

An RGBN Benchmark : Supplementary Material

Sema Berkiten Szymon Rusinkiewicz
Princeton University

3.1. Synthetic Datasets

3.1.1 User Interface

We provide a user interface which is tailored for creating photometric data from arbitrary 3D models automatically. The user can load, rotate, translate, scale 3D models and assign materials and textures to 3D models. The user also can set or load an environment map as well as adjust the amount of ambient illumination, light configurations, and Mitsuba rendering parameters such as number of samples per pixel etc. The user can see a low resolution preview of the rendering and once the result is satisfying the user can produce necessary files ready to be rendered with Mitsuba Renderer.

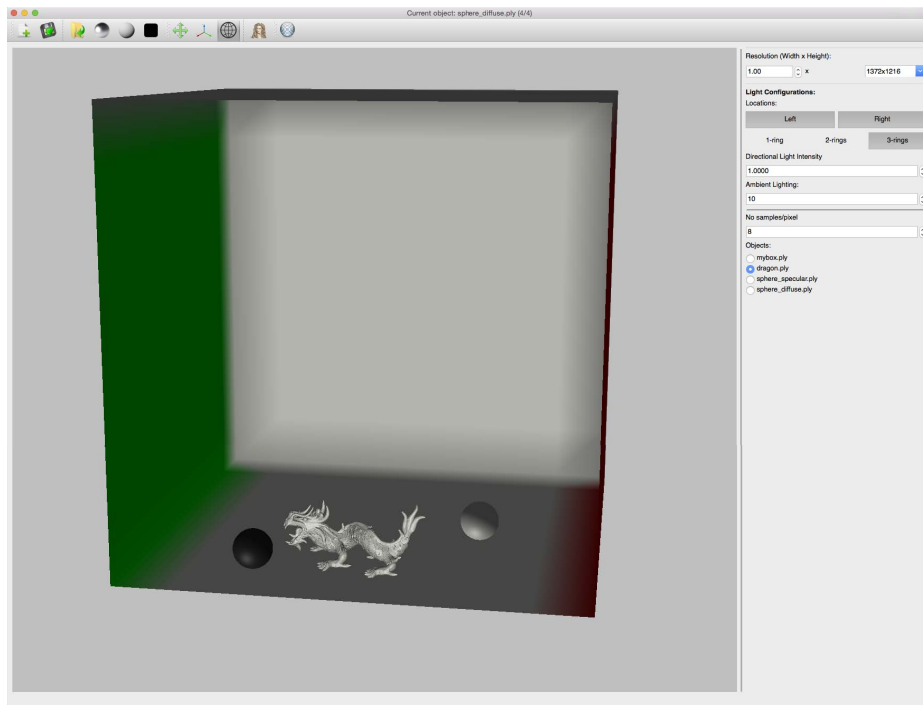


Figure 1: User interface used for producing realistic photometric datasets with ground truths.

3.1.2 A selection from the benchmark

We created 93 synthetic photometric datasets and a selection of them is shown in Figure 2. We would like to note that we did not intend to create a large database but let it grow by accepting new submissions from other users of our system.




Figure 2: Some synthetic datasets from the benchmark.

3.2. Real-world Photometric Datasets

We captured photometric data under varying illumination of the real-world objects in Figure 3.a. We used the same setup as [1]: chrome and white spheres are placed around the object to estimate light direction and intensity because a hand-held flash is used as the light source. 3D models of those object are reconstructed by scanning with a NextEngine laser scanner as shown in Figure 3.b. Other real-world datasets that we use for justifying the synthetic benchmark are shown in Figure 4.



Figure 3: (a) Real-world objects captured under varying illumination with calibration spheres included in the scene, (b) 3D scanned models of the real-world objects on the left.




Figure 4: Real-world objects that are used for various experiments to justify the synthetic datasets, refer to Section 4. Five spheres with various materials and cardboards with different angles which are indicated under the images.

4. Comparison of the Real and Synthetic Datasets

We compare error patterns for 2 more objects in addition to the penguin included in the paper. We picked a sphere and a folded cardboard, because it is easy to create synthetic 3D models of them. The results are shown in Figure 5. We can see the similar error patterns between columns 1 and 2 and between columns 3 and 4.

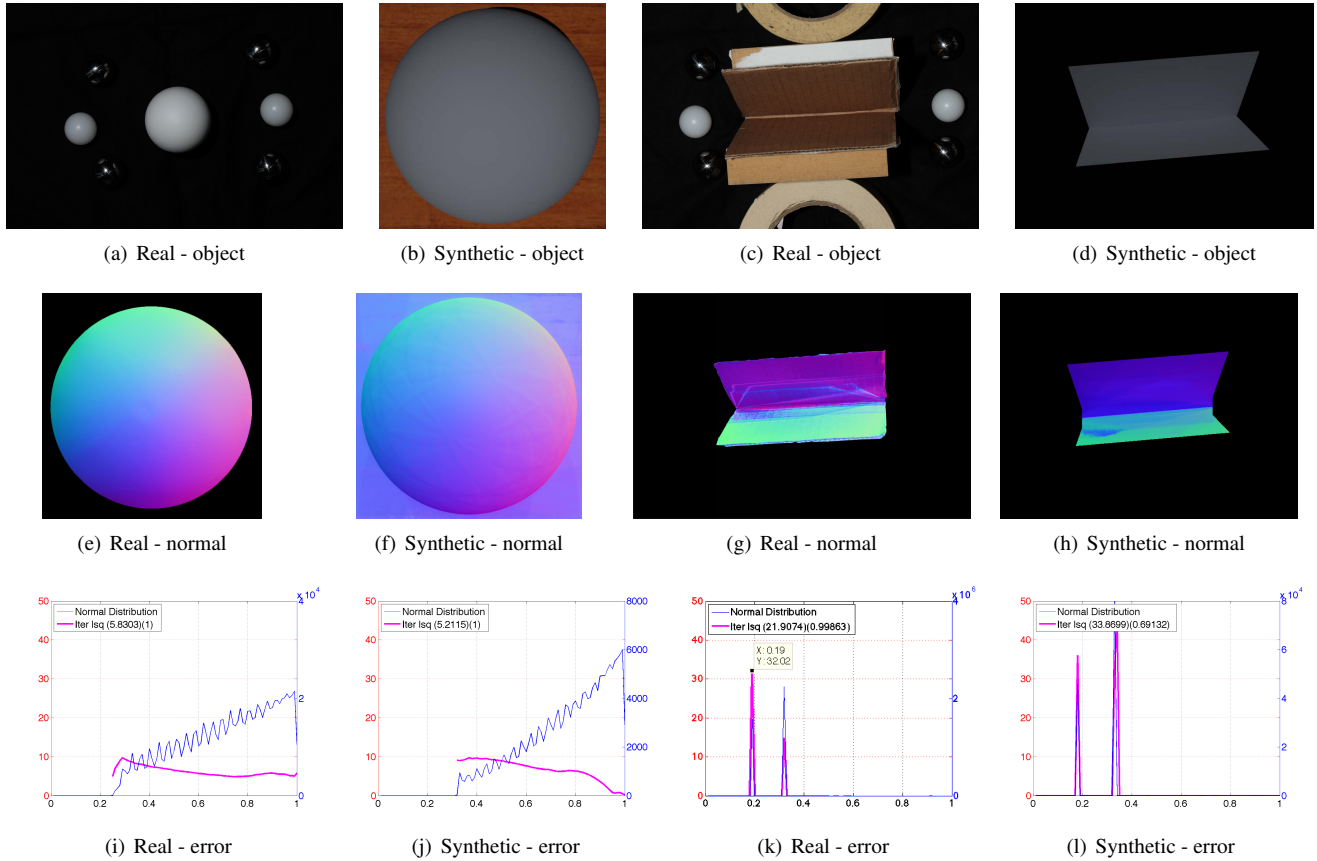
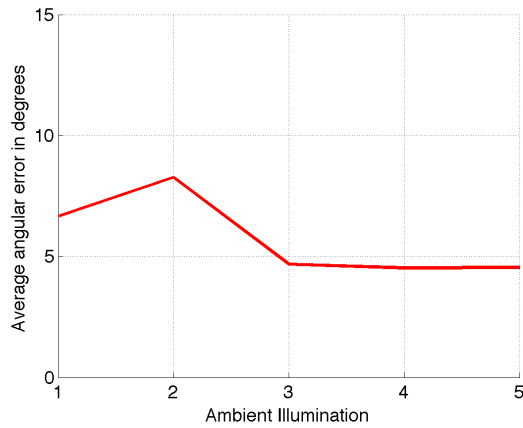


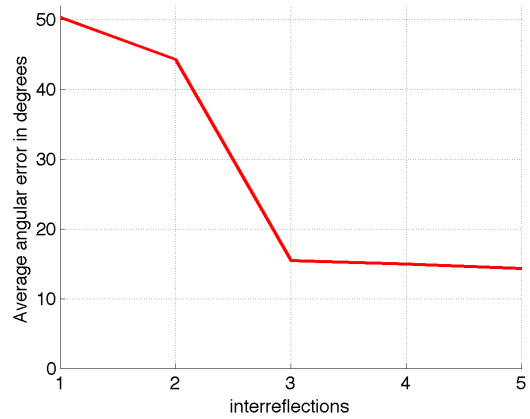
Figure 5: Evaluations with read-world datasets. We captured some real-world objects and created their synthetic versions and compare error patterns.

We also tried to conduct several experiments with real-world datasets similar to the synthetic ones in Section 6. However, it was hard to control and isolate single factor at a time, that is why there are some variations at the results. The main factor which might lead to additional errors in the experiments could be the uneven light positions because we used a hand-held flash as the light source and it is impossible to replicate the same light positions in every capture.

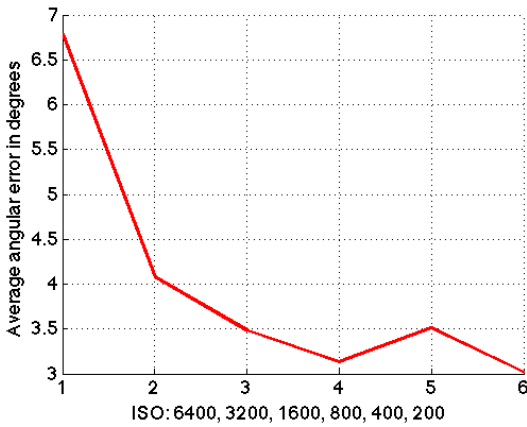
To test the effect of uncompensated illuminations, we captured a white diffuse sphere with increasing amount of ambient lighting and plot the average angular errors as the ambient lighting increases as shown in Figure 6.a. The effect of different amounts of inter-reflection explored by capturing a folded cardboard with different angles and the errors are shown in Figure 6.b. Image noise is created by setting the camera different ISO values and capture a white diffuse sphere. As the ISO value decreases the image noise decreases as shown in Figure 6.c. We try to capture a white sphere with different light-ring positions, although it is not very accurate because we do not have fixed light-ring setup but instead we tried to simulate light-ring configurations with hand-held flash (Figure 6.d).



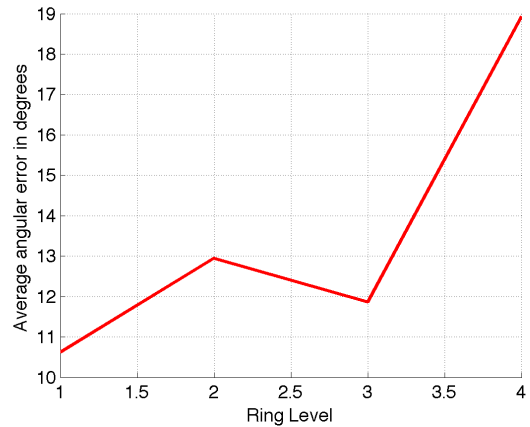
(a) Uncompensated Ambient Illumination



(b) Amount of Inter-reflections



(c) Image Noise



(d) Light-ring Position

Figure 6: Evaluations with read-world datasets.

6. Evaluation of Photometric Stereo Algorithms

We evaluate several PS algorithms as mentioned in the paper. Error statistics including minimum, maximum, mean, median, and standard deviation are shown in Figure 7. Separate error plots are shown in the subsequent subsections.

		CPS				UPS		
		Lsq	RPCS	Iter Lsq	HSH	LDR	Entropy	TV
Ambient - angular	min	0.75	7.99	0.77	13.07	7.23	14.19	12.38
	max	47.35	31.40	21.22	30.64	94.44	73.47	22.90
	mean	20.40	20.48	8.50	25.27	27.42	52.39	17.72
	median	18.73	20.74	6.95	26.78	24.64	65.92	17.19
	std	16.12	7.27	6.62	5.07	21.27	24.62	2.79
Interreflection - angular	min	11.01	22.35	8.45	38.23	75.44	4.28	67.80
	max	67.24	67.26	67.23	67.54	111.20	135.30	152.80
	mean	41.34	42.09	40.91	52.82	91.81	67.29	102.30
	median	44.67	44.57	44.51	55.40	87.65	83.16	93.33
	std	13.76	12.14	14.30	8.37	10.69	41.85	24.27
Material - angular	min	2.96	7.48	3.03	9.37	8.03	9.51	17.62
	max	16.23	9.08	4.55	13.03	22.80	93.11	22.44
	mean	5.69	8.11	3.30	11.55	14.33	53.39	18.92
	median	3.33	8.17	3.09	11.79	13.75	56.32	18.42
	std	4.58	0.55	0.51	1.09	6.31	31.94	1.68
Number of Images - angular	min	0.75	1.77	0.77	13.02	2.67	2.78	12.38
	max	5.74	38.08	6.02	82.75	90.00	76.85	53.56
	mean	1.15	10.18	1.15	25.53	11.64	38.14	20.98
	median	0.87	4.76	0.83	13.43	5.87	8.52	18.51
	std	0.88	13.25	0.94	23.30	20.25	35.41	8.10
Noise - angular	min	0.45	6.03	0.54	84.61	8.73	14.40	17.26
	max	1.77	7.32	1.22	98.39	11.94	61.33	18.07
	mean	1.01	6.74	0.83	91.17	10.40	46.76	17.72
	median	0.96	6.74	0.80	91.00	10.20	60.95	17.76
	std	0.45	0.39	0.23	4.04	0.95	22.28	0.26
Ring Position - angular	min	2.26	2.35	2.29	53.74	2.73	3.16	12.77
	max	44.33	38.10	20.76	102.00	88.21	82.42	97.58
	mean	8.48	16.57	8.43	71.93	36.70	49.07	45.46
	median	2.97	10.84	3.39	68.00	29.80	60.94	45.79
	std	13.67	15.03	12.40	16.35	34.76	34.70	25.67
Roughness - angular	min	3.73	6.75	3.62	12.61	9.20	12.65	19.67
	max	9.75	8.06	3.93	46.57	17.92	85.03	39.82
	mean	5.79	7.28	3.78	19.54	13.41	48.45	29.17
	median	5.74	7.21	3.77	14.95	13.53	46.37	27.65
	std	1.73	0.46	0.11	10.39	2.87	26.53	7.39
Shadow - angular	min	0.46	0.01	0.33	1.27	53.30	29.60	48.44
	max	21.12	64.71	73.46	74.55	105.30	121.10	73.66
	mean	13.13	44.85	38.30	50.51	93.45	55.70	64.95
	median	13.70	51.94	43.24	58.16	99.58	53.65	67.39
	std	5.94	20.14	27.95	23.22	14.50	20.84	7.54
Specularity - angular	min	0.81	6.69	1.07	10.19	5.95	68.73	11.77
	max	5.61	6.90	1.45	11.99	8.82	78.17	14.43
	mean	3.12	6.78	1.20	10.92	7.53	71.52	13.22
	median	2.67	6.78	1.16	10.74	7.50	70.08	13.33
	std	1.76	0.07	0.13	0.63	0.79	3.20	0.86

Figure 7: Statistics of angular error for various experiments.

6.1. Uncompensated Ambient Illumination

In real acquisition setups, it is hard to eliminate the ambient lighting most of the time; so we test the effect of *unknown* ambient illumination on the PS results. Some images with increasing ambient illumination are depicted in Figure 8. From the Figure 9, we can conclude that the error levels become unacceptable even in the presence of small ambient illumination.



Figure 8: Pears with increasing ambient illumination.

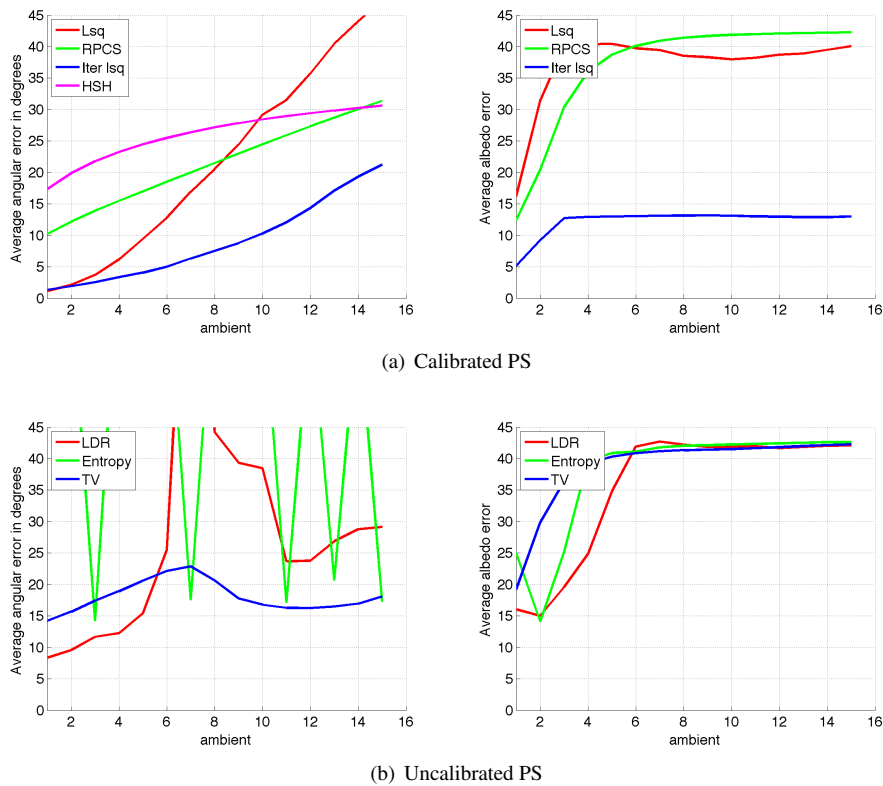


Figure 9: The first column is the angular error and the second column is the error in albedo as the *uncompensated ambient illumination* increases.

6.2. Inter-reflections

We use several cones with the same radius but varying heights and render the views from the bottom of the cones. This configuration is designed to have more inter-reflection as the height of the cone increases, as shown in Figure 10. One can argue that it is not possible to separate the effect of inter-reflections from the shadows in this case. That is why, we keep the light positions closer to the camera to avoid extreme self-shadows, however there is still some influence of shadows on the results.

The plots in Figure 11 demonstrate the angular and albedo errors. Although the photometric stereo algorithms ignore the effect of inter-reflections, we observe that it has an important influence on the results.




Figure 10: Cones with increasing heights, to explore *inter-reflections*.




Figure 11: First column is the angular error, and the second column is the error in albedo as the *inter-reflections* increases.

6.3. Material

We created 8 datasets of the same object with varying materials: some of these are shown in Figure 12. The 8 materials are diffuse, specular, diffuse with subsurface scattering, specular with subsurface scattering, specular-texture with subsurface scattering, diffuse texture, and specular texture in the same order with the error plots shown in Figure 13.




Figure 12: Three example materials on the Pear.

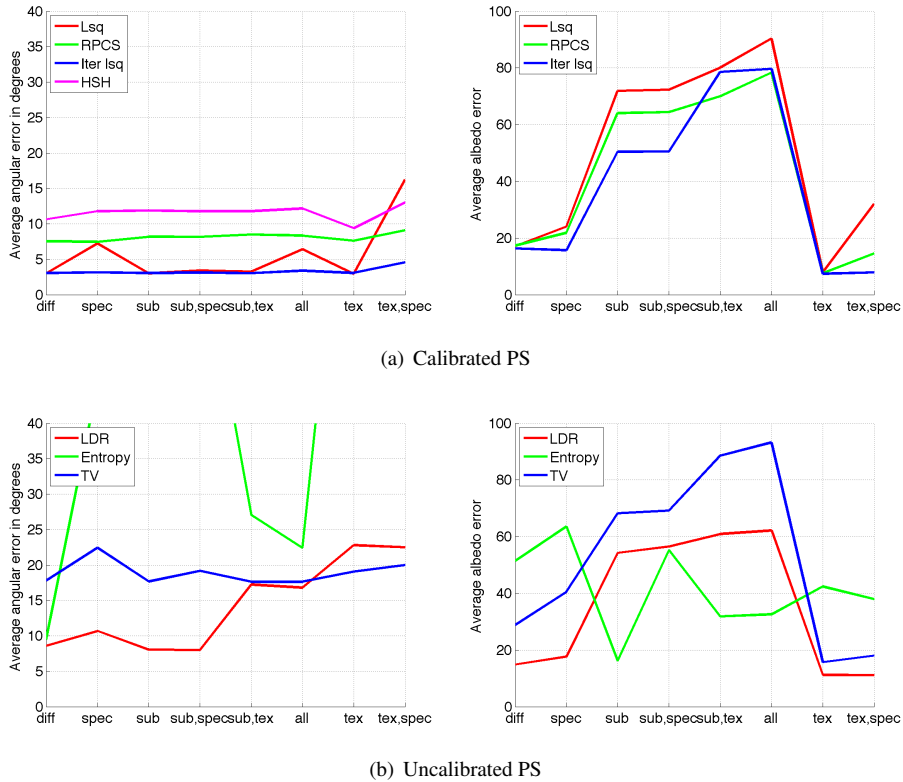


Figure 13: First column is the angular error, and the second column is the error in albedo vs varying materials.

6.4. Number of Images

Although the minimum number of images required by photometric stereo algorithms is generally 3, in practice we need more than 3 to have better accuracy. In Figure 14, we observe that the most algorithms perform reasonably on around 10 images.

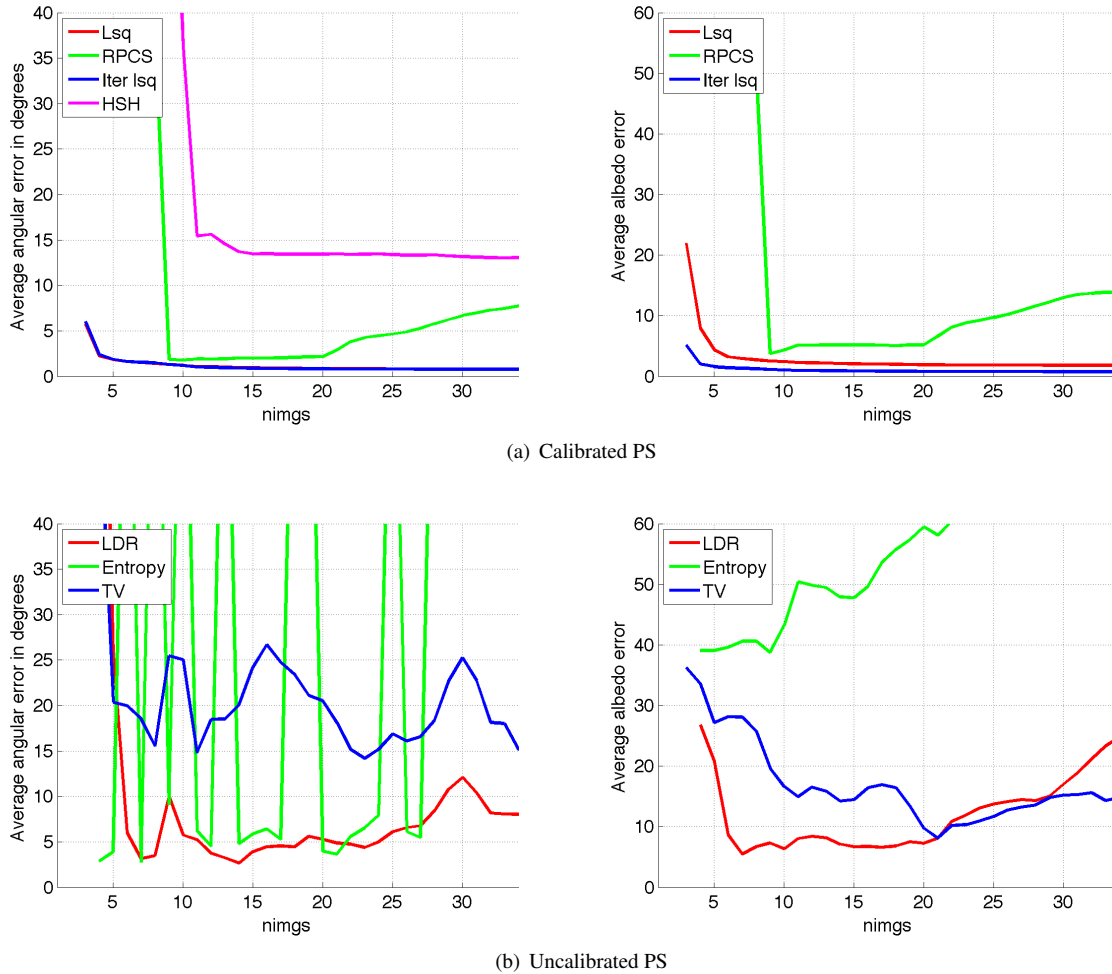


Figure 14: Average errors in normals (the first column) and albedo (the second column) as the *number of images* increases.

6.5. Noise

We explore the effect of Poisson image noise, which is the best known representation of camera noise, on the quality of the photometric stereo results. We artificially create Poisson noise and add it to all rendered images of a pear (see Figure 15) in increasing amounts. The errors in the recovered normal maps and the albedo for different algorithms are shown in Figure 16. We observe that the most of the algorithms are resilient to the Poisson image noise when calculating the normal maps. However, the performance of the least square methods gets worse than others as the amount of the noise increases.




Figure 15: Sample images with increasing *Poisson image noise*. (See the online version for better quality.)




Figure 16: First column is the angular error, and the second column is the error in albedo as the *Poisson image noise* level increases.

6.6. Light-ring position

We also run experiments to determine the optimal position for the common configuration of a ring of lights surrounding the camera. In Figure 17, the light directions are shown on a hemisphere. We used 9 light-ring positions, each with a fixed number of lights, 24. Figure 18 shows the total average errors versus the light-ring configuration. 1 is the light configuration closest to the camera (which is in the center of the hemisphere at the top), while 9 is the closest configuration to the base ground. The conclusion of this experiment is that an intermediate light configuration results in the lowest errors: lights too near the camera result in not enough variation in light positions for a stable estimation, while too large a ring results in significant shadowing.




Figure 17: Shows all the **light configurations**, where in each experiment only one ring is on and each ring contains 24 lights.

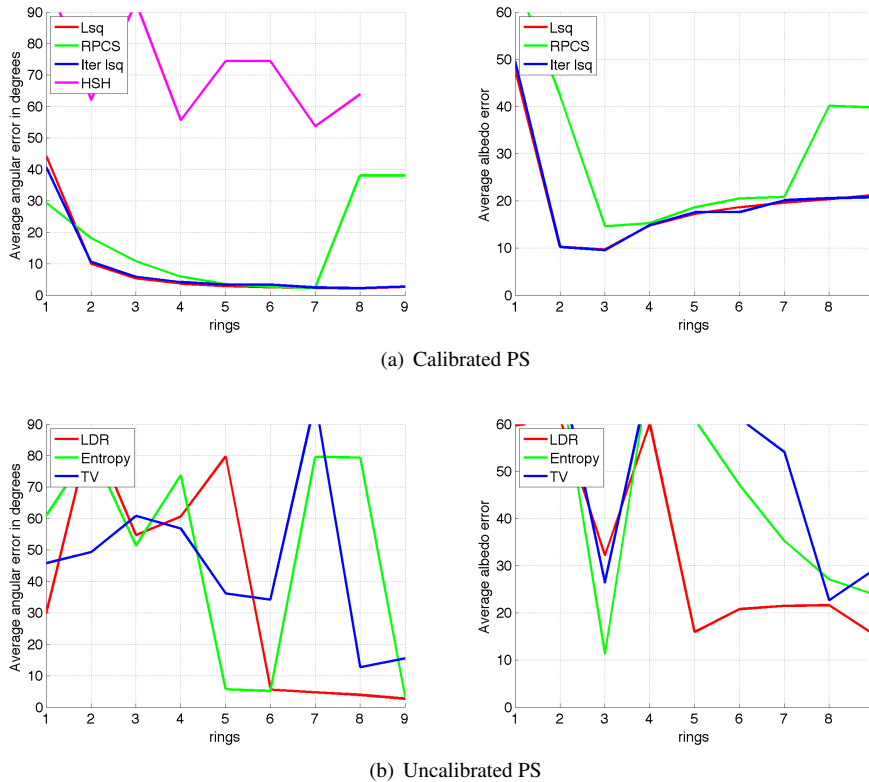


Figure 18: Average errors in normals (the first column) and albedo (the second column) vs different **light-ring configurations**.

6.7. Roughness of the Surface

Figure 19 depicts some 3d pear models with increasing surface roughness. As the surface roughness increases, the size of the highlight increases, making elimination of specularity harder. However, if the roughness increases enough, the surface begins to look diffuse. We can observe this phenomenon in Figure 20: errors first increase and then decrease. Nevertheless, the heuristics used in the iterative least squares method appear to work better than others.

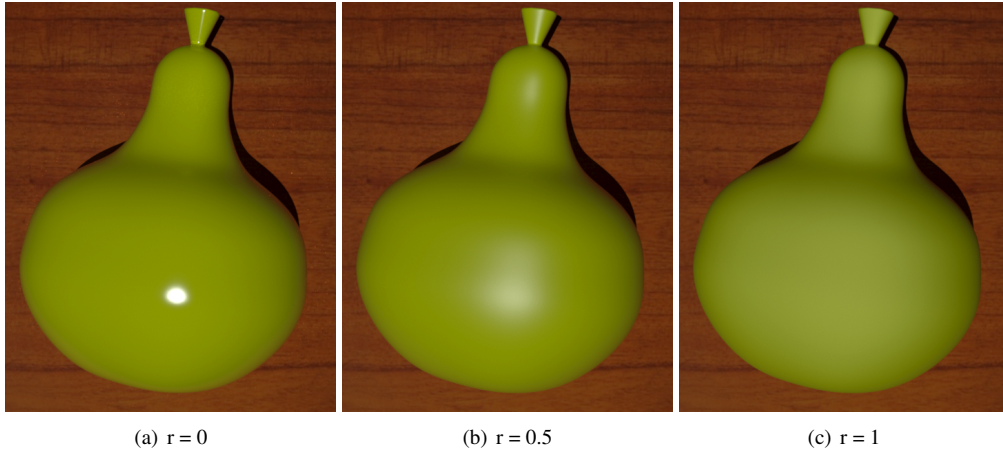


Figure 19: Pears with varying roughnesses.

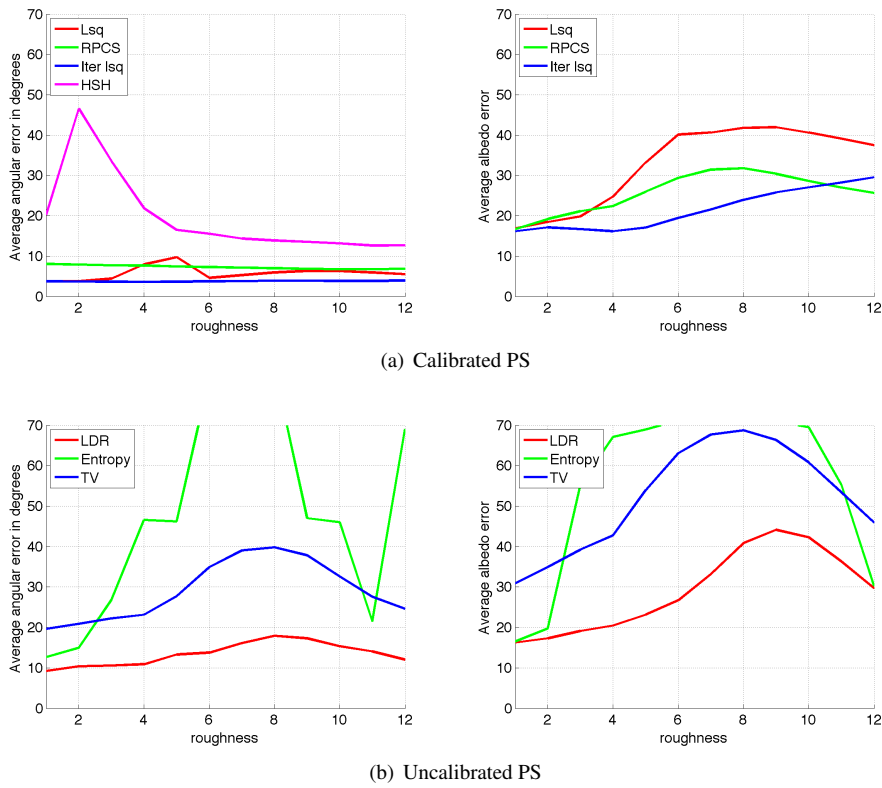


Figure 20: First column is the angular error, and the second column is the error in albedo as the surface roughness increases.

6.8. Shadow

To evaluate the effect of the shadow amount in the scene, we create sinusoidal surfaces, S , with varying heights as follows:

$$S = (x, y, f(x, y)) \tag{1}$$

$$f(x, y) = k * \sin(x) * \sin(y) \tag{2}$$

where the coefficient k determines the maximum height of the surface (and hence the amount of shadowing), as shown in Figure 21.

In Figure 22, the error plots are shown. The conclusion we draw is that all calibrated methods suffer from heavy shadows, however surprisingly using a constant hard threshold on shadow performs best among all the methods. We also observe that the uncalibrated methods behave unstably when they are solving the convexity for the object.




Figure 21: Sinusoidal surfaces with varying heights to explore the effect of the shadows.

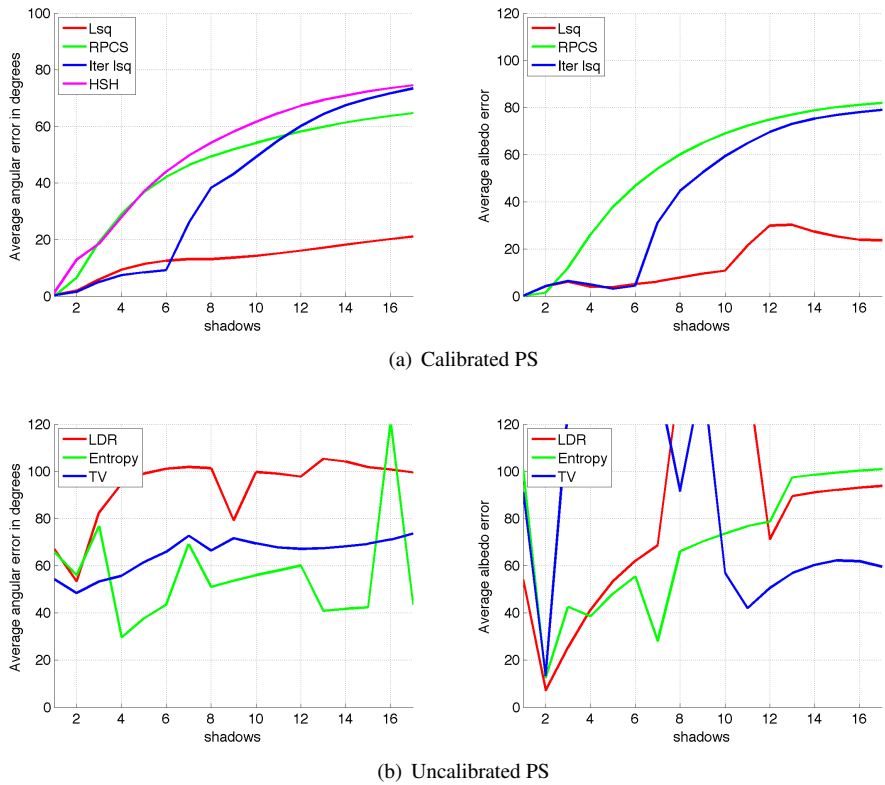


Figure 22: First column is the angular error, and the second column is the error in albedo as the amount of shadows increases.

6.9. Specularity

By only changing the specularity parameter of the *rough plastic* material in the Mitsuba Renderer, we create several datasets of the same pear as shown in Figure 23. The evaluations of the algorithms in Figure 24 reveal that increasing specularity does not change the average error per algorithm much. However, some algorithms perform better than others on this particular material, such as iterative least squares and LDR.




Figure 23: Pears with increasing specularity.

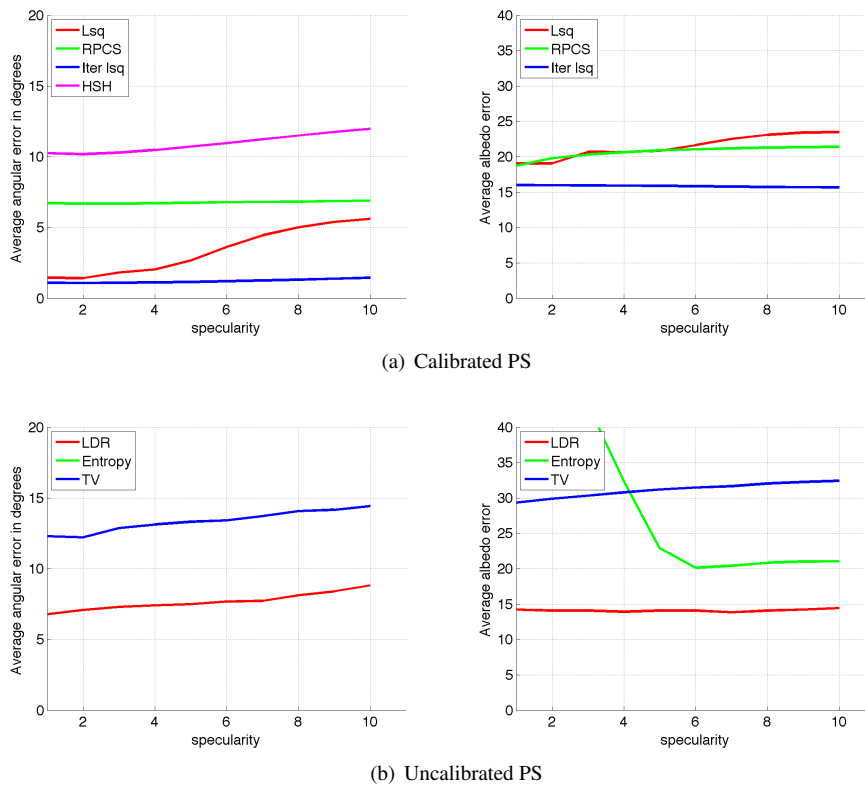


Figure 24: The first column is the angular error and the second column is the error in albedo as the specularity increases.

References

- [1] C. Toler-Franklin, A. Finkelstein, and S. Rusinkiewicz. Illustration of complex real-world objects using images with normals. In *NPAR '07: Proceedings of the 5th international symposium on Non-photorealistic animation and rendering*, pages 111–119, New York, New York, USA, Aug. 2007. ACM Request Permissions. 3