] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, Mar. 2014. 4
- [11] R. Grosse, M. K. Johnson, E. H. Adelson, and W. T. Freeman. Ground truth dataset and baseline evaluations for intrinsic image algorithms. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 2335–2342. IEEE, 2009. 1
- [12] D. Hauage, S. Wehrwein, K. Bala, and N. Snavely. Photometric ambient occlusion. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 2515–2522. IEEE, 2013. 1, 2, 3, 4, 5, 6
- [13] J. Holloway, A. Sankaranarayanan, A. Veeraraghavan, and S. Tambe. Flutter shutter video camera for compressive sensing of videos. In *Computational Photography (ICCP), 2012 IEEE International Conference on*, pages 1–9, April 2012. 2
- [14] I. Ihrke. Matlab color calibration toolbox. 6
- [15] Y. Ji, J. Ye, and J. Yu. Reconstructing gas flows using light-path approximation. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 2507–2514. IEEE, 2013. 2
- [16] M. Kowalski. Sparse regression using mixed norms. *Applied and Computational Harmonic Analysis*, 27(3):303–324, 2009. 4
- [17] E. H. Land and J. McCann. Lightness and retinex theory. *JOSA*, 61(1):1–11, 1971. 2
- [18] M. Levoy and P. Hanrahan. Light field rendering. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 31–42. ACM, 1996. 2

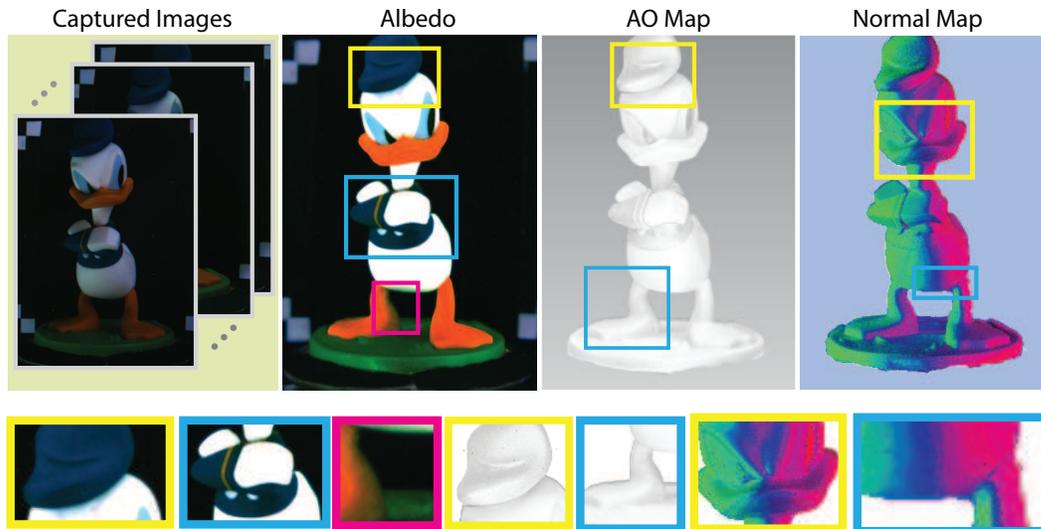


Figure 9: The AO, normal and albedo estimation results on a real Donald duck. The bottom row is the closeup view.

- [19] Y. Li and M. S. Brown. Single image layer separation using relative smoothness. *CVPR*, 2014. 2
- [20] F. Lu, Y. Matsushita, I. Sato, T. Okabe, and Y. Sato. Uncalibrated photometric stereo for unknown isotropic reflectances. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 1490–1497. IEEE, 2013. 1, 2
- [21] M. Lustig, D. Donoho, and J. M. Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic resonance in medicine*, 58(6):1182–1195, 2007. 2
- [22] K. Marwah, G. Wetzstein, Y. Bando, and R. Raskar. Compressive light field photography using overcomplete dictionaries and optimized projections. *ACM Transactions on Graphics (TOG)*, 32(4):46, 2013. 2
- [23] V. Masselus, P. Peers, P. Dutré, and Y. D. Willems. Relighting with 4d incident light fields. In *ACM Transactions on Graphics (TOG)*, volume 22, pages 613–620. ACM, 2003. 2
- [24] M. O’Toole, F. Heide, L. Xiao, M. B. Hullin, W. Heidrich, and K. N. Kutulakos. Temporal frequency probing for 5d transient analysis of global light transport. *ACM TRANSACTIONS ON GRAPHICS*, 33(4), 2014. 7
- [25] M. O’Toole and K. N. Kutulakos. Optical computing for fast light transport analysis. *ACM Trans. Graph.*, 29(6):164, 2010. 1
- [26] P. Peers, D. K. Mahajan, B. Lamond, A. Ghosh, W. Matusik, R. Ramamoorthi, and P. Debevec. Compressive light transport sensing. *ACM Transactions on Graphics (TOG)*, 28(1):3, 2009. 2
- [27] S. Qaisar, R. Bilal, W. Iqbal, M. Naureen, and S. Lee. Compressive sensing: From theory to applications, a survey. *Communications and Networks, Journal of*, 15(5):443–456, Oct 2013. 2
- [28] S. D. Roth. Ray casting for modeling solids. *Computer graphics and image processing*, 18(2):109–144, 1982. 1
- [29] Y. Y. Schechner, S. K. Nayar, and P. N. Belhumeur. Multiplexing for optimal lighting. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 29(8):1339–1354, 2007. 4
- [30] D. Schneider. New camera chip captures only what it needs. In *IEEE JOURNALS and MAGAZINES*, volume 50, pages 13–14. IEEE, 2013. 2
- [31] S. M. Seitz, Y. Matsushita, and K. N. Kutulakos. A theory of inverse light transport. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, volume 2, pages 1440–1447. IEEE, 2005. 7
- [32] K. Sunkavalli, T. Zickler, and H. Pfister. Visibility subspaces: uncalibrated photometric stereo with shadows. In *Computer Vision—ECCV 2010*, pages 251–264. Springer, 2010. 1, 2
- [33] M. F. Tappen, E. H. Adelson, and W. T. Freeman. Estimating intrinsic component images using non-linear regression. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 2, pages 1992–1999. IEEE, 2006. 2
- [34] S. Wang and N. Rahnavard. Binary compressive sensing via sum of l1-norm and l(infinity)-norm regularization. In *Military Communications Conference, MILCOM 2013 - 2013 IEEE*, pages 1616–1621, Nov 2013. 4
- [35] Y. Weiss. Deriving intrinsic images from image sequences. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, volume 2, pages 68–75. IEEE, 2001. 1, 2
- [36] G. Wetzstein, R. Raskar, and W. Heidrich. Hand-held schlieren photography with light field probes. In *Computational Photography (ICCP), 2011 IEEE International Conference on*, pages 1–8. IEEE, 2011. 2
- [37] G. Wetzstein, D. Roodnick, W. Heidrich, and R. Raskar. Refractive shape from light field distortion. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 1180–1186. IEEE, 2011. 2