Eurographics Symposium on Rendering (2005)
Kavita Bala, Philip Dutré (Editors)

Table-top Computed Lighting for Practical Digital Photography

Ankit Mohan[1†], Jack Tumblin[1], Bobby Bodenheimer[2], Cindy Grimm[3], Reynold Bailey[3]

[1]Northwestern University, [2]Vanderbilt University, [3]Washington University in St. Louis

Abstract

We apply simplified image-based lighting methods to reduce the equipment, cost, time, and specialized skills required for high-quality photographic lighting of desktop-sized static objects such as museum artifacts. We place the object and a computer-steered moving-head spotlight inside a simple foam-core enclosure, and use a camera to quickly record low-resolution photos as the light scans the box interior. Optimization guided by interactive user sketching selects a small set of frames whose weighted sum best matches the target image. The system then repeats the lighting used in each of these frames, and constructs a high resolution result from re-photographed basis images. Unlike previous image-based relighting efforts, our method requires only one light source, yet can achieve high resolution light positioning to avoid multiple sharp shadows. A reduced version uses only a hand-held light, and may be suitable for battery-powered, field photography equipment that fits in a backpack.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture I.3.3 [Computer Graphics]: Picture/Image Generation

1. Introduction

Modern digital cameras have made picture-taking much easier and more interactive. However, lighting a scene for good photography is still difficult, and practical methods to achieve good lighting have scarcely changed at all. We show that sketch-guided optimization and simplified forms of image-based lighting can substantially reduce the cost, equipment, skill, and patience required for small-scale studio-quality lighting.

Good studio lighting is difficult because it is a 4D inverse problem that photographers must solve by making successive approximations guided by years of experience. For non-experts, good studio lighting can be surprisingly frustrating. Most people can specify the lighting they want in screen space (e.g., “get rid of this obscuring highlight, make some shadows to reveal rough texture here, but fill in the shadows there”), but determining what kind of lights to use, where to place them, and how to orient them is never easy.

We are especially interested in camera-assisted lighting for human-scale, desktop-sized static objects. We want lighting that accurately reveals the shape, texture, materials, and most visually meaningful features of the photographed item. In particular, we seek a method to help museum curators as they gather digital photographic archives of their vast collections of items.

Pioneering work in image-based lighting [DHT*00, HCD01, DWT*02, MPDW03] offers promising approaches that can help with the photographic lighting problem. Unfortunately, most require too many precise measurements and adjustments for day-to-day use outside the laboratory. Precision is required to address more ambitious goals such as recovering shape, BRDF, and appearance under arbitrary viewing and lighting conditions. For the much smaller, yet more widespread problem of photographic lighting, we need a method that requires less time, expense, and complexity, yet allows users who are not lighting experts to quickly find the lighting they want.

This paper offers three contributions. We extend existing

† ankit@cs.northwestern.edu

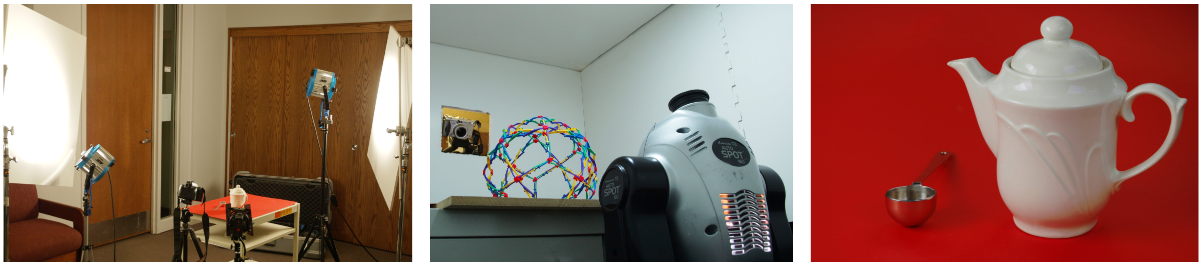

Figure 1: *Light placement for obtaining high quality photographs can be extremely tedious and time consuming (left). Our system use a simple setup with a steerable spotlight and an uncalibrated enclosure (center) to obtain results comparable to professional lighting even when used by novice users (right).*

image-based lighting ideas to reduce the required equipment to a single light source and single camera; we replace trial-and-error light repositioning with optimization and on-screen painting; and we reduce the need for high dynamic range photography, thus reducing the capture time. The result is a novel and inexpensive system that a novice can use to intuitively describe and obtain the desired lighting for a photograph.

2. Related Work

Lighting has long been recognized as a hard problem in computer graphics and many papers have explored optimization for light placement and other parameters [SDS*93, KPC93, PF95, CSF99, SL01]. Some of these systems used painting interfaces to specify desired lighting in a 3D scene [SDS*93, PF95, PRJ97], and we use a similar approach to make lighting for photography more intuitive. The system by Shacked et al. [SL01] was even able to provide fully automatic lighting by applying image quality metrics. Marschner et al. [MG97] used inverse rendering techniques to estimate and alter the directional distribution of incident light in a photograph. However, all these systems require 3D information unavailable in our photographic application.

Several commercial photographic products have also used lighting enclosures similar to ours, but they achieve very soft lighting with limited user controls. Moreover, they do not help users solve light placement problems. These systems include diffusive tents [Pho], photo-boxes [MK] and translucent back-lit platforms with an array of individually dimmed light sources [Ast].

Image-based methods have also been used to permit arbitrary relighting of well-measured objects. Most methods, including ours, perform relighting using a weighted sum of differently lit basis images, done first by [NSD94]. However, prior efforts used more elaborate and expensive equipment because their goals were different from ours. These include measurement of a 4D slice of the reflectance field of the human face [DHT*00], museum artifacts measured by a rotating-arm light stage [HCD01], an ingenious but extensive system by Debevec et al. [DWT*02] for real-time video playback and measurement of light fields, a dome of electronic flashes for real time image relighting [MGW01], a free form light stage to enable portable gathering of light-field data with some calibration [MDA02], and full 4D incident light measurements by Masselus et al. [MPDW03]. In all of these cases, data-gathering required either customized equipment or collection times much longer than would be practical for photographic lighting.

Three recent systems also offered novel sketch guided relighting from basis images. Akers et al. [ALK*03] used a robotic light-positioning gantry to gather precisely lit images, and like us, provided a painting interface to guide relighting. But unlike us they used spatially varrying weights that could produce physically impossible lighting. Digital Photomontage [ADA*04] used sketch guided graph-cut segmentation coupled with gradient domain fusion to seamlessly merge several photographs. They demonstrated merging differenlty lit photographs to create novel illumination conditions. Though their interaction scheme worked well for a small number of images ($\sim$10), it may be impractical for the hundreds of images required for complete control over lighting directions. Also, their system does nothing to help the user with light placement, and may produce physically unrealizable results. Anrys and Dutre [AD04] used a Debevec-style light stage with around 40 fixed, low powered light sources and a painting interface to guide lighting. Their optimization only found light intensities, and light placement was still left up to the user. Also, their point light sources could cause multiple shadows and highlights which might be undesirable for archival purposes. The data capture time was high since they captured high-dynamic-range (HDR) photos for every light location.

Unlike previous attempts, our system does not require users to decide on correct or complete light source placement. This is possible because our capture process is significantly dif-

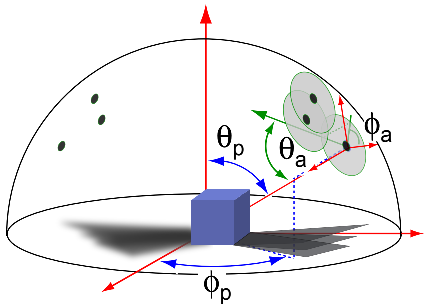

Figure 2: *All possible lighting angles parameterized by light position* (θ_p, ϕ_p) *and direction* (θ_a, ϕ_a). *Point light sources (on the left side of the hemisphere) result in multiple hard shadows, while overlapping area (on the right) light sources can be used to simulate a larger light source.*

ferent, and better suited for the task of photography. We require less than five minutes to complete the initial image capture and a few more minutes to get the final result. The equipment required is minimal and portable, and our hand-held version can be carried in a backpack. Also, HDR capture is reduced to a minimum in our system.

3. Simplifications: HDR and 2D lighting

Our goal is to do what a good photographer does, but with computational help. We want to light a scene for a particular photograph, *not* build a calibrated 4D data set to reconstruct every possible form of illumination. Photographers make consistent choices about which types of lights to use, how to adjust them, and where to place them. We will show how our streamlined image-based method follows these same choices.

Like most previous image-based lighting methods, we apply the observations formalized by Nimeroff [NSD94] that lights and materials interact linearly. If a fixed camera makes an image I_i from a fixed scene lit only by a light L_i, then the same scene lit by many lights scaled by weights w_i will make an image $I_{out} = \sum_i w_i I_i$. Adjusting weights lets us "relight" the image, as if the weights modulate the lights rather than the images. As we collect more images I_i, we can simulate more lighting possibilities.

How many images do we really need to gather? We only need enough images to span the kind of lighting a skilled photographer might explore to get good results in a photo studio. Several common practices in studio lighting can help us.

First, professional photographers choose lamps with broad, nearly uniform beams of light, often with a reflector and lens to help direct more light forward. Second, they adjust light placement angles carefully, but not their distances from the object. Distance to the light affects foreshortening of shadow shapes, but these effects are subtle and rarely noticed. Third, they adjust lights to control shadow softness versus sharpness. Light sources (or more accurately, the shadows they form) become 'softer' by increasing the angular extent as measured from the lit object. Fourth, they seek out lighting arrangements that produce a simple set of shadows and highlights that best reveal the object's shape, position, and surface qualities. They avoid complex overlapping shadows, lack of shadows due to overly-soft light, and contrast extremes due to large specular highlights or very dark shadows. Simpler shadows usually mean fewer lights, and thus fewer basis images.

Accordingly, we use commercially available light sources instead of custom or special-purpose devices. We place light sources at a moderate distance (typically around 1 meter) from the object. We use small-to-moderate area 'soft' light sources instead of the much sharper point-like sources often used in earlier approaches. Overlapped soft shadows blend far less noticeably than sharp shadows from the same light positions (as shown in Figure 2), thus requiring fewer images to avoid multiple shadow artifacts. Also, overlapping area light sources can be combined to produce a larger area light source.

Note that we do not need to know the light positions or their absolute intensities for our images; we select weights w_i and images I_i by their ability to match the lighting target images a user sketches for us. Instead of calibration, we only need consistency in the aiming direction of a single, commercially available steerable light, and consistency in the light response curve of a commercially available digital camera.

We also avoid the use of HDR photographs where possible, as these typically require multiple calibrated exposures and computation to merge them [DM97]. Instead, we rely on the camera's automatic exposure adjustments to capture what we call *light-aiming images* suitable for interactive lighting design. We photograph high resolution basis images afterwards, for construction of the output image, and only resort to HDR capture methods for a basis image with large over-exposed regions. Under-exposed regions can be ignored, as their contributions are already invisible, and are further reduced as their weights are less than one ($w_i \leq 1$).

Formally, arbitrary external illumination is four-dimensional for a desktop scene: $L(\theta_p, \phi_p, \theta_a, \phi_a) = L(\Theta)$. Suppose that the photographed object receives all its light from a hemisphere of tiny, invisible, inward-pointing video projectors, each at a distance r from the object. Each projector's position in desktop polar coordinates is (θ_p, ϕ_p). Each projector's centermost pixel $P(\theta_a = 0, \phi_a = 0)$ forms a ray that illuminates the center point of our desktop, and in the projector's polar coordinates the other pixels are $P(\theta_a, \phi_a)$, as shown in Figure 2. All projectors' light output is the 4-D incident light field, and describes all possible lighting. To simulate all possible lighting, we would need a new image I_i to capture light from each pixel of each video projector! Instead, we use only broad beams of light ($P(\theta_a, \phi_a) \cong \cos(\theta_a)cos(\phi_a)$), regular sampling of light placement angles (θ_p, ϕ_p), and specify 'softer' to 'sharper' shadows by varying the angular extent (θ_p, ϕ_p) as measured from the lit object. This angular extent should not be confused with the lamp's beam width (θ_a, ϕ_a); in our 'hemisphere of video projectors' analogy, beam width sets the image from a projector, but angular extent sets the number of adjacent projectors that emit this same image.

In summary, rather than recreate arbitrary 4D incident light fields, we use weighted sums of basis images that represent the type of lighting used by professional photographers. This method is much more practical and efficient, with little, if any, loss of useful generality.

4. Method

We construct a high quality user-guided picture in three steps. First the system automatically captures low-resolution light-aiming photos for densely sampled lighting angles around the photographed object. These quick photos are used only to guide the lighting design, not to form the final output. Second, the user iteratively paints the desired lighting by simple lighten-darken operations to generate a target image. The system finds weights w_i for each light-aiming photo such that their weighted sum matches the target image in the least-squares sense. Finally the system takes a few selected high resolution basis images by relighting the scene from light source positions that have weights w_i greater than a threshold. A weighted sum of these high resolution images gives the final result. If the result is not satisfactory, the user can sketch on the current result for use as the next iteration's target image.

4.1. Enclosed Light Source & Aiming Images

Freed from photometric and angular calibration requirements as discussed in Section 3, we are able to build a much simpler and cost-effective controlled light source. We place the object and a gimbal-mounted moving-head spotlight inside an enclosure of almost any convenient size, shape and material. The powerful computer-aimed light pivots to any desired pan and tilt angle with good repeatability ($\leq \pm 0.5°$) to light any desired spot inside our enclosure. The enclosure acts as a reflector, and effectively provides a controllable 2D area light source around the object. The size and shape of the enclosure is almost irrelevant as long as the light is close enough to the object to keep parallax low, and the light is powerful enough for the camera to get a reasonable exposure.

We built a $1 \times 1 \times 1.5m^3$ sized box of white $1/2$" foam-core board as our enclosure, and chose an inexpensive moving-head spotlight. The 150-watt *American DJ Auto Spot 150* disco-light, shown in Figure 1 can tilt 270°, pan 540°, and includes 9 color filters, gobos and several other fun features.

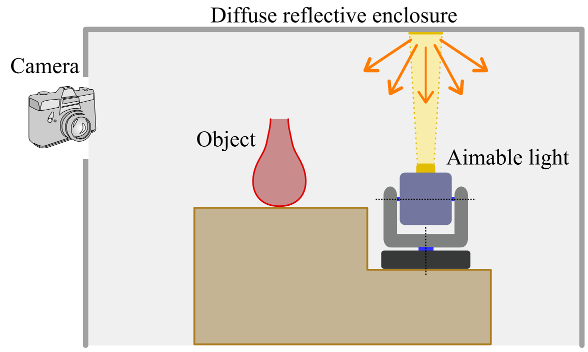

Figure 3: *The disco-light setup. The object and disco light are both enclosed in a white foam box, with the camera looking in through a window in the enclosure wall farthest from the light.*

Computer control by the DMX512 protocol is easy to program with the SoundLight USB DMX controller. Our foam-core enclosure resembles a hemi-cube around a pair of tables. We place the gimbal light on a small table that lowers its rotation center to the plane of an adjacent taller table holding the photographed object, as shown in Figure 3. Using adjacent but separate tables reduces vibration, permits gimbal angles to approximate hemisphere angles, and separates the object from the swiveling lamp. We place the camera behind a small opening cut in the enclosure wall on the end farthest from the light source.

The system gathers aiming images rapidly and automatically. Through the DMX512 controller we direct the gimbal light to scan the upper hemisphere of light aiming directions in equal-angle increments as we record low-resolution aiming images, either by collecting viewfinder video (320×240@10Hz) or by individual computer-triggered photographs using auto-exposure. We are able to record hundreds of individual aiming images per minute, and can complete all the data gathering in less than five minutes using a Pentium 2GHz computer, and a Canon Powershot G3 camera.

To the best of our knowledge, no other image-based lighting work exploits these movable and controllable lights. Enclosed pivoting lights retain many advantages of the more sophisticated lighting systems, avoid multiple sharp shadows, can offer variable 'softness' by spot size adjustment, and are much simpler and cheaper to construct. Of course, they do not easily provide accurate lighting direction calibration or point-light illumination, but these features are not needed for our goals.

After recording, we linearize each captured frame (RGB) by applying the camera's inverse response curve, recovered by the method of Debevec et al. [DM97], and converted to luminance values. Linear response ensures weighted sums of whole images are accurate representations of physically realizeable lighting. We then down-sample the linearized aiming image dataset to 64×64 for use as the *aiming basis set* for the following optimization step.

4.2. Sketch-Guided Lighting Optimization

After gathering aiming images, users can interactively specify and refine lighting by sketching the desired intensity on a *target intensity image*. This grey-scale image (examples in Figure 5) approximates the final output image the user would like to see. For editing the target image, the user starts off either with a simple grey wash (such as uniform grey, or light grey fading to dark grey across the image, etc.), or the previous iteration's result. The user then carries out a series of lighten and darken operations in the different regions of the image to approximate the desired results. The process is extremely simple and intuitive, and takes a few of minutes at most.

Given a target image, the optimization finds weights w_i for each aiming image that produces the best match to the target image. We take a constrained least-squares approach, solving for weights w_i for each of the small, luminance-only aiming basis images. Let N be the number of images in the aiming image set, each of size $m \times n$. We formulate the optimization problem as follows:

$$\min_{w} \; |Aw - t|^2$$

$$\text{subject to} \quad 0 \leq w_i \leq 1 \quad \forall\, i \in (1 \ldots N)$$

where w is the N-dimensional vector of weights, A is an $(m \times n) \times N$ matrix of basis images (that is, each basis image is treated as a vector), t is the $(m \times n)$ vector representing the target image painted by the user, and $|.|$ represents the L^2 norm of the vector. We solve this bound constrained quadratic optimization problem using an active set method [NW99]. The optimization is quite fast and takes around 1-2 minutes on a 2GHz Pentium 4 desktop machine.

The result is a least-squares optimal match to the supplied target image. As the objective function is quadratic, weights for images with weak contributions are rapidly driven to zero. In our experience, the number of significant nonzero weights is consistently small $(5 - 15)$. This greatly reduces the number of images needed for the final lighting solution.

After finding the w_i weights, we apply them to the linearized color aiming images, then re-apply the camera response function to display a preview of the output image. The user then has the option of replacing the target with a grayscale version of this result and can repeat the sketching and optimization cycle until satisfied with the color preview of the output image.

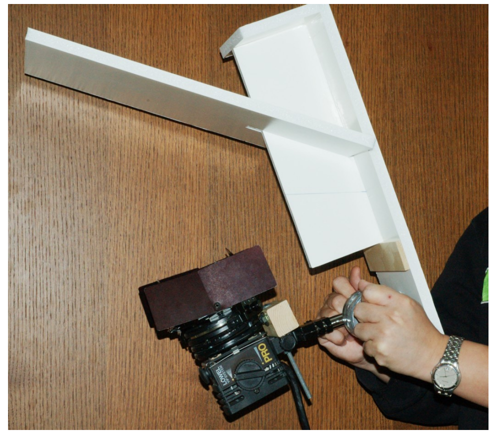

Figure 4: *Light source with attached foam-core diffuse reflector used for hand-held data gathering.*

4.3. Output Assembly

The user now has the desired visually pleasing, but low-resolution, image that is a weighted sum of a small subset of the linearized aiming images. For high-quality results, we wish to replace each of these aiming images with an image taken at the maximum resolution available from the camera. We re-take *just* those photos that correspond to the aiming images with significant weights w_i, again using auto-exposure on the camera, and record a set of high-resolution photos called *basis images*. Recall that we can exactly replicate the lighting using the gimballed spotlight; the only things that change are the camera settings.

We capture HDR photographs for images that contain large over-exposed regions as a result of the camera's autoexposure. As discussed in Section 3, under-exposed regions do not require HDR photos. We then linearize each basis image to remove effects of the camera response curve. As before, we construct a linear output image as a weighted sum of basis images, using the weights determined by the optimization to match the target image. Finally, we re-apply the camera's response function to the linear output image to get the desired high resolution result.

5. Portable, Hand-held Method

Even a foam-core box and a moving-head spotlight are impractical to carry around everywhere. However, the 'Free-form light-stage' [MDA02] showed that it is possible to gather calibrated image sets suitable for 2D relighting with nothing more than four small light-probe-like spheres, a digital camera on a tripod, a hand-held point-light source, possibly battery-powered, and approximately 30 minutes of time to take several hundred digital photographs. Pang et al. [PWH04] also used a similar approach by mounting a camera on the light source and used camera calibration techniques to estimate lighting directions with reasonable accuracy. While these methods try to meet the ambitious goal of incident light field capture, they would tax anyone's patience to record more than just a few items. We present a faster and simpler variant that serves our purposes better.

In the method of Section 4, we required repeatable light source positioning. However, if we record all of our 'aiming images' at the final output resolution, and if we either ignore over-exposed specular highlights or record high dynamic range images when needed, then *repeatability is not needed*. This allows us to use a hand-held light source instead. As shown in Figure 4, we use a small 250W hand-held light intended for television news cameras, attached to a diffuse reflector (foam core again), and limit the beam width with barn-doors to form a well-defined area light source.

To gather all photos, we hold the light outstretched and "dance" (see video). We sample the hemisphere of lighting directions by a polar-coordinate scan in ϕ-major order as the camera takes sequential photographs. A Nikon D70 camera, takes a steady stream of photos at about 3 frames per second using autoexposure for each frame. The user stands facing the object, and holds the light at arms' length while moving the lamp in an arc that passes directly over the object. The user moves the lamp from one side of the table to the other, scanning by π radians in θ axis with constant ϕ, and the natural alignment of their shoulders helps aim the light's centerline directly at the object. After each pass over the object with the light, the user steps sideways to change the ϕ angle for the next scan, and makes enough of these passes to cover $0 \leq \phi < \pi$ radians. In practice the user can be more careless with the light, as long as the hemisphere of light directions is well-sampled and the images are not over-exposed. After the image capture *dance* is complete, we downsample all images to construct aiming photos, and proceed with the sketch guided lighting design as before.

We find this process is quite simple and pleasing, and in under three minutes we can gather around 150 high-quality aiming/basis photos. An experienced user might not need to scan the whole hemisphere, but can quickly illuminate just from the useful and interesting lighting directions.

6. Results

Images in Figure 5 show results from our sketch guided lighting system. Both the moving-head light and the handheld methods are equally successful at creating arbitrary cleanly-lit images of desktop-sized objects. The data sets gathered by either method is sufficiently dense to allow easy lighting design. Additionally, our system yields reasonable results even when presented with unrealistic targets or highly reflective objects.

Figure 5(a), demonstrates a user interaction sequence with the system. Starting from a uniform grayscale image as the target, the user guides the optimization, iteratively improving the target until she gets the desired output. Figure 5(b) shows how simple approximate sketching on the target image can give an interesting sidelighting effect. Figure 5(c) shows how the highlight can bring out the underlying texture in a surface.

Figure 5(d) shows lighting for a highly specular object. Good lighting for such smooth, highly reflective objects is always difficult, as the light source itself is visible in the reflection. Our system produces results similar to the target image without large, objectionable saturated regions. In future systems we may hide the enclosure seams by constructing wide smooth rounded corners resembling a photographer's 'cyc'.

Figure 5(f) shows results from the handheld method of Section 5. The data gathering time was under 3 minutes, and the results are comparable to the moving-head light method. While the handheld method is not practical for photographing a large collection of objects, it can be an invaluable tool for well-lit photography in the field.

7. Discussion and Future Work

The ability to have large area light sources is crucial for photographing highly specular objects. Light source size also affects the sharpness of shadows and highlights. Our system has a unique advantage in that larger area light sources can be simulated by combining pictures illuminated with overlapping light sources. We could extend our optimization to penalize each distinct light source cluster, thus preventing disjoint highlights. The softness of the light can also be controlled by varying the beam width between a point-source and a large area source as it quickly sweeps over the hemisphere of lighting directions. More advanced moving-head spotlights usually provide controllable spot sizes suitable for this purpose.

Even though our system is aimed primarily at non-professional photographers, a few simple additions can make it a flexible tool for a creative expert to experiment with different lighting designs more easily. For example, the user might specify a simple weighting mask to set the importance of different image regions and influence the optimization process. While weighting masks would make the system more flexible, they would complicate the target sketching process. We do not know yet if the results would warrant the increase in complexity. Also, tools to directly tweak the light position and size on a virtual hemisphere around the object might also aid expert users.

There are several possible ways of dealing with the ambient light in the reflective enclosure. Underexposing all images using exposure compensation on the camera, using a larger enclosure or one made of materials with special reflective properties would greatly minimize the ambient component. Finally, it might also be possible to explicitly subtract the ambient term from the basis images.

This paper takes the problem of good lighting for desktop photography and finds a simple and practical solution using image-based relighting techniques. More sophisticated image-based measurements might also be achievable while maintaining the simplicity and elegance of the system. For example, we could estimate the incoming light direction by calibrating the ad-hoc enclosure setup with a light-probe, or by using dimensionality reduction [WMTG05] for the hand-held case. Combined with surface normals, such calibration might suffice for image-based estimates of BRDF.

8. Acknowledgements

We would like to thank the computer graphics group at Northwestern University for their support and feedback and the anonymous reviewers for their helpful comments and suggestions. Thanks to Jingjing Meng and Xing Hu for help in early stages of the project, Kelli Johnson and Nathan Matsuda for help with the construction of the enclosure, Vincent Masselus and Amy Gooch for help with Figure 2, and Holger Winnemöller for help with the video. This research was funded in part by NSF grants 0238062 and 0237621.

References

[AD04] ANRYS F., DUTRÉ P.: Image based lighting design. In *The 4th IASTED International Conference on Visualization, Imaging, and Image Processing* (2004). 2

[ADA*04] AGARWALA A., DONTCHEVA M., AGRAWALA M., DRUCKER S., COLBURN A., CURLESS B., SALESIN D., COHEN M.: Interactive digital photomontage. In *SIGGRAPH* (2004), vol. 23, pp. 294–302. 2

[ALK*03] AKERS D., LOSASSO F., KLINGNER J., AGRAWALA M., RICK J., HANRAHAN P.: Conveying shape and features with image-based relighting. In *IEEE Visualization* (2003). 2

[Ast] ASTRON SYSTEMS: Orbiculight lighting system. http://www.astronsys.com. 2

[CSF99] COSTA A. C., SOUSA A. A., FERREIRA F. N.: Lighting design: a goal based approach using optimisation. In *Proceedings of Eurographics Workshop on Rendering* (1999), vol. 10, pp. 317–328. 2

[DHT*00] DEBEVEC P., HAWKINS T., TCHOU C., DUIKER H.-P., SAROKIN W., SAGAR M.: Acquiring the reflectance field of a human face. In *SIGGRAPH* (2000), pp. 145–156. 1, 2

[DM97] DEBEVEC P. E., MALIK J.: Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH* (1997), vol. 31, pp. 369–378. 3, 4

[DWT*02] DEBEVEC P., WENGER A., TCHOU C., GARDNER A., WAESE J., HAWKINS T.: A lighting reproduction approach to live-action compositing. In *SIGGRAPH* (2002), pp. 547–556. 1, 2

[HCD01] HAWKINS T., COHEN J., DEBEVEC P.: A photometric approach to digitizing cultural artifacts. In *Proceedings of conference on Virtual reality, archeology, and cultural heritage* (2001), pp. 333–342. 1, 2

[KPC93] KAWAI J. K., PAINTER J. S., COHEN M. F.: Radioptimization: goal based rendering. In *SIGGRAPH* (1993), pp. 147–154. 2

[MDA02] MASSELUS V., DUTRÉ P., ANRYS F.: The free-form light stage. In *Proceedings of the 13th Eurographics workshop on Rendering* (2002), pp. 247–256. 2, 5

[MG97] MARSCHNER S. R., GREENBERG D. P.: Inverse lighting for photography. In *Proceedings of IS&T/SID Fifth Color Imaging Conference* (1997), pp. 262–265. 2

[MGW01] MALZBENDER T., GELB D., WOLTERS H.: Polynomial texture maps. In *SIGGRAPH* (2001), ACM Press, pp. 519–528. 2

[MK] MK DIGITAL: Photo light box. http://www.mkdigitaldirect.com. 2

[MPDW03] MASSELUS V., PEERS P., DUTRÉ P., WILLEMS Y. D.: Relighting with 4d incident light fields. In *SIGGRAPH* (2003), vol. 22, pp. 613–620. 1, 2

[NSD94] NIMEROFF J. S., SIMONCELLI E., DORSEY J.: Efficient re-rendering of naturally illuminated environments. In *Proceedings of the Fifth Eurographics Workshop on Rendering* (1994), pp. 359–373. 2, 3

[NW99] NOCEDAL J., WRIGHT S. J.: *Numerical Optimization*. Springer-Verlag New York, 1999. 5

[PF95] POULIN P., FOURNIER A.: Painting surface characteristics. In *Proceedings of Eurographics Workshop on Rendering* (1995), pp. 160–169. 2

[Pho] PHOTEK: Digital lighthouse. http://www.photekusa.com. 2

[PRJ97] POULIN P., RATIB K., JACQUES M.: Sketching shadows and highlights to position lights. In *Proceedings of Conference on Computer Graphics International* (1997), pp. 56–63. 2

[PWH04] PANG W.-M., WONG T.-T., HENG P.-A.: Estimating light vectors in real time. *IEEE Computer Graphics and Applications 24*, 3 (2004), 36–43. 5

[SDS*93] SCHOENEMAN C., DORSEY J., SMITS B., ARVO J., GREENBURG D.: Painting with light. In *SIGGRAPH* (1993), pp. 143–146. 2

[SL01] SHACKED R., LISCHINSKI D.: Automatic lighting design using a perceptual quality metric. *Computer Graphics Forum 20*, 3 (2001), 215–226. 2

[WMTG05] WINNEMÖELLER H., MOHAN A., TUMBLIN J., GOOCH B.: Light waving: Estimating light positions from photographs alone. *Computer Graphics Forum 24*, 3 (2005), to appear. 7

(a) Sequence showing successive sketching/optimization iterations to get the desired lighting. The first result uses a constant grayscale target, while the others use previous results as starting points for the target image.

(b) Strategic placement of highlights in the target result in an interesting side-lit image.

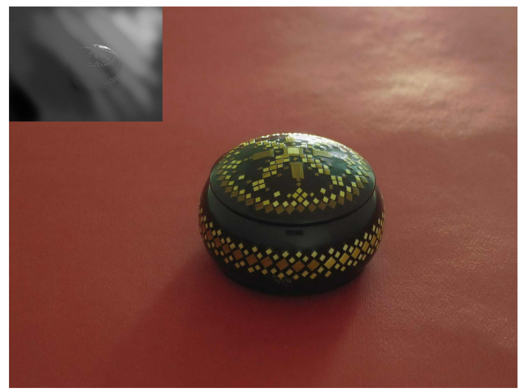

(c) Positioning of highlights reveals underlying texture in the surface.

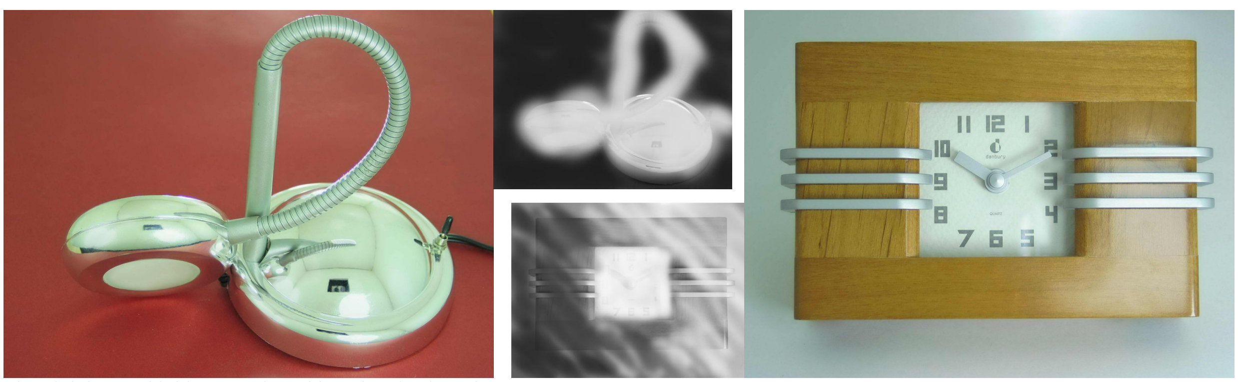

(d) Lighting a highly specular object by forcing the background to be dark.

(e) Target results in image suggesting illumination from the right.

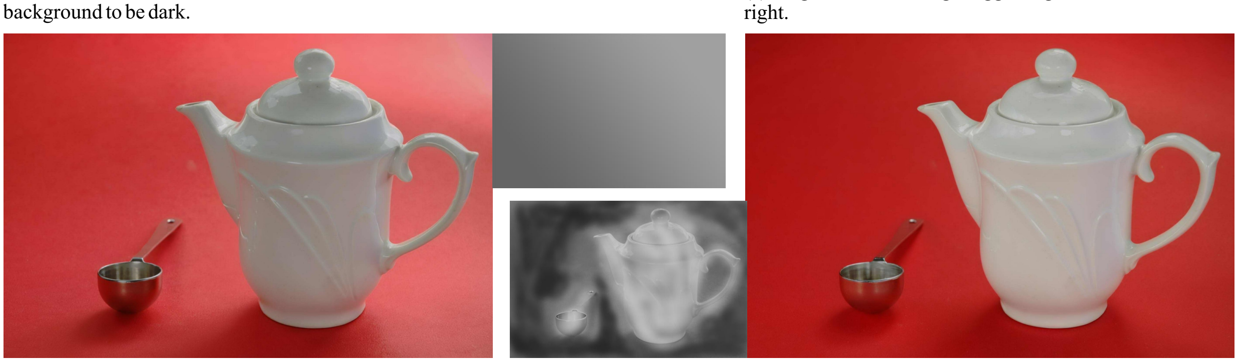

(f) Data captured by the handheld method. Image on the left uses a smooth grayscale gradient as the target image.

Figure 5: *Sample target images and lit photographs.*