Photometric Ambient Occlusion
Supplemental Material

Daniel Hauagge Scott Wehrwein Kavita Bala Noah Snavely
Cornell University
{hauagge, swehrwein, kb, snavely}@cs.cornell.edu

Contents

1. Test objects

In Figure 1 we show an image of our two test objects TENTACLE and LIGHTWELL. We include the 3D model for both as PLY files together with the supplemental material.

In Figure 2 we show a schematic drawing of one block of LIGHTWELL together with the local visibility angle α. The top left hole is the deepest with $\alpha = 10°$, while $\alpha = 90°$ corresponds to a flat surface.

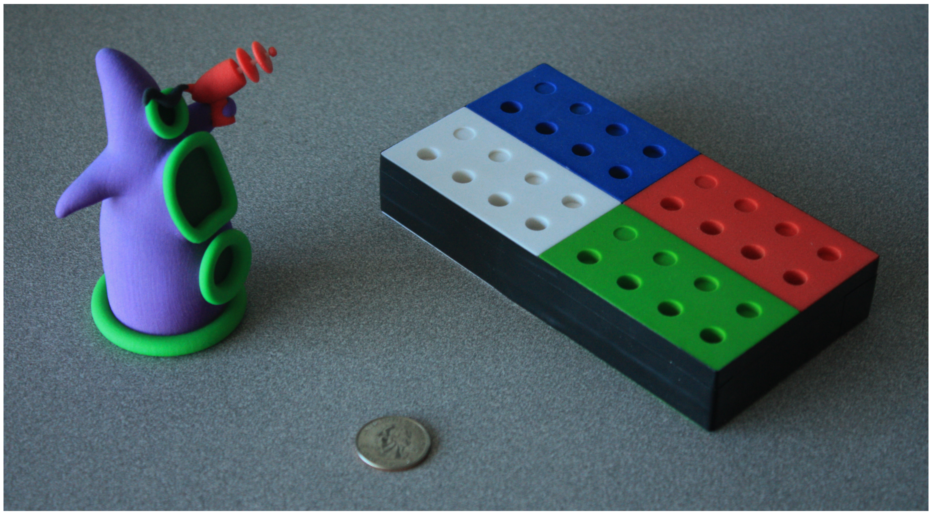

Figure 1. Images of TENTACLE and LIGHTWELL used for tests. Black tape surrounding LIGHTWELL was added to reduce sub-surface scattering that resulted form light shining on the side of the object.

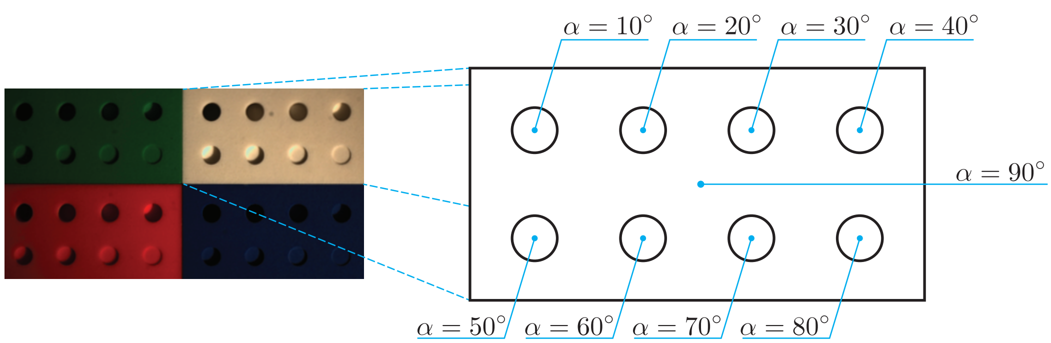

Figure 2. One block from LIGHTWELL together with local visibility angle α (smaller values correspond to deeper holes while $\alpha = 90°$ corresponds to a flat surface).

2. Results for LIGHTWELL

In Figure 3 we show results for LIGHTWELL for our algorithm (using the model with both direct and ambient term) and Weiss+Retinex. First row shows a sample image from the sequence (the same for both columns). Second row shows the ambient occlusion (Weiss+Retinex does not compute ambient occlusion). Ideally the ambient occlusion for the different color blocks would be identical. We can see that the flat regions are roughly the same brightness and that corresponding crevices have similar gray levels. For a precise measurement of the error in the ambient occlusion see the plot in Figure 7 in the paper. In the third row we show the reflectance ρ estimated by both algorithms. One can see that in general the albedo is flatter for our algorithm as the crevices are less noticeable. Nevertheless deep crevices (top left of each block) are still problematic, which was also observed in plot in Figure 7 of the paper. The impact of the incorrect values for ρ can be seen in the fourth row, where we show the computed illumination image L for the image shown in the first row. We see that where our model successfully obtained the albedo for the bottom of the holes these show up as white in L, while for Weiss+Retinex we see the color of the object as ρ for these regions was estimated as darker than it is in reality.

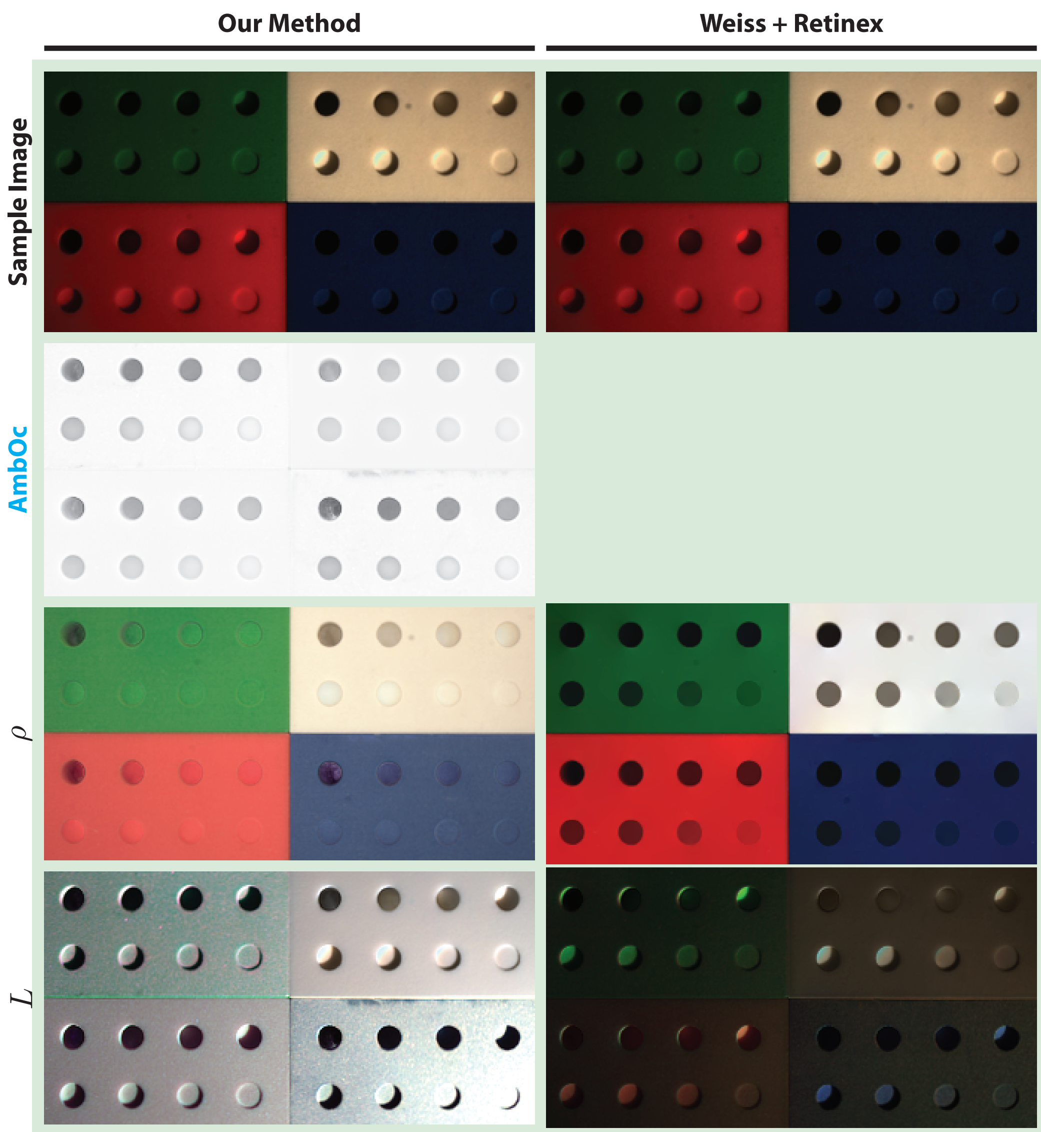

Figure 3. Results for LIGHTWELL.

3. Comparison of our results with others reported in the literature

In §6.3 of the paper we report results on the MIT intrinsic benchmark [2], shown in detail in the plot in Figure 9. More recent work in intrinsic image decomposition [1, 3, 4] also report results on the benchmark but use different subsets of the images. For lack of space in the paper we moved a more detailed comparison with these works to the supplemental material. In table 1 we show the local mean squared error (LMSE) for individual images of our algorithm for the 1st (only direct light) and 2nd estimates (direct and ambient term), together with results from other work when available. In last three columns of the last row we show our average on the same subset of images as reported by [1, 3, 4]. On all cases our algorithm outperforms these works.

Table 1: Detailed results for MIT intrinsic image benchmark.

	k-D (Direct light)			**k-DA (Direct + Ambient light)**					
	ρ	**L**	**Avg**	**ρ**	**L**	**Avg**	**Barron and Malik**	**Shen and Yeo**	**Shen, Yang, and Li**
apple	0.0060	0.0060	0.0060	0.0060	0.0060	0.0060			0.0102
box	0.0040	0.0040	0.0040	0.0050	0.0050	0.0050		0.0018	0.0115
cup1	0.0030	0.0020	0.0020	0.0030	0.0020	0.0020		0.0042	0.0055
cup2	0.0030	0.0010	0.0020	0.0030	0.0010	0.0020	✓	0.0050	0.0073
deer	0.0270	0.0160	0.0210	0.0370	0.0210	0.0290	✓		0.0319
dinosaur	0.0150	0.0120	0.0140	0.0160	0.0070	0.0120			0.0212
frog1	0.0200	0.0180	0.0190	0.0290	0.0260	0.0270		0.0526	0.0287
frog2	0.0560	0.0120	0.0340	0.0530	0.0170	0.0350	✓	0.0435	0.0238
panther	0.0080	0.0060	0.0070	0.0240	0.0140	0.0190		0.0078	0.0049
paper1	0.0040	0.0040	0.0040	0.0100	0.0080	0.0090		0.0014	0.0125
paper2	0.0070	0.0040	0.0060	0.0090	0.0060	0.0080	✓	0.0027	0.0161
pear	0.0060	0.0050	0.0050	0.0060	0.0040	0.0050	✓		0.0102
phone	0.0110	0.0080	0.0100	0.0350	0.0130	0.0240			0.0112
potato	0.0110	0.0080	0.0090	0.0060	0.0060	0.0060	✓		0.0140
raccoon	0.0110	0.0090	0.0100	0.0150	0.0110	0.0130	✓	0.0048	0.0077
squirrel	0.0190	0.0240	0.0220	0.0200	0.0250	0.0230			0.0374
sun	0.0040	0.0050	0.0050	0.0070	0.0050	0.0060	✓	0.0023	0.0070
teabag1	0.0070	0.0160	0.0120	0.0120	0.0330	0.0230	✓	0.0268	0.0631
teabag2	0.0030	0.0110	0.0070	0.0120	0.0200	0.0160		0.0151	0.0307
turtle	0.0170	0.0200	0.0190	0.0200	0.0260	0.0230	✓	0.0174	0.0247
average	0.0121	0.0095	0.0109	0.0164	0.0128	0.0150	0.0190	0.0150	0.0190
Our average on the same subset							0.0123	0.0101	0.0109

4. Full set of results for the MIT intrinsic image benchmark

In Table 2 we show results for individual objects on the MIT intrinsic image benchmark. The set of objects corresponds to the one used to compute the averages shown in Figure 8 of the paper. We show here the gray scale results, which is what the benchmark uses. In the paper we show results in color in Figures 5 and 8.

Table 2: Full set of results for the MIT intrinsic image benchmark.

Original Image		κ**-D** **Direct Light**	κ**-DA** **Direct + Ambient Light**	**Weiss + Retinex**
	ρ		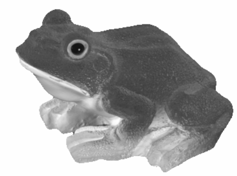	
	L		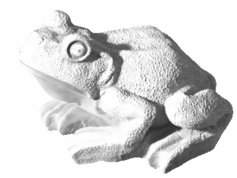	
	ρ			
	L			
	ρ			

Continued on next page

Table 2 – *Continued from previous page*

Original Image		κ-**D** **Direct Light**	κ-**DA** **Direct + Ambient Light**	**Weiss + Retinex**
	L			
	ρ			
	L			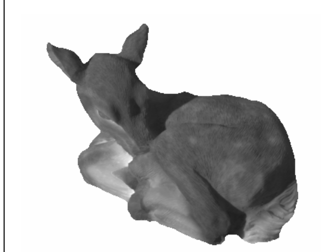
	ρ			
	L			

Continued on next page

Table 2 – *Continued from previous page*

Original Image		κ-**D** **Direct Light**	κ-**DA** **Direct + Ambient Light**	**Weiss + Retinex**
	ρ	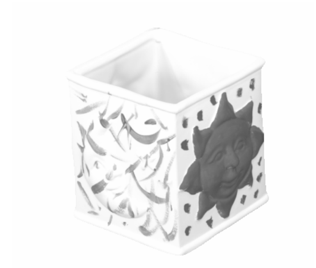	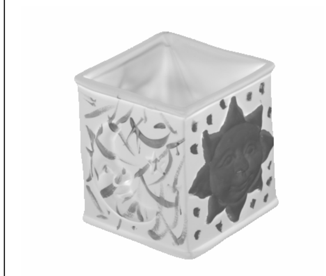	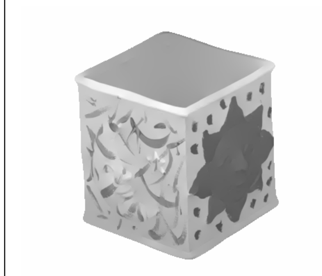
	L		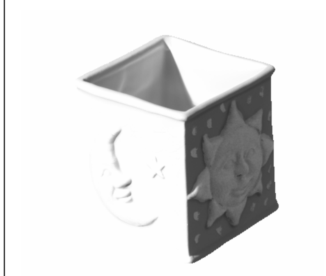	
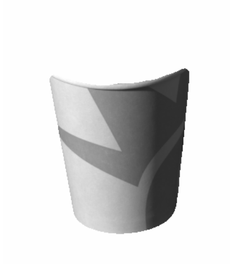	ρ		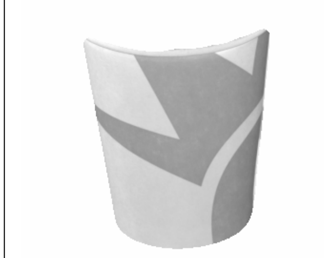	
	L			

Continued on next page

Table 2 – *Continued from previous page*

Original Image		κ**-D** **Direct Light**	κ**-DA** **Direct + Ambient Light**	**Weiss + Retinex**
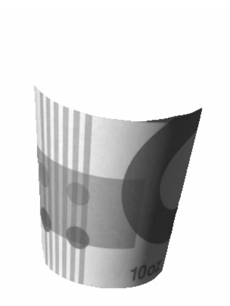	ρ			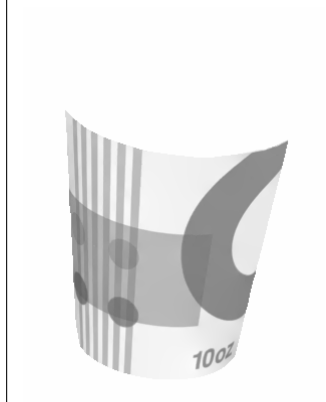
	L	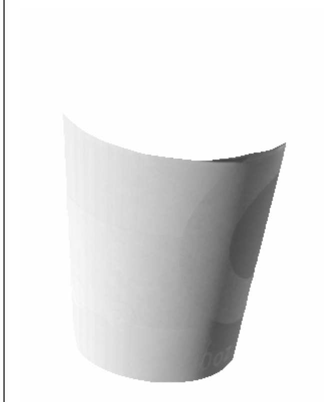		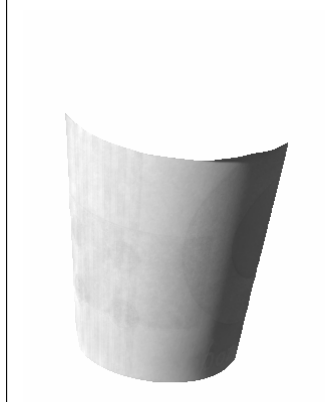
	ρ	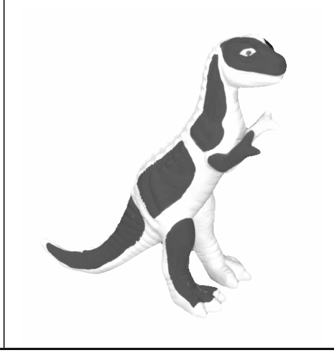	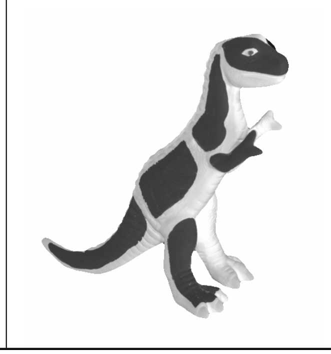	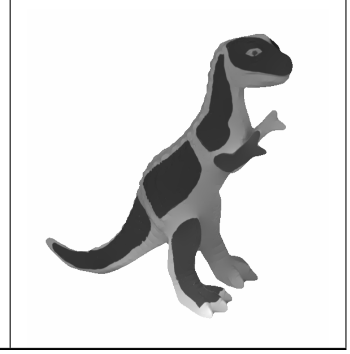

Continued on next page

Table 2 – *Continued from previous page*

Original Image		κ-D Direct Light	κ-DA Direct + Ambient Light	Weiss + Retinex
	L			
	ρ			
	L			
	ρ			
	L			

Continued on next page

Table 2 – *Continued from previous page*

Original Image		κ-D Direct Light	κ-DA Direct + Ambient Light	Weiss + Retinex
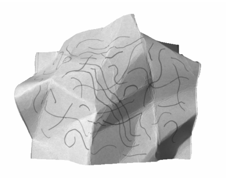	ρ	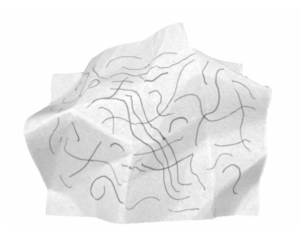	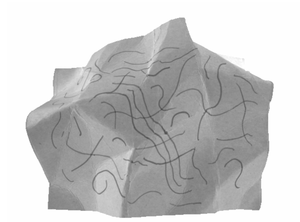	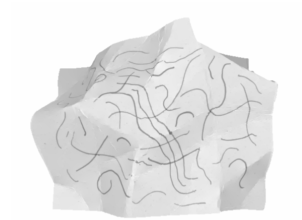
	L	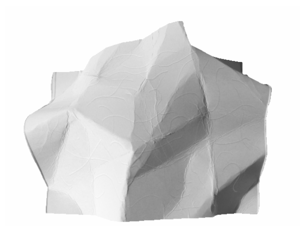	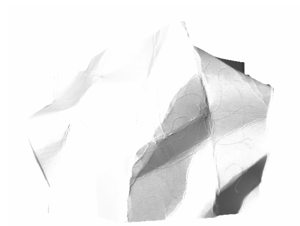	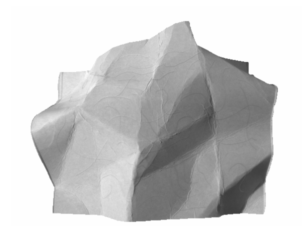
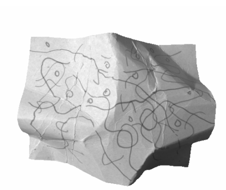	ρ			
	L	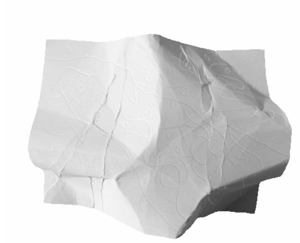	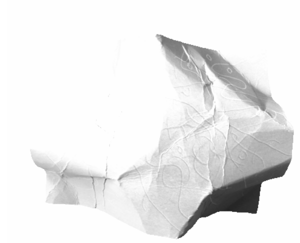	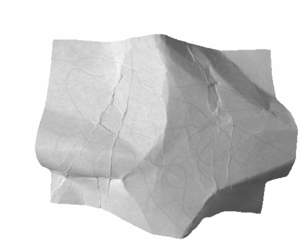
	ρ			

Continued on next page

Table 2 – *Continued from previous page*

Original Image		κ-D Direct Light	κ-DA Direct + Ambient Light	Weiss + Retinex
	L			
	ρ			
	L			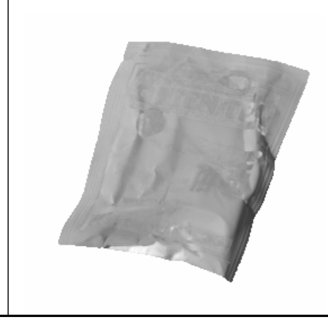

References

[1] J. T. Barron and J. Malik. High-frequency shape and albedo from shading using natural image statistics. In *CVPR*, 2011.

[2] R. Grosse, M. K. Johnson, E. H. Adelson, and W. T. Freeman. MIT Intrinsic Images, 2009. `http://people.csail.mit.edu/rgrosse/intrinsic/`.

[3] J. Shen, X. Yang, X. Li, and Y. Jia. Intrinsic image decomposition using optimization and user scribbles. *Trans. Systems, Man, and Cybernetics*, 2012.

[4] L. Shen and C. Yeo. Intrinsic images decomposition using a local and global sparse representation of reflectance. In *CVPR*, 2011.