

A Theory of Multiplexed Illumination

Yoav Y. Schechner

Dept. Electrical Engineering
Technion - Israel Inst. Technology
Haifa 32000, ISRAEL
yoav@ee.technion.ac.il

Shree K. Nayar and Peter N. Belhumeur

Dept. Computer Science
Columbia University
New York, NY 10027
{nayar,belhumeur}@cs.columbia.edu

Abstract

Imaging of objects under variable lighting directions is an important and frequent practice in computer vision and image-based rendering. We introduce an approach that significantly improves the quality of such images. Traditional methods for acquiring images under variable illumination directions use only a single light source per acquired image. In contrast, our approach is based on a multiplexing principle, in which multiple light sources illuminate the object simultaneously from different directions. Thus, the object irradiance is much higher. The acquired images are then computationally demultiplexed. The number of image acquisitions is the same as in the single-source method. The approach is useful for imaging dim object areas. We give the optimal code by which the illumination should be multiplexed to obtain the highest quality output. For n images corresponding to n light sources, the noise is reduced by $\sqrt{n}/2$ relative to the signal. This noise reduction translates to a faster acquisition time or an increase in density of illumination direction samples. It also enables one to use lighting with high directional resolution using practical setups, as we demonstrate in our experiments.

1 Introduction

Imaging objects under different source directions is important in computer vision and computer graphics [1, 2, 3, 5, 7, 8, 13, 14, 15, 18, 19, 20, 22, 24, 25, 26]. It is used for various purposes: object recognition and identification [2, 6, 8, 14, 21, 25, 26], image based rendering of objects and textures [3, 4, 5, 13, 14, 19, 20, 22], and shape recovery [10, 11, 12, 13].

In the above mentioned research directions and applications, images have been acquired under a single light source at a time. Frequently, however, a single source does not illuminate all the object parts with sufficient intensity to produce images with a high signal-to-noise ratio (SNR). While this problem may be overcome using long exposures, such an approach significantly lengthens the acquisition time.

In contrast to using single sources, we show that illuminating objects by multiple sources has significant benefits. Given a set of desired illumination directions, our approach enables capturing the required information with a much higher quality without adding to the acquisition

time. The approach reduces problems associated with dynamic range, e.g., due to shadows and specular highlights, although the exposure settings are the same in all the acquired images. We formalize these statements in a theory of *multiplexed illumination*. We describe the optimal scheme for multiplexing illumination sources from different directions during image acquisition, and the computational demultiplexing which follows the acquisition. We stress that the result of the demultiplexing is not an approximation; *all the features that one obtains with a single source (shadows, specularities, shading) are fully recovered*.

Beside giving a theoretical analysis of the benefits of the method, we also describe the limitations of the multiplexed illumination approach. Finally, we present a novel design for an easily programmable light source. A projector creates patterns on a white wall. The patterns reflecting off the wall towards the object serve as light sources. This fast and flexible lighting apparatus was used to demonstrate the multiplexing theory in our experiments.

We stress that this work is unrelated to structured light methods. While structured light deals with spatial patterns projected directly onto the object, we multiplex the *direction* from which the entire object is illuminated.

The proposed approach yields dramatically better results than those of the single-source illumination method. If n images are taken for imaging the object under n light sources, the method we devise improves the SNR by $\sqrt{n}/2$. It shortens the acquisition time by the same factor, relative to methods which enhance the SNR by long exposures. The results of this paper have implications for a broad range of vision and graphics algorithms

2 Standard Lighting Methods

Almost all current methods used for illumination research and gathering databases of objects under variable lighting directions are based essentially on single light sources. Such a setup is schematically depicted in Fig. 1. Many implementations have been based on a fixed constellation of light sources (strobes),¹ operated one at a time [1, 5, 8, 14, 26]. Such setups suffer from a low efficiency of light power, since almost all light sources (or illumination directions)

¹Other systems use mechanical scanning of the lighting direction [3, 10, 13, 15, 18, 19]. Obviously, such scanning methods are very slow.




Figure 1. An object is viewed under varying illumination directions. When only a single light source used at a time, the captured images may be dark and noisy.

are “off,” for any image acquisition. This inefficiency may translate into long exposure times (making acquisition of moving objects very difficult), very poor angular sampling of illumination directions, or poor SNR. In the following list we describe several cases in which these problems appear.

Shadows and dark and bright albedos typically coexist in the same frames. Let the system expose the brightest points without saturation. Then, image regions corresponding to low radiance will be captured with low SNR.

Specular highlights are limited to small parts of the image, but may be much brighter than most image parts. Suppose the systems is set to avoid saturation of the highlights, then for the rest of the image, the signal readings are low.

Low power illumination sources may not cast enough light to sufficiently illuminate even the bright scene points. Can’t we always use a brighter source? In practice we cannot due to a tradeoff between the directional resolution of the setup and the power of each source. We want to have a high directional resolution of the illumination setup, with hundreds, or even tens of thousands of sources illuminating the objects. It becomes a practical problem to make the sources dense enough and keep each of them at high power. It is much easier to create systems having a very high directional resolution made up of low power sources, as we show in Sec. 6.

Problems of low object radiance may be overcome using long exposures [4, 15, 17] for each illumination direction. However, long exposures significantly increase the total acquisition time. In addition, dark current noise increases with exposure time.

3 Solution by Multiplexing

3.1 The Case of 3 Sources

For a moment, let us consider a special case in which the number of light sources is three, typically the case with photometric stereo. We label the light sources as 1, 2 and 3.

We denote the acquired measurements by a . The image irradiance under one of the sources is denoted by i , and an estimate of i is denoted by \hat{i} . Suppose that for each acquired image, only one source is “on.” The estimated intensity at a pixel due to any of the sources is trivially given by

$$\begin{bmatrix} \hat{i}_1^{\text{single}} \\ \hat{i}_2^{\text{single}} \\ \hat{i}_3^{\text{single}} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \end{bmatrix}. \quad (1)$$

In this case, on average only a 1/3 of the illumination resources are exploited for any single measurement (pixel). A more efficient method uses two sources for each acquired image. Each of the three acquired measurements exploits in average 2/3 of the illumination resources. The values acquired by the detector are now

$$\begin{bmatrix} a_{1,2} \\ a_{2,3} \\ a_{1,3} \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} i_1 \\ i_2 \\ i_3 \end{bmatrix}. \quad (2)$$

The multiplexing of illumination sources causes more light to be sensed at any acquired measurement. While the intensities corresponding to the individual light sources are not obtained as trivially as in the method of Eq. (1), they can be easily demultiplexed from the measurements:

$$\begin{bmatrix} \hat{i}_1^{\text{decoded}} \\ \hat{i}_2^{\text{decoded}} \\ \hat{i}_3^{\text{decoded}} \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 1 & -1 & 1 \\ 1 & 1 & -1 \\ -1 & 1 & 1 \end{bmatrix} \begin{bmatrix} a_{1,2} \\ a_{2,3} \\ a_{1,3} \end{bmatrix}. \quad (3)$$

What has been gained from the multiplexing process? Suppose each measurement (e.g., $a_1, a_3, a_{2,3}, a_{1,2}, \dots$) includes an independent additive noise having variance σ^2 . This noise level is the same for all images obtained by Eq. (1). However, it is easy to show that the noise variance reduces to $(3/4)\sigma^2$ in the images extracted from the lighting-multiplexed acquired measurements, using Eq. (3).

Thus, for the same number of measurements (three), the multiplexing scheme yields a better signal to noise ratio in the final output. The only cost is a negligible demultiplexing calculation. It can be said that at practically no additional cost, multiplexing leads to better results (up to limitations described in Sec. 4). Similar considerations exist in domains completely unrelated to illumination: some color cameras use cyan, magenta, and yellow filters, in order to extract better red-green-blue color images [23].

3.2 General Light Multiplexing

Consider the setup depicted in Fig. 2. The object is illuminated by many light sources simultaneously, using a multiplexing code. This creates a strong irradiance of the object, leading to bright, clear acquired images. The acquired images are later decoded (demultiplexed) on a computer.

Let the vector $\vec{\Theta}$ parameterize the direction from which the source illuminates an object point. The vector $\vec{\Theta}$ is measured in the global coordinate system of the illumination system.² Let $i_{\vec{\Theta}}(x, y)$ denote the value of a specific im-

² $\vec{\Theta}$ is unrelated to the surface normals of the object.




Figure 2. In Hadamard-multiplexed illumination, about half of the light sources are on simultaneously, creating brighter, clearer captured images. These images are later demultiplexed (decoded) on a computer. The setup can be created by projecting light patterns containing bright segments on a wall/screen. Each segment behaves as an independent source.

age pixel (x, y) under a single, narrow light source at $\vec{\Theta}$. We require this value for a range of lighting directions, thus $\mathbf{i}(x, y)$ denotes the vector of values of the image irradiance at that single pixel, with varying $\vec{\Theta}$. The length of $\mathbf{i}(x, y)$ is n , corresponding to the number of distinct lighting directions. We denote by $\mathbf{a}(x, y)$ the vector of measurements acquired under different lighting settings (typically with multiple simultaneous illumination sources). The number of acquired images (the length of $\mathbf{a}(x, y)$) equals n . We note that, while we deal with image intensity measurements, they are equivalent to the object radiance.³

The images acquired with different lighting conditions represent light energy distributions. Therefore, $\mathbf{i}(x, y)$ and $\mathbf{a}(x, y)$ are additive quantities, and are related to each other by a linear superposition:

$$\mathbf{a}(x, y) = \mathbf{W} \mathbf{i}(x, y) , \quad (4)$$

where \mathbf{W} is a weighting matrix. Each row m in the matrix \mathbf{W} denotes which of the light sources are “on” and which are “off” when the image is acquired. Each column s in this matrix corresponds to a specific illumination source, equivalent to a specific $\vec{\Theta}$. For this reason, in this paper we use the notation $i_s(x, y)$ interchangeably with $i_{\vec{\Theta}}(x, y)$. We estimate $\mathbf{i}(x, y)$ by

$$\hat{\mathbf{i}}(x, y) = \mathbf{W}^{-1} \mathbf{a}(x, y) . \quad (5)$$

When only a single light source is “on” at any time, $\hat{\mathbf{i}}(x, y)$ is equal to a raw measured value. Thus $\mathbf{W} = \mathbf{I}$, where \mathbf{I} is the identity matrix. However, the intensities per illumination direction can be multiplexed in the acquired

³The transformation between the radiance and the image $i_{\vec{\Theta}}(x, y)$ is a multiplication by a factor depending on the camera parameters, independent of the lighting and the object.

measurements, i.e., \mathbf{W} can be general. Then, each measurement $\mathbf{a}(x, y)$ simultaneously acquires energy corresponding to lighting from multiple directions. Thus, the energy in $\mathbf{a}(x, y)$ can be made larger than in single-source lighting, potentially increasing the quality of the estimated $\hat{\mathbf{i}}(x, y)$.

3.3 The Optimal Multiplexing Code

Suppose that statistically independent additive noise η having zero mean and variance σ^2 is present in the measurements. The estimation (Eq. 5) propagates this noise to the final output $\hat{\mathbf{i}}(x, y)$. The output noise vector is $\mathbf{W}^{-1}\eta$. At each pixel (x, y) , the covariance matrix Σ of $\hat{\mathbf{i}}$ is

$$\begin{aligned} \Sigma &= \mathcal{E} \left\{ \left[\hat{\mathbf{i}}(x, y) - \bar{\mathbf{i}}(x, y) \right] \left[\hat{\mathbf{i}}(x, y) - \bar{\mathbf{i}}(x, y) \right]^t \right\} \\ &= \sigma^2 (\mathbf{W}^t \mathbf{W})^{-1} , \end{aligned} \quad (6)$$

where \mathcal{E} denotes expectation and $\bar{\mathbf{i}}(x, y) = \mathcal{E}[\mathbf{i}(x, y)]$. The mean squared error of $\hat{\mathbf{i}}(x, y)$ at each pixel is then

$$\text{MSE} = \frac{1}{n} \text{Trace}(\Sigma) = \frac{\sigma^2}{n} \text{Trace}[(\mathbf{W}^t \mathbf{W})^{-1}] . \quad (7)$$

In acquisition under a single source, $\mathbf{W} = \mathbf{I}$, thus

$$\text{MSE}_{\text{single}} = \sigma^2 . \quad (8)$$

We aim to maximize the signal to noise ratio of $\hat{\mathbf{i}}(x, y)$. Thus, the multiplexing matrix \mathbf{W} should minimize the MSE.

An analogous mathematical problem was encountered in the 1970's in the fields of spectrometry and X-ray astronomy [9]. Let the elements of the matrix \mathbf{W} be $w_{m,s}$, where⁴ $0 \leq w_{m,s} \leq 1$. The matrix \mathbf{W} that has these characteristics and optimizes the MSE is called an **S**-matrix [9, 27]. If $(n+1)/4$ is an integer, the rows of the **S** matrix are based on *Hadamard codes* of length $n+1$. Ref. [9] details recipes for creating **S**. Briefly, the characteristics [9] of **S** are:

- The value of each of its elements $w_{m,s}$ is either 0 or 1. Thus each light source is either “on” or “off.”
- Each row or column has n elements: $(n+1)/2$ have the value 1, and $(n-1)/2$ have the value 0. Thus, the light energy corresponding to a little more than half of the sources is captured in each acquired multiplexed measurement.
- Inverting **S** is simple, no matter how large it is. Defining $\mathbf{1}_n$ as an $n \times n$ matrix, all of whose elements are 1's,

$$\mathbf{S}^{-1} = [2/(n+1)](2\mathbf{S}^t - \mathbf{1}_n) . \quad (9)$$

Thus except for the global factor of $2/(n+1)$, each of the elements of \mathbf{S}^{-1} is either 1 or -1.

The matrix \mathbf{W} (or **S**) describes the binary state of the illumination sources (“on” or “off”), and is thus independent of the pixel coordinates (x, y) . As an example [9], an **S** matrix for $n = 7$ is

⁴Incoherent light energy from a source is not subtracted by multiplexing. It is also not amplified (optical amplification occurs only in specialized media). For this reason, $0 \leq w_{m,s} \leq 1$.

$$\mathbf{S} = \begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 1 \\ 0 & 1 & 1 & 1 & 0 & 1 & 0 \end{bmatrix}. \quad (10)$$

In this case, the rows of \mathbf{S}^{-1} are cyclic permutations of the row vector $(1/4)[1 \ 1 \ 1 \ -1 \ 1 \ -1 \ -1]$.

For the \mathbf{S} matrix,

$$\text{MSE}_{\text{Hadamard}} = \frac{4n\sigma^2}{(n+1)^2} \rightarrow \frac{4\sigma^2}{n} \text{ for large } n. \quad (11)$$

We measure the noise by the root-mean-squared (RMS) error. Following Eqs. (8,11), the SNR of the two methods are related as

$$\frac{\text{SNR}_{\text{Hadamard}}}{\text{SNR}_{\text{single}}} = \frac{\sqrt{n} + (1/\sqrt{n})}{2} \approx \frac{\sqrt{n}}{2}. \quad (12)$$

This increase in accuracy has been termed in the context of spectrometry [9] as the *multiplex advantage*. For example, if we have $n = 255$ light sources (direction samples), then multiplexing will increase the accuracy of the estimated object radiance by a factor of about 8. If we have 10^4 light sources (as will be described in Sec. 6), then the multiplex advantage will be 50. Please note that this improvement is obtained although the *number of acquired images is the same for single-source and multiplexed measurements*.

4 Pros and Cons of Multiplexing

Consider a case in which a single source is sufficient to create images with high intensity measurement values, for example, a value 200 in an 8-bit camera. If the object is diffuse, then other sources may yield a similar value. Then, obviously, turning on several such sources simultaneously will cause the image to saturate, thereby ruining the data at the saturated points. This indicates that multiplexing has limitations, which we describe in this section.

For a diffuse object, let i be a typical image readout when the acquisition is done under a single source. This readout is assumed to occur when the exposure time is t_0 , which we term as the *baseline time*. For example, we may set $t_0 = 33\text{ms}$ as in video. In general, the exposure time is t , thus the signal is it/t_0 . Suppose this value is estimated by multiplexing N sources of similar irradiance over the object, where $N \leq n$. Following Eqs. (11,12), the SNR of \hat{i} is

$$\text{SNR} \sim \frac{i}{\sigma} \frac{t}{t_0} \frac{\sqrt{N} + (1/\sqrt{N})}{2} \sqrt{M}. \quad (13)$$

Here M is the number of frames taken with the same exposure settings for each illumination pattern. Noise reduction by simple averaging of redundant frames is accounted for by \sqrt{M} . Note that Eq. (13) assumes a diffuse object with a constant shadowing state, i.e., each illumination source yields a similar signal in the image.

We aim to recover the images under n individual lighting directions. The total acquisition time and the total baseline acquisition time are

$$T = ntM \quad \text{and} \quad T_0 = nt_0, \quad (14)$$

respectively. This is true also for illumination multiplexing, since it uses the same number of frames n . Eq. (14) applies also when $N < n$, i.e., when only part of the sources are multiplexed. Thus,

$$\text{SNR} \sim \frac{i}{\sigma} \frac{1}{\sqrt{M}} \frac{T}{T_0} \frac{\sqrt{N} + (1/\sqrt{N})}{2}. \quad (15)$$

For a single source ($N = 1$),

$$\text{SNR}_{\text{single}} \sim \frac{i}{\sigma} \frac{1}{\sqrt{M}} \frac{T}{T_0}. \quad (16)$$

The above derivation assumes that the acquired images are not saturated. However, saturation limits the ability to multiplex, since $(N+1)/2$ sources are “on” per frame. If v is the saturation value of the camera, then we must bound t and N so that

$$i(N+1)t/(2t_0) \leq v \Rightarrow (N+1)t < 2(v/i)t_0. \quad (17)$$

We now look at special cases of interest.

Dim objects or sources: The acquired images are far from saturation if $i \ll v$. According to Eq. (15,17), we can increase the SNR by extending the exposure time t , and hence T . Alternatively (or in conjunction), the SNR can be increased by increasing N , i.e., by multiplexing more sources. Looking at Eqs. (15,16), we see that:

- 1) For a fixed acquisition time T (e.g., $T = T_0$), the decoded images should have $\approx \sqrt{N}/2$ better SNR than those acquired under a single source.
- 2) Rather than increasing the SNR by increasing T , we may acquire all the images faster by keeping $T = T_0$ constant and increasing N . Acquisition of illumination-multiplexed images is then $\sqrt{N}/2$ faster than single-source acquisition. For example, if $N = 255$, we boost the speed by a factor of ≈ 8 , while if $N \approx 1000$ then the factor is ≈ 16 .
- 3) Let T extend so that under a single-source the SNR matches the SNR obtained through multiplexed illumination. Then, in the total amount of time needed to capture N single-source images, we can capture $N^2/4$ illumination-multiplexed images. This enables capturing a larger number of illumination direction ($\vec{\Theta}$) samples.

These benefits apply to image regions which are dim due to the situations listed in Sec. 2: significant spatial variation in image radiance due to albedo variations, shadows and highlights, or sources which are individually dim.

Bright objects and sources: When the acquired images saturate, Eq. (17) takes effect. At the limit,

$$t/t_0 = 2v/[i(N+1)]. \quad (18)$$

Using this limit in Eqs. (14,15),




Figure 3. [Top] Multiplexing does not apply to bright diffuse objects, which may saturate. [Bottom] A highlight forces low exposure settings to avoid its saturation. Multiplexing brightens most of the image, but does not saturate the highlights, thanks to their locality.

$$\text{SNR} \sim \frac{\sqrt{iv}}{\sigma} \sqrt{\frac{T}{T_0} \frac{1 + (1/N)}{\sqrt{2}}} \quad \text{for large } N. \quad (19)$$

At the saturation limit, we would like to *avoid* multiplexing, since the SNR somewhat decreases when N increases. The reason for this behavior is that at the limit, we can only increase N at the expense of the single frame exposure time t (Eq. 18). This undermines our goal of capturing as much light as possible per frame.

To conclude, illumination multiplexing should be done, as long as it does not hit the saturation bound (See the top of Fig. 3). The number of light sources to multiplex

$$N = \min \{n, \text{int} [2vt_0/(it) - 1]\} \quad (20)$$

may depend on the image value i . For dimmer object, we can multiplex more sources for benefit.

If the image contains a large variability of brightness values, then for some image parts we would want to extract the information without multiplexing at all (single-source images). For others we would want to extract the information from the full multiplexing of the sources ($N = n$). For some image parts, we may wish to multiplex only several sources ($N < n$) per frame. This suggests that high dynamic range data of all scene points can be obtained by taking several illumination sequences, each with a different level of multiplexing. This combined method might resemble other methods which use multiple exposures with varying exposure times [15, 17]. Nevertheless, the multiplexing method, with a constant exposure time $t = t_0$, takes a much shorter acquisition time than methods which use long exposures.

Finally, recall that this limitation analysis was made on the assumption of a diffuse object. When differences in brightness are due to specular highlights, illumination multiplexing is much more efficient, as we show next.

5 Robustness to Specular Highlights

Studies of appearance with specularities [6, 13, 15, 16] can benefit from our multiplexing scheme. Each measurement is described by a row m in Eq. (4). The acquired value is

$$a_m(x, y) = \sum_{s=1}^n w_{m,s} i_s(x, y) \quad (21)$$

We represent the intensity $i_s(x, y)$ as a sum of a diffuse component and a specular component:

$$a_m(x, y) = \sum_{s=1}^n w_{m,s} [i_s^{\text{diffuse}}(x, y) + i_s^{\text{specular}}(x, y)]. \quad (22)$$

The acquired image is composed of such components too:

$$a_m(x, y) = a_m^{\text{diffuse}}(x, y) + a_m^{\text{specular}}(x, y) \quad (23)$$

where

$$a_m^{\text{diffuse}}(x, y) = \sum_{s=1}^n w_{m,s} i_s^{\text{diffuse}}(x, y) \quad (24)$$

A highlight due to specular reflection at a pixel (x, y) does not occur for most source directions $\vec{\Theta}$. Rather, it occurs if the illumination comes from a very narrow solid angle around a single direction. For a highly specular surface we thus say that, only one source $\tilde{s}(x, y)$ produces a specular highlight at (x, y) which can be seen from the position of the camera. Therefore,

$$i_s^{\text{specular}}(x, y) = i_{\tilde{s}}^{\text{specular}}(x, y) \delta[s, \tilde{s}(x, y)] \quad (25)$$

It follows that

$$a_m^{\text{specular}}(x, y) = w_{m, \tilde{s}(x, y)} i_{\tilde{s}}^{\text{specular}}(x, y) \quad (26)$$

Suppose that in a single-source image, the light source is “on” in a direction corresponding to the highlight, i.e., $w_{m,s} = \delta(s, [\tilde{s}(x, y)])$. The acquired image is then

$$a^{\text{single}}(x, y) = i_{\tilde{s}}^{\text{diffuse}}(x, y) + i_{\tilde{s}}^{\text{specular}}(x, y) \gg i_{\tilde{s}}^{\text{diffuse}}(x, y) \quad (27)$$

In such cases, we get the familiar situation in which the specular highlight at (x, y) is much brighter than most of the image pixels, which measure only the diffuse component (See the bottom of Fig. 3). This creates a problem of dynamic range (see Sec. 2).

In contrast, in our multiplexed method, when the light source corresponding to the highlight is “on,” half of the rest of the sources, which do not create a highlight in (x, y) are “on” as well. Then,

$$a^{\text{multiplexed}}(x, y) \sim [n/2] i_{\tilde{s}}^{\text{diffuse}}(x, y) + i_{\tilde{s}}^{\text{specular}}(x, y) \quad (28)$$

The diffuse component in the acquired image $a^{\text{multiplexed}}$ is significantly brighter than in a^{single} , while the specular component is (almost) not amplified. This is illustrated in the bottom of Fig. 3. This greatly reduces the dynamic range problem.




Figure 4. Results of an experiment. All images are contrast stretched for display purposes. [Left] The images are acquired with multiplexed illumination. [Middle] Two images decoded by our method, each showing the objects as if illuminated by a single light source at a different direction. [Right] Corresponding images acquired by single-source illumination. The single-source images have a significantly lower SNR than their corresponding decoded images, and low gray level information.

6 Experiments

6.1 The Implementation Setup

The setup includes three elements: a PC-controlled projector, a white diffuse wall, and a camera. The projector projects patterns of bright and dark segments on the wall. The illuminated segments on the wall diffusely reflect light into the room, acting as separate light sources, as depicted in Figs. 1,2. This novel design allows a convenient and accurate computer control of the high resolution light “sources.” The light from the wall illuminates the object. We imaged the object with a Sony NTSC monochrome camera having a linear radiometric response.

The wall is divided into $n = 255$ segments, each of which turns “on” and “off” according to the encoded patterns, as in Fig. 2. This setup can easily be scaled to produce tens of thousands of dense samples of the illumination directions, since the projector display has millions of pixels. For each of the demonstrations, we acquire images using our multiplexed illumination method. In addition, we also acquire 255 images under the corresponding individual sources (segments), using the same setup parameters (projector brightness, exposure time, lens aperture and camera gain).

As in any implementation, imperfections occur. In particular, the dark sources (wall segments) are not completely dark due to stray light and inter-reflections in the illumination system. As we prove in the Appendix, acquisition under multiplexed lighting is more robust to this problem than acquisition under a single source. Following the results in the Appendix, we image the objects when a “dark pattern” is projected. This “darkness image” is subtracted from all the images acquired under active lighting.

6.2 Superior Signal to Noise Ratio

The viewed scene in this demonstration is composed of several simple shapes. Two of the images acquired under multiplexed illumination are displayed on the left part of Fig. 4. The captured images are bright, making the noise insignificant. Based on all of the acquired multiplexed frames, we derive the images with demultiplexed illumination, as if




Figure 5. The noise in the decoded images [Solid], which is ~ 0.08 graylevels, is significantly smaller than the noise in the single-source images [Dotted], which is ~ 0.6 graylevels.

each image is illuminated by a single, small source. Two of the decoded images are displayed in the middle part of Fig. 4. The corresponding two images taken under a single-source are displayed on the right part of Fig. 4. The single-source images are very dark, and are linearly contrast stretched in post-processing for display purposes.

The fact that any decoded image reproduces the single-source image is easily seen in the cast shadows. Yet, the decoded images have a much better quality than the images acquired under a single source. For a quantitative analysis, we examine three patches over the flat shape faces. The noise in each patch is estimated as the standard deviation of its values at any image. In total, there are $255 \times 3 = 765$ measurements taken under single-source conditions, and corresponding measurements of the decoded images. Fig. 5 plots samples of these corresponding noise measurements. It is easily seen that the noise in the decoded images is much smaller than in the single-source images. On average, the ratio between corresponding noise levels is 7.97, consistent with the theoretical multiplex advantage of $\sqrt{256/2} = 8$, predicted by Eq. (12).



Figure 6. Experimental results of imaging shiny objects.




Figure 7. Rendering the object based on the decoded images, as if illuminated by a wide angle scene. Specular reflections are densely rendered thanks to the large number of sources.

6.3 Specular Highlights and Rendering

Fig. 6 shows results of an experiment done with specular objects. For display purposes, we saturated the specular highlights in order to uncover the object details. Yet we stress that none of the raw captured images has saturated pixels. As in Sec. 6.2, the decoded image fits the single-source image, but its noise is significantly reduced.

Thanks to the ability of our method to handle low signals, the available projector light power is divided into many small sources. This high density of illumination samples enables a more realistic image-based rendering of specular reflections. The cup shown in Fig. 6 is assumed to have a gray reflectance. In Fig. 7 it is rendered as if it is illuminated by a window, out of which is a natural outdoor scene. The scene-window is specularly reflected from various places on the cup. The resolution of the reflected image is crude, having 255 “pixels” to represent it, corresponding to the illuminating segments in the image acquisition. Yet, this resolution is sufficient for the example in Fig. 7, since the reflection occupies a couple of hundred pixels on the acquired image. It is easy to scale the method to thousands of illumination segments and more, by using the appropriate Hadamard codes.

6.4 Color Multiplexing

Projecting color patterns enables us to capture the object colors with the monochrome camera used in the previous demonstrations. However, projecting red/green/blue colors means that each measurement captures $\approx 1/3$ of the spectrum. It is more efficient to multiplex the color channels, in analogy to some camera mosaics [23], and then decode the primary colors from them. Hence, in conjunction to multiplexing the illumination *direction*, we also multiplex the illumination *color*. We project cyan, magenta, and yellow patterns, each capturing $\approx 2/3$ of the color bandwidth. Af-




Figure 8. [Top] Monochrome images taken under illumination by cyan, magenta and yellow patterns, which multiplex the illumination color as well as direction. [Bottom] Decoded color images of a face mannequin corresponding to different illumination directions.

ter decoding the illumination direction and color, we obtain true-color images as if illuminated by a single white source. For an object, this demonstration uses a face mannequin. Decoded images corresponding to different illumination directions are shown in Fig. 8.

7 Discussion

Imaging objects under variable lighting conditions is an important aspect of a broad range of computer vision and image based rendering techniques. Multiplexed illumination can be highly beneficial to these techniques. It enables more accurate imaging of dim objects. It is faster than methods that increase the exposure time, and it facilitates easy implementation of high resolution lighting systems. We expect this approach to find a wide range of applications.

There are still questions open for further research, especially about the tradeoffs of the approach. It may be possible to increase the upper bound limit on the number of multiplexed sources N , beyond those roughly estimated in Sec. 4. The reason for this hypothesis is that the brightness of even a diffuse scene point can vary dramatically with illumination direction; it may be in a shadow for some of the light sources. We also explore the application of the multiplexing principle in other domains of imaging and vision.

A Appendix: Non-Zero “Darkness”

Frequently, light comes from illumination sources which are supposed to be completely dark. This is caused by stray light in the illumination system or ambient light from possible auxiliary sources. We now show that our multiplexed illumination scheme is much more robust to such disturbances than single-source illumination.

Typically, such disturbances increase the radiance of all the illumination sources, no matter if they are “on” or “off” for a specific measurement m . Thus, the values $w_{m,s}$ are perturbed by δw_s . This perturbation propagates to the acquired measurements, which are perturbed by

$$\delta a_m(x, y) = \sum_{s=1}^n \delta w_s i_s(x, y) . \quad (29)$$

Suppose that the recovery stage (5) uses the ideal, unperturbed matrix \mathbf{W}^{-1} , ignoring the unknown perturbations that occurred in the acquisition stage. The measurements, perturbed as in Eq. (29) affect the recovered estimation by

$$\delta \hat{i}_k(x, y) = \sum_{m=1}^n (\mathbf{W}^{-1})_{k,m} \delta a_m(x, y) , \quad (30)$$

where $(\mathbf{W}^{-1})_{k,m}$ is the element at row k and column m of \mathbf{W}^{-1} . Combining Eqs. (29) and (30),

$$\delta \hat{i}_k(x, y) = \sum_{s=1}^n \delta w_s i_s(x, y) \sum_{m=1}^n (\mathbf{W}^{-1})_{k,m} . \quad (31)$$

If Hadamard coding is used, then $\mathbf{W}^{-1} = \mathbf{S}^{-1}$. According to Eq. (9), $\sum_m (\mathbf{S}^{-1})_{k,m} = (2/n+1)$. On the other hand, in single-source imaging $\mathbf{W}^{-1} = \mathbf{I}$, thus $(\sum_m \mathbf{W}^{-1})_{k,m} = 1$. Therefore,

$$\delta \hat{i}_{\text{single}} = \sum_{s=1}^n \delta w_s i_s \quad \delta \hat{i}_{\text{hadamard}} = \frac{2}{n+1} \delta \hat{i}_{\text{single}} . \quad (32)$$

Thus, when our multiplexed illumination is used, the effect of illumination perturbations is much smaller than in single-source illumination.

By taking an image of the object with all the sources in the “off” state, we get some low estimate of the perturbation image $\sum_{s=1}^n \delta w_s i_s(x, y)$. This “darkness image” can then be subtracted from $a(x, y)$. This compensation is partial, since part of the ambient light is due to inter-reflections in the illumination apparatus, originating from the light which does hit its parts. Moreover, this dark image of the object can be expected to be relatively noisy.

Acknowledgments

Yoav Schechner is a Landau Fellow - supported by the Taub Foundation, and an Alon Fellow. The work was supported by the Ollendorff Center in the Elect. Eng. Dept. at the Technion. This work was also supported by National Science Foundation grants ITR IIS-00-85864, PECASE IIS-9703134, EIA-02-24431, IIS-03-08185 and KDI-99-80058.

References

- [1] Y. Adini, Y. Moses and S. Ullman, 1997, “Face recognition: the problem of compensating for changes in illumination direction,” IEEE Trans. PAMI 19:721-732.
- [2] R. Basri and D. Jacobs, 2003, “Lambertian reflectance and linear subspaces,” IEEE Trans. PAMI 25:218-233.
- [3] K. J. Dana, B. van Ginneken, S. K. Nayar, and J. J. Koenderink, 1999, “Reflectance and texture of real-world surfaces,” ACM TOG 18:1-34.
- [4] P. Debevec, 2002 “Image-based lighting” IEEE Computer Graphics & Applications 22:26-34.
- [5] P. Debevec and D. Lemmon, 2001, “Image-based lighting” Course #14, SIGGRAPH, www.debevec.org/IBL2001/.
- [6] R. Epstein, P. W. Hallinan and A. L. Yuille, 1995, “ 5 ± 2 eigenimages suffice: An empirical investigation of low-dimensional lighting models,” Proc. Physics-Based Modeling in Comp. Vis., 108-116.
- [7] H. Farid and E. H. Adelson, 1999, “Separating reflections and lighting using independent components analysis,” Proc. CVPR, Vol. I, 262-267.
- [8] A. Georgiades, P. Belhumeur, D. Kriegman, 2001, “From few to many: Illumination cone models for face recognition under variable lighting and pose,” IEEE PAMI, 643-660.
- [9] M. Harwit and N. J. A. Sloane, 1979, *Hadamard transform optics* (Academic Press, New York).
- [10] M. Hatzitheodorou, 1998, “Shape from shadows, a Hilbert space setting” J. Complexity 14:63-84.
- [11] G. Healey and T. O. Binford, 1987, “Local shape from specularly,” Proc. ICCV 151-160.
- [12] B. P. K. Horn, 1986, *Robot vision*, Ch. 10 (MIT Press).
- [13] M. Koudelka, P. Belhumeur, S. Magda and D. Kriegman, 2001, “Image-based modeling and rendering of surfaces with arbitrary BRDF’s,” Proc. CVPR, 568-575.
- [14] K. C. Lee, J. Ho and D. Kriegman, 2001, “Nine points of light: acquiring subspaces for face recognition under variable lighting,” Proc. CVPR, 519-526.
- [15] H. Lensch, J. Kautz, M. Gosele, W. Heidrich and H. Seidel, 2001, “Image-based reconstruction of spatially varying materials,” EGWR 1014-115.
- [16] Q. T. Luong, P. Fua and Y. Leclerc, 2002, “Recovery of reflectances and varying illuminants from multiple views,” Proc. ECCV.
- [17] S. Mann and R. W. Picard, 1995, “On being ‘Undigital’ with digital cameras: extending dynamic range by combining differently exposed pictures,” IS&T Annual Conf., 422-428.
- [18] S.R. Marschner, S.H. Westin, E.P.F. Lafortune and K.E. Torrance. 2000, “Image-based bidirectional reflectance distribution function measurement,” App. Opt. 39:2592-2600.
- [19] W. Matusik, H. Pfister, A. Ngan, P. Beardsley, R. Ziegler and L. McMillan, 2002, “Image-based 3D photography using opacity hulls” SIGGRAPH 427-437.
- [20] K. Nishino, Z. Zhang and K. Ikeuchi, 2001, “Determining reflectance parameters and illumination distribution from a sparse set of images for view-dependent image analysis,” Proc. ICCV Vol. I, 599-606.
- [21] M. Osadchy and D. Keren, 2001, “Image detection under varying illumination and pose,” Proc. ICCV, Vol. II, 668-673.
- [22] R. Ramamoorthi and P. Hanrahan 2002, “Frequency space environment map rendering,” ACM TOG 21:517-526.
- [23] S. F. Ray, 1994, *Applied photographic optics*, 2nd ed., 563 (Focal Press, Oxford).
- [24] I. Sato, Y. Sato and K. Ikeuchi, 2003, “Illumination from shadows,” IEEE Trans. PAMI, 25:290-300.
- [25] A. Shashua, 1997, “On photometric issues in 3D visual recognition from a single 2D image,” Int. J. of Comp. Vis. 21:99-122.
- [26] T. Sim, S. Baker and M. Bsat, 2002, “The CMU pose, illumination and expression (PIE) database,” Proc. Int. Conf. on Automatic Face and Gesture Recognition 53-58.
- [27] J. F. Turner II and P. J. Treado, 1996, “Near-infrared acousto-optic tunable filter Hadamard transform spectroscopy,” App. Spectroscopy 50(2):277-284.