

Discovering Point Lights with Intensity Distance Fields

Supplementary Material

Edward Zhang
University of Washington
Seattle, WA
edzhang@cs.washington.edu

Michael F. Cohen
Facebook, Inc.
Seattle, WA
mcohen1@fb.com

Brian Curless
University of Washington
Seattle, WA
curless@cs.washington.edu

1. Specular Materials

In our derivation, we defined the shading function to be

$$S(q, p) \equiv \frac{V(x, p) f(x, \frac{p-x}{\|p-x\|}, \omega) (n(x) \cdot (p-x))_+}{\|p-x\|^3} \quad (1)$$

and the Intensity function to be

$$I(q, p) \equiv \frac{B(q)}{S(q, p)}. \quad (2)$$

In Figure 1a, we show an example of $I(q, p)$ for a Cook-Torrance microfacet BRDF [1]. In comparison to the diffuse $I(q, p)$ in Figure 1d, the darker area of the lobe is oriented such that it will more strongly limit the potential intensity of lights in the mirror direction from observation q .

To more intuitively reason about how our method interacts with specularity, it is helpful to assume a minimum intensity for point lights illuminating a scene. This is reasonable because we usually expect the number of point emitters in a scene to be fairly small; for a fixed amount of total power in a scene, fewer lights imply that each individual light should have higher power. Therefore, lights are unlikely to lie in areas with low IDF values. This implies that the observations limiting these areas are pruning away these regions from consideration.

Now, we consider specularity. Traditional specularity triangulation methods interpret a single bright specular highlight as a strong indicator of a light source in the mirror direction. However, with our formulation, a single bright observation is not assumed to be due to specularity – it could just indicate a nearby point light. Instead, our formulation effectively uses the contrapositive of this reasoning: a *lack* of specular highlight means that it is *unlikely* for a light to fall along the mirror direction.




Figure 1: Here we compare our scene transforms by comparing two scenes with the same light locations, but different materials. The top row corresponds to a scene where the left wall is specular and the remaining walls are diffuse; the bottom row corresponds to the scene where all walls are diffuse. In 1a, we show $I(q, p)$ for a microfacet BRDF, where observation q is the red point viewed from the direction of the arrow. For comparison, a diffuse $I(q, p)$ is shown in 1d. In 1b, we show the full IDF for the specular scene, where all the points on the left wall are viewed from the red point. Due to the specular observations, the center of the scene is darker than the same region in a diffuse scene with the same light positions in 1e. The Limiter field behaves similarly at the true light positions in both cases (1c and 1f).

This implies that additional observations of the same point greatly help our method. For example, if we viewed the point in the specular highlight from a different angle for which it was not in the specular highlight, that would provide evidence against the presence of a nearby point light (because the observation would limit the brightness of nearby points), but leave open the possibility of a light in the original mirror direction. Viewing non-specular highlights thus allows our

method to prune out more possible light positions.

We show an example of a full IDF for a scene with one specular wall (the left wall) in Figure 1b, where each point on the specular wall is viewed just once. We note that $\mathcal{L}(p)$, and by extension, $\mathcal{V}(p)$, qualitatively functions the same way as in the diffuse case; however, in specular scenes, we prefer having a denser set of observations in order to make the maxima clearer.

2. Occlusion

Occlusions do not significantly impact the performance of $\mathcal{V}(p)$ as a candidate position proposal scheme (Figure 2), since the voting function does not explicitly identify occlusion boundaries. $\mathcal{D}(p)$ is often discontinuous along shadow boundaries, and sometimes can directly result in light locations at the intersections of the discontinuities (Figure 2a). This is equivalent to tracing rays from shadow edges to the associated occluders and finding the intersection.




Figure 2: In an extreme case, occluders entirely block some surfaces from being lit (2a). More commonly, occlusion boundaries result in discontinuities in the IDF (2b). The shadow edges (white dashed lines) are visible in the Limiter field (2c) as the limiters of larger uniformly-colored regions. The inset shows a region in $\mathcal{L}(p)$ where shadow edges can result in false peaks in the $V(p)$. These can be removed (2d) by determining whether or not occlusion was responsible for the local variation in the limiter field.




Figure 3: Consider point p and $\mathcal{L}(p)$ on the scene surface, with the line from p to $\mathcal{L}(p)$ tangent to some occluder. Note that $B(\mathcal{L}(p))$ decreases to the left, in the opposite direction from the normal to the occluder at the tangent. If we examine another point p' to the left of p , without the occluder we would also expect $\mathcal{L}(p')$ to be left of $\mathcal{L}(p)$ (since p' is the same distance from $\mathcal{L}(p')$ as p is from $\mathcal{L}(p)$, and $B(\mathcal{L}(p')) < B(\mathcal{L}(p))$). However, due to the occluder, instead $\mathcal{L}(p')$ is the first unoccluded point to the right of $\mathcal{L}(p)$. As p' moves to the left, $\mathcal{L}(p')$ continues being forced to the right due to the occluder, and if the occluder is far from $\mathcal{L}(p)$ the amount by which $\mathcal{L}(p')$ moves can become very large. This gets detected as a peak in $\mathcal{V}(p)$.

For a point near an occluder, a small change in position can result in a large change in the surface regions visible to that point. If the limiter happens to lie in these disoccluded regions, then we will get false peaks in $\mathcal{V}(p)$ (Figure 2c-2d). These false peaks are caused when two conditions are met (as illustrated in Figure 3):

- the line from $\mathcal{L}(p)$ to p is tangent to an occluding surface, and
- $B(\mathcal{L}(p))$ increases in the direction of the normal of the occluding surface.

These conditions can be detected when computing $\mathcal{L}(p)$; however, in practice, we find that discarding candidate positions that are too close to scene surfaces is enough to filter out these false peaks.

The main practical issue with working with occlusions is that each $I(q, p)$ computation requires an occlusion test against the entire scene. Our 2D test sets had few enough primitives to allow such full scene intersection tests to be feasible. In the 3D case, we create a depth environment map for each p so that testing for occlusion between x and p is simply a depth lookup. In practice, we found that the surfaces used to compute $\mathcal{V}(p)$ in our real-world 3D data were effectively unoccluded from the entire scene, so to save computation time we did not factor in occlusion in our results.

3. Real-world 3D Datasets

We scanned two real world scenes to demonstrate that our method works even when some of the assumptions are violated. In both scenes, we set up two bare

light bulbs (approximating isotropic point sources) and measured their positions by hand.

The process for each scene is as follows:

1. First, we scan the scene using the Lenovo Phab2 Pro Tango phone, reconstructing the scene geometry using the Tango software.
2. We then perform exposure correction on the RGB images, following Zhang et al. [2], project the images onto the geometry, and take the robust minimum (average of the values between the first quartile and the median) intensity as the $B(q)$ (with a diffuse assumption).
3. We manually segment out some surfaces in the scan that are believed to have the same diffuse reflectance (generally the walls).
4. Using only $B(q)$ at these surfaces, we run our iterative algorithm, computing $\mathcal{V}(p)$ and extracting candidate positions.
5. During candidate position extraction, we filter out unlikely light positions, namely those that lie very close to scene surfaces. We also filter out light positions p where $\mathcal{D}(p)$ is smaller than some minimum light intensity threshold, following the reasoning in Section 1.
6. During the optimization portion of our algorithm, we solve for light intensities, positions, as well as an unknown constant ambient light intensity.
7. To evaluate the positional accuracy of our system, we compare the optimized positions with the hand-measured ground truth positions.
8. To evaluate the accuracy of our estimated intensities, we compared the relative intensities of the two lights by fixing one of the intensities to 1 (since the captured images could not provide absolute intensities). We obtained ground truth relative light intensities by taking an image of a diffuse scene lit only with one light, then replacing that light with the second light in the exact same position. Due to linearity of light, we can take the average ratio between the two images to be the relative light intensity.

The first scene, shown in the main paper and reproduced in Figure 4, is a small, relatively bare room. One light gives a clear maximum in $\mathcal{V}(p)$ (Figure 4c), and after a few iterations, another maximum appears (Figure 4d), at which point our optimization succeeds. No other maxima are detected. The error in the estimated positions, relative to hand-measured ground




Figure 4: Results of running our method on a scanned 3D dataset. In 4a-4b, we show the scanned scene colored with the averaged observations $B(q)$. The recovered light locations are rendered as red spheres; compared with the ground truth, the position error is 4% of the scene radius. In the bottom two images, we only show the $B(q)$ used in our computation. In 4c, the initial medial axis and $\mathcal{V}(p)$ are shown, with a clear maximum (circled in green) at the location of one of the light sources. After a few iterations (4d), the second light position is also revealed.

truth, were 0.05m and 0.08m in a 2m by 2.6m by 1.6m room, for a relative error of about 5%. The estimated intensity for the second light was 1.43, while the measured relative intensity was 1.39.

The second scene (Figure 5) is a harder case: the room was larger, had more complex geometry (a vaulted ceiling), included clutter, and was incomplete and not closed. One light gives a clear maximum in $\mathcal{V}(p)$ (Figure 5c); after several iterations, we see several new maxima appear. Most of them are filtered out using our heuristics, and we are left with four candidates, including the original maximum (Figure 5d). After optimization, the maxima that did not correspond to true lights had very small intensities compared to those that did represent true lights (less than 5% of the intensity). To get our final solution, we optimize again but remove the lights with low intensities. The error in the estimated positions, relative to hand-measured ground truth, were 0.26m and 0.31m in a 3.5m by 3.4m by 2.6m room, for a relative error of about 10%. The estimated intensity for the second light was 1.42, while the measured relative intensity was 1.39.




Figure 5: Results of running our method on a scanned 3D dataset. In 5a-5b, we show the scanned scene colored with the averaged observations $B(q)$. The original light locations are rendered as green spheres, while our recovered light locations are rendered as red spheres. In the bottom two images, we only show the $B(q)$ used in our computation. In 5c, the initial medial axis and $V(p)$ are shown, with a clear maximum (circled in green) at the location of one of the light sources. After a few iterations (5d), several more maxima appear (rendered as red spheres); two among these are close enough to the true light positions for the optimization to succeed.

References

- [1] R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. *ACM Trans. Graph.*, 1(1), 1982. 1
- [2] E. Zhang, M. F. Cohen, and B. Curless. Emptying, re-furnishing, and relighting indoor spaces. In *SIGGRAPH Asia*, 2016. 3