Multiplexing for Optimal Lighting

Yoav Y. Schechner, *Member*, *IEEE*, Shree K. Nayar, *Member*, *IEEE*, and Peter N. Belhumeur

Abstract—Imaging of objects under variable lighting directions is an important and frequent practice in computer vision, machine vision, and image-based rendering. Methods for such imaging have traditionally used only a single light source per acquired image. They may result in images that are too dark and noisy, e.g., due to the need to avoid saturation of highlights. We introduce an approach that can significantly improve the quality of such images, in which multiple light sources illuminate the object simultaneously from different directions. These illumination-multiplexed frames are then computationally demultiplexed. The approach is useful for imaging dim objects, as well as objects having a specular reflection component. We give the optimal scheme by which lighting should be multiplexed to obtain the highest quality output, for signal-independent noise. The scheme is based on Hadamard codes. The consequences of imperfections such as stray light, saturation, and noisy illumination sources are then studied. In addition, the paper analyzes the implications of shot noise, which is signal-dependent, to Hadamard multiplexing. The approach facilitates practical lighting setups having high directional resolution. This is shown by a setup we devise, which is flexible, scalable, and programmable. We used it to demonstrate the benefit of multiplexing in experiments.

Index Terms—Physics-based vision, image-based rendering, multiplexed illumination, Hadamard codes, photon noise.

1 Introduction

Imaging objects under different light source directions is important in computer vision and computer graphics [2], [8], [23], [24], [25], [30], [33], [37], [38], [42], [43]. It is used for various purposes: object recognition and identification [2], [7], [24], [34], [36], [42], [43], image-based rendering of objects and textures [5], [23], [24], [33], [34], [37], and shape recovery [11], [16], [17], [18], [23], [47]. In these research directions and applications, images have typically been acquired under a single operating light source. Often, however, a single source may not illuminate all object parts with intensity sufficient to produce high signal-to-noise ratio (SNR) images. While this problem may be overcome using long exposures, such an approach significantly lengthens the acquisition time.

In contrast to using single sources, we claim that illuminating objects by multiple sources has significant benefits. For a given acquisition time and a set of desired illumination directions, our approach enables capturing the required information with enhanced quality. The approach reduces problems associated with dynamic range (DR), e.g., due to shadows and specular highlights. We formalize these statements in a theory of *multiplexed illumination*. For signal-independent noise, we describe the optimal scheme for multiplexing illumination sources from different directions during image acquisition and the computational demultiplexing which follows the acquisition. This scheme is based on Hadamard codes. We stress that the demultiplexing result is not an approximation: All the features that one obtains with a single source (shadows, specularities, and shading) are fully recovered.

Illumination multiplexing is not free of limitations. Thus, a great deal of this paper is devoted to the study of limitations and the consequence of various disturbing effects. Starting from a simple additive noise model, we explore each disturbing effect separately: nonideal light sources, saturation, photon noise, and γ-correction. Due to photon noise, there is a threshold gray level, under which lighting by Hadamard multiplexing is beneficial. We derive this threshold for diffuse objects and show that it increases with the number of illumination sources. In addition, we describe a design for an easily programmable lighting setup. A projector creates patterns on a white wall. The patterns reflecting off the wall serve as the object's light sources. We used this setup in lighting experiments. Partial results were presented in [39], [40]. They have been followed by studies of other researchers into additional multiplexing possibilities [3], [26], [41], [45].

2 Standard Lighting Methods

Typical methods used for illumination research and for gathering object databases under variable lighting directions are based essentially on single light sources. Many systems have been based on a fixed constellation of sources, operated one at a time [24], [28], [33], [34], [43]. Other systems use mechanical scanning[1] of the lighting direction [4], [5], [16], [23], [25], [30], [31]. In such setups, wide areas in the field of view (FOV) may be too dim, for the following reasons:

Specular highlights are limited to small image regions. The setup should be set to avoid saturation of highlights, to enable subsequent quantitative analysis. The rest of the image regions are thus much dimmer, with low signal readings.

1. Obviously, such scanning methods are very slow.

- *Y.Y. Schechner is with the Department of Electrical Engineering, Technion-Israel Institute of Technology, Haifa 32000, Israel. E-mail: yoav@ee.technion.ac.il.*
- *S.K. Nayar and P.N. Belhumeur are with the Computer Science Department, Columbia University, 450 Mudd Hall, 500 West 120 Street, New York, NY 10027. E-mail: {nayar, belhumeur}@cs.columbia.edu.*

Manuscript received 4 Dec. 2005; revised 9 May 2006; accepted 24 Aug. 2006; published online 18 Jan. 2007.
Recommended for acceptance by K. Kutulakos.
For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-0677-1205.
Digital Object Identifier no. 10.1109/TPAMI.2007.1151.

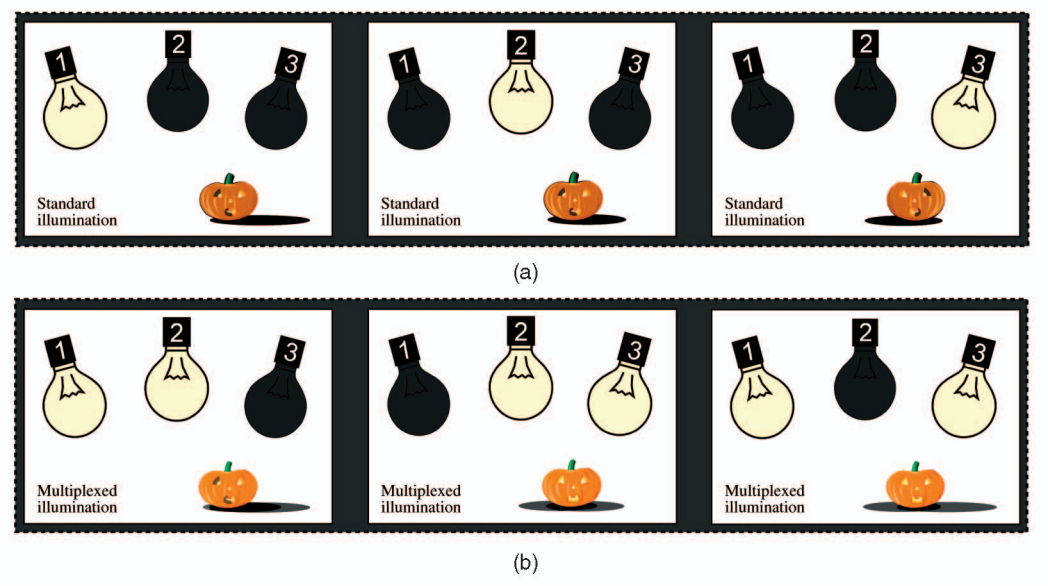

Fig. 1. Three images are taken using three light sources. (a) In a standard illumination scheme, a single light source is used at a time. (b) Multiplexed illumination.

Shadows and dark albedos typically coexist with bright albedos in any frame. The image values corresponding to shadows may be extremely low under single light source illumination. Let the system expose the brightest points without saturation. Then, image regions corresponding to low radiance are captured with low SNR.

Low power illumination sources may not cast enough light. Is it always possible to use a brighter source? This is difficult in practice, due to a tradeoff between the directional resolution of the setup and the power of each source in a constellation. We want to have a high-directional resolution of the illumination setup, with hundreds, or many more 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 Section 10. In such setups, operating only a single dim source at a time is highly inefficient. 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.

Low object radiance may be overcome by using long exposures [5], [6], [25], [29] for each illumination direction. However, long exposures significantly increase the total acquisition time. In addition, dark current noise increases with exposure time, as we detail in Section 8.

3 Solution by Multiplexing

"Two are better than one." —Ecclesiastes 4:9.

Before looking at the general case, consider for a moment 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," as depicted in Fig. 1a. The estimated intensity at a pixel due to any of the sources is trivially given by

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

On average, only a 1/3 of the illumination resources are used at a time. Moreover, any single source may create strong cast shadows and possibly specular highlights.

Now, suppose that two sources are used per acquired image (multiplexing), as in Fig. 1b. Each of the three acquired measurements exploits an average 2/3 of the lighting resources. The values acquired by the detector are now

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

While the intensities corresponding to the individual light sources are not obtained as trivially as in the method of (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? First, the DR of the raw images a has been compressed: The variability of illumination directions increases the diffuse reflection component relative to highlights (more on this in Section 6) and fills-in shadows. There is also an SNR advantage. Let each measurement (e.g., a_1, a_3, $a_{2,3}$, $a_{1,2}$..) include an independent additive noise having variance σ^2. This noise level is the same for all images obtained by (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 measurements, using (3). Thus, for the same

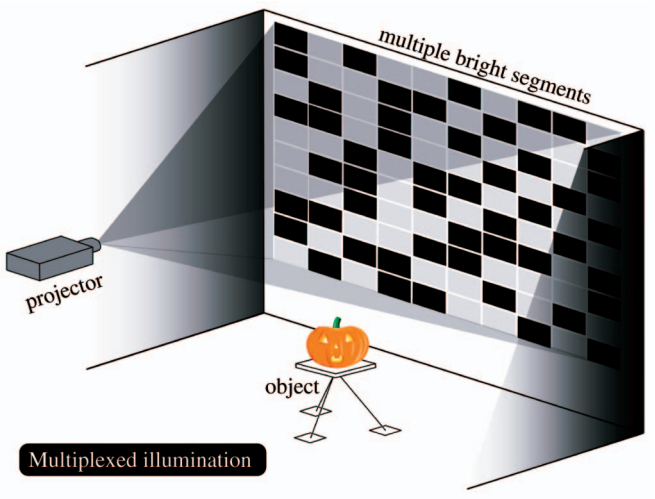

Fig. 2. In Hadamard-multiplexed lighting, $\approx 1/2$ of the sources are on simultaneously. The acquired frames are then computationally demultiplexed. A setup can use patterns containing bright segments, projected on a screen. Each segment functions as a distinct source.

number of measurements (three), multiplexing improves the SNR in the final output. The only cost is a negligible demultiplexing calculation.

3.1 General Light Multiplexing

In the setup depicted in Fig. 2, the object is illuminated by many light sources simultaneously, using a multiplexing code. The acquired images are later decoded (demultiplexed) on a computer. Let $\vec{\Theta}$ parameterize the direction from which a source illuminates an object point. This vector is measured in the global coordinate system of the illumination system.[2] Let $i_{\vec{\Theta}}(x,y)$ denote the value of a specific image pixel (x,y) under a single, narrow light source at $\vec{\Theta}$. This value is required for a range of lighting directions, thus $\mathbf{i}(x,y)$ denotes the vector of gray-level values[3] 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 light settings (typically with multiple simultaneous illumination sources). The number of acquired images (the length of $\mathbf{a}(x,y)$) equals n.

The acquired images 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

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

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 an image is acquired. Each column s in this matrix corresponds to a specific illumination source, and is equivalent to $\vec{\Theta}$. For this reason, we use the notation $i_s(x,y)$ interchangeably with $i_{\vec{\Theta}}(x,y)$. We estimate $\mathbf{i}(x,y)$ by

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

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

3. The gray level is equivalent to the image irradiance, and thus to the object radiance. The transformation between the radiance and the image irradiance is a multiplication by a factor depending on the camera parameters, independent of the lighting and the object.

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 during the measurements, i.e., $\mathbf{W}$ can be general. Then, each $\mathbf{a}(x,y)$ simultaneously acquires energy corresponding to lighting from multiple directions.

3.2 Optimal Multiplexing

Suppose that statistically independent additive noise η having zero mean and variance σ^2 is present in the measurements. The estimation (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 $\mathbf{\Sigma}$ of $\hat{\mathbf{i}}$ is

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

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

$$\mathrm{MSE} = (1/n)\mathrm{Trace}(\mathbf{\Sigma}) = (\sigma^2/n)\mathrm{Trace}\left[(\mathbf{W}^{\mathrm{t}}\mathbf{W})^{-1}\right]. \tag{7}$$

We assume for the time being, that σ is independent of the measurements $\mathbf{a}$. In single source acquisition, $\mathbf{W} = \mathbf{I}$, thus

$$\mathrm{MSE}_{\mathrm{single}} = \sigma^2. \tag{8}$$

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 1970s in the fields of spectrometry and X-ray astronomy [12], [15], [44]. Let the elements of the matrix $\mathbf{W}$ be $w_{m,s}$, where[4] $0 \le w_{m,s} \le 1$. The matrix $\mathbf{W}$ that has these characteristics and optimizes the MSE is called an S-matrix [15], [44]. If $(n+1)/4$ is an integer, the rows of the $\mathbf{S}$ matrix are based on *Hadamard codes* of length $n+1$. Harwit and Sloane [15] detail recipes for creating $\mathbf{S}$, one of which is described in the Appendix. Briefly, the characteristics [15] of $\mathbf{S}$ are:

- Each of its elements $w_{m,s}$ has a value of either 0 or 1, i.e., 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 $\mathbf{S}$ is simple, no matter how large it is. Defining $\mathbf{1}_{n\times n}$ as an $n \times n$ matrix, all of whose elements are 1s,

$$\mathbf{S}^{-1} = [2/(n+1)](2\mathbf{S}^{\mathrm{t}} - \mathbf{1}_{n\times n}). \tag{9}$$

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 $\mathbf{S}$) describes the binary state of the illumination sources ("on" or "off") and is therefore independent of the pixel coordinates (x,y). As an example [15], [39], for $n = 7$, each row of the $\mathbf{S}$ matrix is given by a cyclic permutation of the row vector

$$\mathbf{s} = [1\ 1\ 1\ 0\ 1\ 0\ 0]. \tag{10}$$

4. Incoherent 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, $0 \le w_{m,s} \le 1$.

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

For the $\mathbf{S}$ matrix,

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

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

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

This increase in accuracy has been termed in the context of spectrometry [15] as the *multiplex advantage* (MA). For example, if we have $n = 255$ light sources (direction samples), then multiplexing increases the accuracy of the estimated object radiance by a factor of about 8. This improvement is obtained although the *number of acquired images is the same for single-source and multiplexed measurements*.

The MA can be exploited in various ways, other than an increase of SNR. For instance, the exposure time of each frame (and, thus, the overall acquisition time) can be shortened by $\approx \sqrt{n}/2$ while keeping the SNR approximately the same. Alternatively, more illumination sources (each being dimmer) may be used at a similar output SNR and total acquisition time. This captures a larger number of illumination direction ($\vec{\Theta}$) samples. This is the option we pursue in Section 4. Nevertheless, the analysis we have made so far should be viewed with caution. There are limitations to multiplexing, which stem from camera saturation, nonideal light sources and photon noise. These aspects are studied in Sections 5, 7, and 8.

4 A Scalable Implementation

A constellation of light sources is difficult to scale to hundreds of sources or more. A large number of sources is useful for high definition of specular reflections and cast shadows. We thus used a lighting setup which is easily scalable to many dense samples of the illumination directions. It is flexible and simple: It includes a white diffuse wall, a camera, and a PC-controlled projector. The latter projects patterns of bright and dark segments on the wall. The illuminated wall segments diffusely reflect light into the room, acting as distinct light sources (see Fig. 2). This design allows a convenient and accurate computer control of the high resolution light "sources." The light from the wall illuminates the object, which is being imaged using a computer-controlled camera. The wall is divided into n segments, each of which turns "on" and "off" according to the encoded patterns, as in Fig. 2.

From now on, whenever we refer to this implementation, a *source* indicates an individual lighting segment. However, we emphasize that almost all the subsequent derivations apply to any lighting implementation, e.g., constellations of distinct bulbs. Some specific issues apply solely to the projector-based setup, in which case, we explicitly say so.

5 Saturation Limit for Diffuse Objects

Consider a single source that is sufficient to create high pixel values, for example, a value of 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 would cause the image to saturate, thereby ruining the data at the saturated areas. This example illustrates the caution we should take when multiplexing illumination over diffuse objects near saturation.

Let B be the number of bits per camera pixel. The gray-level corresponding to saturation is $v = (2^B - 1)$. For a diffuse object, let i be a typical gray level when acquisition is done under a single source. Assume that, for this object, each illumination source yields a similar signal in the image. The readout i is assumed to occur when the exposure time is t_0, which we term as the *baseline time*. For example, we may set $t_0 = 33ms$ as in video. In general, the exposure time is t, thus the signal is it/t_0. We aim to recover the images under n individual lighting directions. This would require n frames. In each of them, we may multiplex $N \leq n$ sources, thus $(N+1)/2$ sources are "on" per frame. Saturation limits the ability to multiplex: We must bound t and N so that

$$i(N+1)t/(2t_0) \leq v \quad \Rightarrow \quad t/t0 \leq 2v/[i(N+1)]. \tag{13}$$

Following (11) and (12), when (13) is satisfied, the SNR of $\hat{i}$ is

$$\mathrm{SNR} \ \sim \ \frac{i}{\sigma}\frac{t}{t_0}\frac{\sqrt{N} + (1/\sqrt{N})}{2}, \tag{14}$$

when σ is independent of the measurements $\mathbf{a}$.

When we hit the saturation bound posed by (13), we may face a dilemma: For the best SNR, should we use a large number of sources N and decrease the exposure time t or vice versa? Inserting this bound to (14), we get

$$\mathrm{SNR} \ \sim \ \frac{v}{\sigma}\frac{1}{\sqrt{N}} = \frac{v}{\sigma}\left(\frac{t_0}{t}\frac{2v}{i} - 1\right)^{-1/2}. \tag{15}$$

Hence, near saturation, the SNR *decreases* with N, while it increases with t. The reason for this behavior is that at the bound, we can only increase N at the expense of the single frame exposure time t (13). This undermines our goal of capturing as much light as possible. Hence, in this case, we would like to *limit* multiplexing to

$$N = \min\{n, \mathrm{int}[2vt_0/(it) - 1]\} \tag{16}$$

light sources. If most of the FOV is unsaturated but bright enough with a single source ($it \sim vt_0$), then (16) implies that no multiplexing ($N = 1$) may be the best option.

In general, N depends on the image value i. For dimmer objects, we can multiplex more sources for benefit. If the image contains a large brightness variability, then some image parts would be best acquired with no multiplexing at all (single-source images). Other parts would optimally be based on full multiplexing of the sources ($N = n$). For some image parts, we may wish to multiplex only several sources ($N < n$) per frame. Thus, high DR data of all scene points can be obtained from several illumination sequences, each with a different level of multiplexing. This method might resemble other methods which use multiple exposures with varying exposure times [5], [6], [25], [29]. Nevertheless, multiplexing with a constant exposure time $t = t_0$ requires a shorter acquisition time than methods that use long exposures.

Finally, recall that this limitation analysis was based 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.

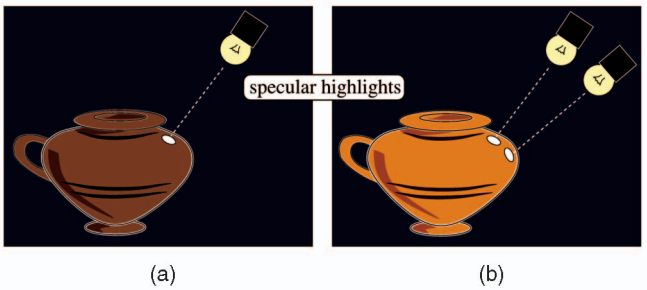

Fig. 3. Avoiding saturation at a highlight forces low exposure settings. Multiplexing brightens most of the image, but does not saturate the highlights, thanks to their locality.

6 Robustness to Specular Highlights

Studies of appearance with specularities [7], [23], [25], [27] can benefit from our multiplexing scheme. Each measurement is described by a row m in (4). The acquired value is

$$a_m(x,y) = \sum_{s=1}^{n} w_{m,s} i_s(x,y). \tag{17}$$

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

$$a_m(x,y) = \sum_{s=1}^{n} w_{m,s}\left[i_s^{\text{diffuse}}(x,y) + i_s^{\text{specular}}(x,y)\right]. \tag{18}$$

The acquired image is also composed of such components:

$$a_m(x,y) = a_m^{\text{diffuse}}(x,y) + a_m^{\text{specular}}(x,y), \tag{19}$$

where

$$a_m^{\text{diffuse}}(x,y) = \sum_{s=1}^{n} w_{m,s} i_s^{\text{diffuse}}(x,y). \tag{20}$$

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), as seen from the position of the camera. Therefore,

$$i_s^{\text{specular}}(x,y) = i_s^{\text{specular}}(x,y)\delta[s,\tilde{s}(x,y)]. \tag{21}$$

It follows that

$$a_m^{\text{specular}}(x,y) = w_{m,\tilde{s}(x,y)} i_{\tilde{s}(x,y)}^{\text{specular}}(x,y). \tag{22}$$

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

$$a^{\text{single}}(x,y) = i_{\tilde{s}}^{\text{diffuse}}(x,y) + i_{\tilde{s}}^{\text{specular}}(x,y) \gg i_{\tilde{s}}^{\text{diffuse}}(x,y). \tag{23}$$

This is a familiar situation: The specular highlight at (x,y) is much brighter than most of the image pixels, which measure only the diffuse component (see Fig. 3a). This creates a DR problem: Powerful sources that illuminate well the non-specularly reflecting areas would saturate the highlights. If the exposure is set to avoid highlight saturation, then the diffuse component may become underexposed.

In contrast, in our multiplexed method, when the highlight-creating source is "on," half of the rest of the sources, which do not create a highlight in (x,y) are "on" as well. Then,

$$a^{\text{multiplexed}}(x,y) \sim [n/2] i_{\tilde{s}}^{\text{diffuse}}(x,y) + i_{\tilde{s}}^{\text{specular}}(x,y). \tag{24}$$

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, as illustrated in Fig. 3. This greatly reduces the DR problem.

Nevertheless, there are two special cases worth mentioning. First, consider an object with a weak but nonzero diffuse reflection, such that the diffuse component in (24) is below the noise level of the camera. This occurs if $i_s^{\text{diffuse}}(x,y) < \mathcal{O}(\sigma/n)$. Then, there is no resulting improvement in the DR. Roughly, the DR is improved only if

$$\frac{2\sigma}{n} < i_s^{\text{diffuse}}, \quad \text{while} \quad [n/2] i_{\tilde{s}}^{\text{diffuse}} + i_{\tilde{s}}^{\text{specular}} < v. \tag{25}$$

Second, consider an object which has no diffuse reflection at all, i.e., $i_s^{\text{diffuse}}(x,y) = 0, \forall s,x,y$. This may occur if the entire FOV is a mirrored surface. In this case, a DR problem does *not* exist even for a single light source: The highlight value can be determined by the single measurement at $\tilde{s}(x,y)$ and there is no concern of having the nonhighlight measurement being underexposed (the latter being nonexistent).

7 Nonideal Light Sources

We now explore consequences of nonideal lighting: stray light, temporal fluctuations in the illuminating sources, and intersource difference of radiance.

7.1 Nonzero "Darkness"

Light often comes from sources that are supposed to be completely dark. This is caused by stray light in the illumination system or ambient light from possible auxiliary sources. 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, $w_{m,s}$ are perturbed by δw_s. This propagates to perturbation in the measurements

$$\delta a_m(x,y) = \sum_{s=1}^{n} \delta w_s i_s(x,y). \tag{26}$$

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 (26) affect the recovered estimation by

$$\delta\hat{i}_k(x,y) = \sum_{m=1}^{n} \left(\mathbf{W}^{-1}\right)_{k,m} \delta a_m(x,y), \tag{27}$$

where $\left(\mathbf{W}^{-1}\right)_{k,m}$ is the element at row k and column m of $\mathbf{W}^{-1}$. Combining (26) and (27),

$$\delta\hat{i}_k(x,y) = \sum_{s=1}^{n} \delta w_s i_s(x,y) \sum_{m=1}^{n} \left(\mathbf{W}^{-1}\right)_{k,m}. \tag{28}$$

In Hadamard coding, $\mathbf{W}^{-1} = \mathbf{S}^{-1}$. According to (9),

$$\left(\mathbf{S}^{-1}\right)_{k,m} = \pm 2/(n+1) \tag{29}$$

and

$$\sum_{m=1}^{n} \left(\mathbf{S}^{-1}\right)_{k,m} = 2/(n+1). \tag{30}$$

On the other hand, in single-source imaging $\mathbf{W}^{-1} = \mathbf{I}$, thus $(\sum_m \mathbf{W}^{-1})_{k,m} = 1$. Therefore,

$$\delta\hat{i}_{\text{single}} = \sum_{s=1}^{n} \delta w_s i_s \qquad \delta\hat{i}_{\text{hadamard}} = \frac{2}{n+1}\delta\hat{i}_{\text{single}}. \tag{31}$$

Thus, using multiplexing, the effect of illumination perturbations is much smaller than in single-source illumination.

Acquiring the object with all the sources in the "off" state yields an estimate of the perturbation image (26). This "darkness image" can then be subtracted from $\mathbf{a}(x, y)$. This compensation is partial since part of the ambient light is due to interreflections 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.

7.2 Noisy Sources

Section 7.1 dealt with a systematic bias in the illumination. In contrast, in this section, we analyze random noise in the sources. Fluctuations in the illumination create multiplicative noise in the image readout. This is contrary to the additive signal-independent noise, which was described in Section 3. We model the fluctuations of source irradiance as proportional to the irradiance itself, say due to instabilities in the electric current. This is consistent with specifications of light sources, in which their stability is stated as a percentage. Thus,

$$\text{STD}(w_{m,s}) = \alpha w_{m,s}, \tag{32}$$

where STD denotes the standard deviation. Here, α is a proportionality factor. Typically, $\alpha \ll 1$. In single source imaging, the pixel intensity fluctuates. The variance (VAR) of this fluctuation is $\text{VAR}_{\text{single}}[\hat{i}_s(x, y)] = \alpha^2 \hat{i}_s^2(x, y)$. Averaging over the source directions s,

$$\text{VAR}_{\text{single}}[\hat{i}(x, y)] = \alpha^2 \langle \hat{i}^{\,2}(x, y)\rangle, \tag{33}$$

where

$$\langle \hat{i}^2(x, y)\rangle = (1/n)\sum_{s=1}^{n} \hat{i}_s^{\,2}(x, y). \tag{34}$$

Let the sources be multiplexed. Due to source noise, any measurement $a_{m,s}$ is perturbed by $\delta a_{m,s}$. The variation $\delta a_m(x, y)$ affects the recovered image according to (27). If Hadamard coding is used, then according to (29),

$$\text{VAR}_{\text{hadamard}}[\hat{i}(x, y)] = \sum_{m=1}^{n}\left(\frac{2}{n+1}\right)^2 \text{VAR}[a_m(x, y)], \tag{35}$$

where we exploit the temporal independence of any source radiance, as sampled at different measurements m. The magnitude $\text{VAR}_{\text{hadamard}}(\hat{i})$ depends on the joint statistics of the light sources. Thus, we now discuss two kinds of setups. In the first one, different light sources have mutually independent fluctuations. In the second one, the scene irradiance fluctuates randomly in time, but all the sources are tightly correlated.

7.2.1 Independent Sources

Let us analyze the noise in setups based on a constellation of distinct bulbs/strobes, each having its own power supply and shielded cabling. Then, the radiance of a light source is independent of the radiance of any other source in the setup. It follows from (4) that

$$\delta a_m(x, y) = \sum_{s=1}^{n} \delta w_{m,s} i_s(x, y). \tag{36}$$

Since $\delta w_{m,s}$ is independent for different sources s, then

$$\text{VAR}[a_m(x, y)] = \sum_{s=1}^{n} \text{VAR}(w_{m,s}) i_s^2(x, y). \tag{37}$$

Combining (32), (35), and (37),

$$\text{VAR}^{\text{independent}}_{\text{hadamard}}[\hat{i}(x, y)] = \alpha^2\left(\frac{2}{n+1}\right)^2 \sum_{s=1}^{n} i_s^2(x, y) \sum_{m=1}^{n} w_{m,s}^2. \tag{38}$$

Recall that, in Hadamard codes $w_{m,s}$ is either 0 or 1, thus $w_{m,s}^2 = w_{m,s}$. Hence,

$$\sum_{m=1}^{n} w_{m,s}^2 = \sum_{m=1}^{n} w_{m,s} = (n+1)/2 \tag{39}$$

and, therefore,

$$\text{VAR}^{\text{independent}}_{\text{hadamard}}[\hat{i}(x, y)] = \alpha^2 \frac{2}{n+1} \sum_{s=1}^{n} i_s^2(x, y). \tag{40}$$

Substituting (33) and (34),

$$\text{VAR}^{\text{independent}}_{\text{hadamard}}[\hat{i}(x, y)] = \frac{2n}{n+1} \text{VAR}_{\text{single}}[\hat{i}(x, y)] \approx 2\alpha^2 \langle \hat{i}^{\,2}(x, y)\rangle. \tag{41}$$

Hence, when n is large, the multiplexing approach amplifies by $\sqrt{2}$ the STD of noise created by irradiance fluctuations.

How significant is this? It depends on the light source technology: Light strobes based on arc lamps [13] have $\alpha = \mathcal{O}(5\%)$, which is rather high and can distort quantitative photometric analysis, with or without multiplexing; strobes based on LED clusters may have a much lower α, thanks to the use of many flashing LEDs at each source position; direct current (DC) sources are very stable,[5] once their temperature reaches saturation, and can be shuttered mechanically or electro-optically during the experiment. In this case, the fluctuations can be unnoticeable with single-source illumination, hence also with multiplexing.

In any case, the output from any source can be monitored online, by bleeding a portion of this output to a light detector placed at the source (strobe). Thus, fluctuations can be measured in real time and then be compensated for in postprocessing, or even stabilized during lighting. This effectively shrinks α to insignificance. Hence, we believe that this noise source should be minor, with the use of proper lighting hardware. Nevertheless, for quantitative computer vision methods (with or without multiplexing), low α should be ensured, and it is better not to take light source stability for granted.

7.2.2 Coupled Sources

Let the sources be coupled, therefore having a strong mutual correlation. This occurs when the radiance of a *single* light source is divided into multiple channels, each of

5. DC incandescent sources are very stable for short time periods, but their performance drifts in time, as their lifetime is ≈ 100 hours. This drift can be easily monitored. DC arc sources have lifetimes of thousands of hours, but in short time periods they may fluctuate a little. This effect is minimized by *quiet* arc bulbs [14].

which irradiates the object from a different direction. This is the case in the projector-based setup, described in Section 4, where a bulb inside the projector is the origin for the light emanated from all the wall segments. Temporal fluctuations of this bulb create the same simultaneous fluctuations in the light emanated from the wall segments.

We start with (32) and (36). Since the sources are coupled to one another,

$$\mathrm{STD}[a_m(x,y)] = \alpha a_m(x,y). \tag{42}$$

Substituting into (35),

$$\mathrm{VAR}^{\mathrm{coupled}}_{\mathrm{hadamard}}[\hat{i}(x,y)] = \alpha^2 \left(\frac{2}{n+1}\right)^2 \sum_{m=1}^{n} a_m^2(x,y). \tag{43}$$

Following (4),

$$\sum_{m=1}^{n} a_m^2 = \sum_{m=1}^{n} \left(\sum_{s=1}^{n} w_{m,s} i_s\right)^2 = \sum_{s=1}^{n}\sum_{c=1}^{n} i_s i_c \sum_{m=1}^{n} w_{m,s} w_{m,c}, \tag{44}$$

where s and c index the sources. For Hadamard codes,

$$\sum_{m=1}^{n} w_{m,s} w_{m,c} = (n+1)/4, \quad \text{if} \quad s \neq c, \tag{45}$$

while this sum equals $(n+1)/2$ if $s = c$, as written in (39). It follows that

$$\sum_{m=1}^{n} a_m^2 = \frac{n+1}{4}\sum_{s=1}^{n} 2i_s^2 + \frac{n+1}{4}\sum_{s=1}^{n}\sum_{c\neq s} i_s i_c. \tag{46}$$

To progress, consider first a diffuse object that is never shadowed. As discussed in Section 5, the intensity i_s is approximately the same for all frames, hence $i_s i_c \approx i_s^2$. Then,

$$\sum_{m=1}^{n} a_m^2 \approx \frac{n+1}{4}\sum_{s=1}^{n} 2i_s^2 + \frac{n+1}{4}\sum_{s=1}^{n}(n-1)i_s^2 = \left(\frac{n+1}{2}\right)^2 \sum_{s=1}^{n} i_s^2. \tag{47}$$

Inserting (47) to (43), and using (33) and (34),

$$\mathrm{VAR}^{\mathrm{coupled}}_{\mathrm{hadamard}}[\hat{i}(x,y)] \approx n\mathrm{VAR}_{\mathrm{single}}[\hat{i}(x,y)]. \tag{48}$$

Equation (48) is not limited to diffuse objects. It can be shown to approximate also the case of a specular highlight at (x,y) caused by a specific source $\tilde{s}$. This follows the constraint $i_{\tilde{s}} < v$ and assuming that i_s is approximately the same for all $s \neq \tilde{s}$, while $\alpha \ll 1$ and n is large.

The noise in (48) is much worse than the one in (41). When the light sources are coupled, the STD of noise originating from light power fluctuations [35] increases with $\sqrt{n}$. This noise may undermine the projector-based setup described in Section 4. Attention should thus be drawn to keep α very low. As in Section 7.2.1, this can be done by bleeding a portion of the projector bulb's radiance to a detector for bulb stabilization or for compensation in postprocessing. We avoided the problem in experiments by exploiting the temporal behavior of the projector (see Section 10).

7.3 Unequal Illumination Sources

Suppose different illumination sources systematically do not illuminate the scene equally, when turned "on." As mentioned in Section 3.1, $i_{\vec{\Theta}}(x,y)$ is equivalent to the object radiance, when illuminated from a specific direction $\vec{\Theta}$. Had there been no shadows and interreflections, the radiance would have been determined by the reflectance function of the object $r_{\vec{\Theta}}(x,y)$, and the intensity of the light source. Since shadows and interreflections exist, the effective reflectance function $\mathbf{r}^{\mathrm{eff}}(x,y)$ indicates how the object pixel corresponding to (x,y) effectively reacts to various illumination directions $\vec{\Theta}$. Can we make quantitative claims about the effective reflectance function?

For a general weighting matrix $\mathbf{W}$, we may interpret the coefficient $w_{m,s}$ as the radiance of a single light source whose index is s, as measured from the position of the object. When the radiance of the illumination sources is unequal, we may interpret it as having $w^{\mathrm{practical}}_{m,s} = f(s)w_{m,s}$, where $0 < f(s) \leq 1$ is the reduction of the intensity of source s, relative to the ideal full brightness value of 1. The function $f(s)$ is unknown for uncalibrated illumination radiance. The acquired measurement is

$$a_m(x,y) = \sum_{s=1}^{n} w^{\mathrm{practical}}_{m,s} r^{\mathrm{eff}}_s(x,y) = \sum_{s=1}^{n} f(s) w_{m,s} r^{\mathrm{eff}}_s(x,y). \tag{49}$$

Comparing this to (17) and using (5), it follows that

$$\hat{i}_s(x,y) = f(s) r^{\mathrm{eff}}_s(x,y). \tag{50}$$

There are two important consequences of (50). First, the inhomogeneity of the sources does not cause crosstalk between demultiplexed images $\hat{i}_s(x,y)$. Each recovered image corresponds to a single light source indexed s. It does not include contributions from other source indices, or from other pixels. Second, the image recovered for a specific source s at $\vec{\Theta}$ is similar to the effective reflectance function $r^{\mathrm{eff}}_s(x,y)$, yet multiplied by a scale factor.

Note that $\hat{i}_s(x,y) = 0$ if and only if $r^{\mathrm{eff}}_s(x,y) = 0$ since $f(s) > 0$. Hence, *shadows* are conserved in the decoded images, regardless of the inhomogeneity of the illumination sources. Thus, cast shadows are recovered as sharply as they should under individual point sources. Finally, we would like to note that precalibration of the radiance of sources in the setup yields $f(s)$, which in turn enables the estimation of $\mathbf{r}^{\mathrm{eff}}(x,y)$ for all pixels.

7.3.1 Adaptive Directional Resolution

When $f(s)$ has significant variations as a function of s, then the output image ensemble $\hat{\mathbf{i}}$ may suffer from DR problems. Consider an object for which $r^{\mathrm{eff}}_s \approx r^{\mathrm{eff}}_c$ for different sources s and c. If $f(s) \gg f(c)$, then $\hat{i}_s \gg \hat{i}_c$, whether the estimation is done through multiplexed lighting or through single source lighting. Thus, some resulting images may be too dark and noisy, relative to others. We minimize this effect by exploiting the flexibility of the projector setup mentioned in Section 4.

Apparently, the DR problem might have been solved by reducing the radiance of the bright sources so as to equalize them to the dim ones. However, this counteracts our desire to use as much light as possible. We thus take a different approach: We locally adapt the illumination directional resolution. This is much more efficient, since no light is lost.

First, we make a rough calibration of $f(\vec{\Theta})$, assuming that the "illuminating wall" is divided into segments of equal area. This is only done once, since the same setup is used for all the experiments. We need to estimate the irradiance *at the object*. This is achieved by replacing the object by a camera, which is placed right at the object's position. The

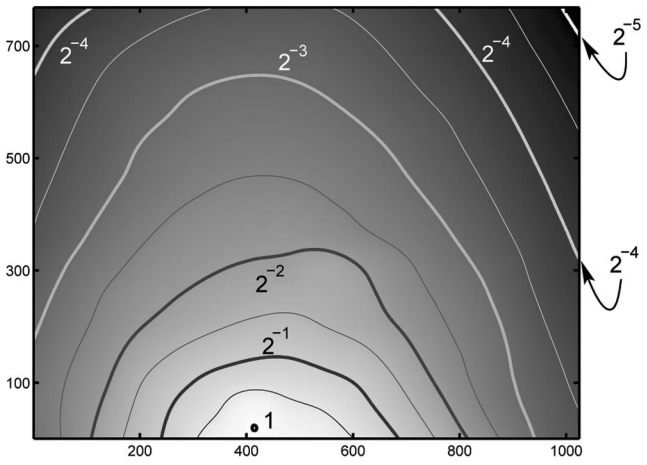

Fig. 4. The falloff $f(\vec{\Theta})$ of object irradiance, due to segments of equal area on the illuminating wall. The wall coordinates correspond to the projector's pixels. The DR of this irradiance is larger than 1:32. Using segments of uniform area to illuminate the object may thus create severe DR problems in the resulting images.

camera then acquires an image of each wall-segment,[6] thus sampling the radiance $L(\vec{\Theta})$. For a small object and a distant segment, $L(\vec{\Theta})$ is related to the object irradiance $f(\vec{\Theta})$ by

$$f(\vec{\Theta}) = bL(\vec{\Theta})\frac{\cos(\theta)}{R^2}, \tag{51}$$

where R is the distance of the segment from the object and b is a normalization factor. Here, θ is the angle at which the wall segment illuminates the object, measured relative to the wall's normal. The $\cos(\theta)$ term accounts for foreshortening a finite-area segment, as seen from the viewpoint of the calibrating camera. We know R and θ for each segment of the illuminating wall, since we know the geometry of our setup. Fig. 4 depicts $f(\vec{\Theta})$ for our setup. The lateral coordinates correspond to pixels of the projector which projects the lighting patters onto the wall. In this figure, the DR of $f(\vec{\Theta})$ is greater than 1:32. Hence, if segments of equal area were used, some images $\hat{i}_s$ would have been ≈ 32 times dimmer than others.

To reduce this DR problem, we use segments of different areas. Each segment has a fixed area for all images and experiments. However, this area is adapted to the irradiance associated with it. Segments for which $f(\vec{\Theta})$ is high, are given a smaller area on the illuminating wall, and vice versa. We achieved this using a quadtree subdivision [40] of the distribution $f(\vec{\Theta})$ shown in Fig. 4. A segment is subdivided if the integral of $f(\vec{\Theta})$ in its area is larger than a fixed threshold. The threshold is related to the final number of segments n: The smaller the threshold, the larger n. Note that, in a quadtree structure, the number of segments is always

$$n_{\text{quadtree}} = 1 + 3k_{\text{quadtree}}, \tag{52}$$

where k_{quadtree} is an integer. Therefore, to obtain $n = 255$ segments, we use $n_{\text{quadtree}} = 256$, corresponding to $k_{\text{quadtree}} = 85$. We then merge two adjacent segments of the quadtree into a single one. The two adjacent segments selected for merger are dim, so that the overall light emanated from their combination does not exacerbate the DR but moderates it. Fig. 5 depicts examples of Hadamard-coded lighting patterns composed of segments, which are created by the quadtree structure tailored to our system.

6. We used a narrow FOV lens. We directed the FOV center at the segment. This reduced problems associated with camera vignetting and foreshortening.

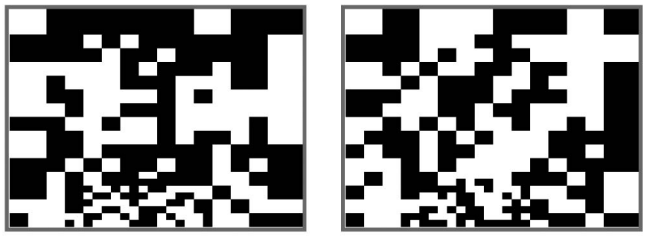

Fig. 5. Two out of $n = 255$ multiplexed lighting patterns used in our lighting system. The patterns are composed of segments. The area of each segment is determined by a quadtree algorithm, designed to reduce DR problems.

8 Photon Noise

Section 7.2 showed that unstable light sources result in signal-dependent noise. Yet, even with stabilized sources ($\alpha = 0$), the noise variance in a measured image generally depends on the measured signal. We now analyze the noise sources of high grade detectors, before studying their consequence on Hadamard multiplexing.

Photon noise. Light absorbed and detected at a pixel during exposure time is not a continuous variable. Fundamentally, it is a quantized, countable variable, having the units of photons. Similarly, the detected photons do not create a continuous electric charge at each pixel. Rather, the fundamental detection signal (charge) is quantized to indivisible *electron* units. This observation leads to a fundamental measurement noise, termed photon noise. The number of electrons generated in the detector is *random*, even if the light source is completely stabilized and the camera circuity induces no noise.

Let us first consider this ideal case. It is simple to show that under assumptions which are reasonable to visible-light photography,[7] the random number $n_{\text{electr}}^{\text{photo}}$ of photo-generated electrons is Poisson distributed [1], [10], [19]. Define $\bar{n}_{\text{electr}}^{\text{photo}} = \mathcal{E}(n_{\text{electr}}^{\text{photo}})$ as the expected number of detected electrons at a pixel, in a constant setup and illumination. This number may be obtained by averaging a very large number of repeated measurements. In this distribution,[8]

$$\text{VAR}(n_{\text{electr}}^{\text{photo}}) = \bar{n}_{\text{electr}}^{\text{photo}}. \tag{53}$$

Hence, the variance increases with the measured electric signal $n_{\text{electr}}^{\text{photo}}$. This aspect counteracts the MA: The use of more light sources at each frame increases the acquisition noise [45], which, in turn, increases the noise of the demultiplexed images. In fact, according to [15], if photon noise were the only noise source, multiplexing would have resulted in deterioration of SNR of the demultiplexed images. However, other noise sources exist, as we now detail.

7. These assumptions may be violated in high energy detectors such as those of γ rays, where any photon can generate multiple electrons at the detector [1].

8. This distribution can be excellently approximated by a Gaussian distribution having this expectation and variance, for $\bar{n}_{\text{electr}}^{\text{photo}} > 10$, which is typically the case for cameras.

TABLE 1
Specifications of Sample Cameras (The Sign "–" Indicates Unknown Data)

Camera	FPS $[\frac{1}{\text{sec}}]$	d_{well} $[e^-]$	ρ_{read} $[e^-]$	D $[\frac{e^-}{\text{sec}}]$	B	g_{electr} $[e^-]$	κ_{gray}
Dalsa 1M75	75	200000	796	-	8	784	1.05
PixeLINK A661	11.4	40000	70	6250	8	157	0.52
FastVision FastCam40	240	25000	50	1333	10	24	2.06
Redlake MotionPro HS1	500	63000	71	3842	10	62	1.18
PtGrey Dragonfly VGA	30	-	-	-	8	53	0.4
PCO Sensicam VGA	30	35000	13	0.1	12	9	1.54

The specifications are used for estimating the parameters g_{electr} and κ_{gray}. The dragonfly parameters are derived in an experiment. The sensicam is a cooled camera used in microscopy and, hence, its noise specifications D and ρ_{read} are much lower than those of the other cameras, which are made for machine vision.

Note that the number of detected electrons $n_{\text{electr}}^{\text{photo}}$ is equivalent to the gray level of the acquired pixel value a

$$a = \frac{n_{\text{electr}}^{\text{photo}}}{g_{\text{electr}}}, \tag{54}$$

where g_{electr} is the number of photo-generated electrons required to change a unit gray level. Typically, $g_{\text{electr}} \gg 1$.

Dark current. In addition to photo-electrons, some electrons are generated in each pixel by thermal excitation, irrespective of the photon irradiance. This phenomenon is termed dark current [10], [22]. The mean dark charge is proportional to the exposure time

$$\bar{n}_{\text{electr}}^{\text{dark}} = Dt. \tag{55}$$

Here, D is the detector dark current. This mean number can be calibrated and compensated for, e.g., using the process described in Section 7.1. However, the actual number of thermally generated electrons $n_{\text{electr}}^{\text{dark}}$ is governed by the Poisson distribution as well, with $\bar{n}_{\text{electr}}^{\text{dark}}$ being its mean and variance. The randomness of $n_{\text{electr}}^{\text{dark}}$ means that it cannot be fully compensated for in post processing. The random component [10], [22] is termed dark current noise.[9]

Additional noise that is considered to be signal independent is induced by electronic circuity in the camera system, as the amplifier [10] and digital quantizer. The variance of this noise can be modeled as $\rho^2 = \rho_{\text{read}}^2 + \rho_{\text{digit}}^2$. Here, ρ_{read} is the amplifier readout noise STD, while $\rho_{\text{digit}} = g_{\text{electr}}/\sqrt{12}$ is the noise of a uniform quantizer,[10] all in units of electrons (e^-). Photon noise, dark current, and ρ are mutually independent, hence [10], [19]

$$\sigma_{\text{electr}}^2 = \rho^2 + Dt + \bar{n}_{\text{electr}}^{\text{photo}}. \tag{56}$$

While the dark current noise and ρ do not depend explicitly on the optical signal, they do depend on the acquisition mode. If single source lighting is very low, one may opt for long exposures rather than multiplexing. This would increase the dark current noise variance Dt. In most uncooled cameras, this noise becomes significant if t is longer than a second.[11] Moreover, higher frame rates tend [19], [22] to increase ρ_{read}, hence cool cameras used in microscopy and astronomy often have pixel readout rates of few MHz or less. For megapixel cameras, this induces frame rates of very few Hz, and even sub Hz. On the other hand, fast cameras are noisier by orders of magnitude (examples are given later, in Table 1). In addition, communication bottlenecks trade bit-depth for frame-rates. Hence, some 10-bit or 12-bit machine vision cameras deliver high frame rates by quantizing the electric signal to 8 bits. This increases ρ_{digit}. In the following discussion, we assume that the camera is meant to be used in applications that require their highest speed. Thus, Dt is usually low.

8.1 Photon Noise Limit to Hadamard Multiplex

The noise (56) includes signal independent components, for which Hadamard multiplexing is designed. It also includes photon noise, for which multiplexing is counterproductive with respect to SNR [45]. Given the camera specifications, what is the limit to the MA?

Let n_{electr} be the number of electrons detected in a pixel. Following (53), the uncertainty of this measurement is

$$\text{STD}_{\text{electr}} = \sqrt{\bar{n}_{\text{electr}}} = \sqrt{g_{\text{electr}} a} \tag{57}$$

electrons, where $\bar{n}_{\text{electr}}$ is the expected number of electrons, which is equivalent to the expected acquired gray level a. Following (54), the electron noise in (57) induces noise on the gray-level readout a having STD of $(\text{STD}_{\text{electr}}/g_{\text{electr}}) = \sqrt{a/g_{\text{electr}}}$. Similarly, dark current noise induces gray-level noise with STD of $(\sqrt{Dt}/g_{\text{electr}})$, while the read noise STD is (ρ/g_{electr}) in these units. Overall, the noise variance of a pixel gray-level value is

$$\sigma_a^2 = \kappa_{\text{gray}}^2 + \frac{a(x,y)}{g_{\text{electr}}}, \tag{58}$$

where

$$\kappa_{\text{gray}}^2 \equiv \rho^2/g_{\text{electr}}^2 + Dt/g_{\text{electr}}^2 \tag{59}$$

encompasses the signal-independent components of the gray-level noise. For a single light source, the noise variance is

$$\sigma_{i,\text{single}}^2 = \kappa_{\text{gray}}^2 + i_s(x,y)/g_{\text{electr}}. \tag{60}$$

For a diffuse object, the values of i_s are assumed to be approximately the same, as in Section 6. Hence, $a(x,y) \approx i_s(N+1)/2$, when multiplexing N out of n sources. Demultiplexing using the $\mathbf{S}$ matrix,

9. The combined contribution of photon noise and dark current noise is termed *shot noise*.

10. Some cameras boast high-bit quantization (e.g., 12 to 16 bits) that does not truly indicate the camera DR [10]. The reason is that the use of many bits decreases ρ_{digit}, but it does not decrease ρ_{read}, which usually dominates ρ.

11. For fast cameras Dt may be significant even in video rates. This is apparent when analyzing the data in Table 1, presented in Section 8.1.1.

$$\mathrm{MSE}_{\hat{i},\mathrm{Hadamard}} = \frac{4N}{(N+1)^2}\sigma_a^2 = \frac{4N}{(N+1)^2}\kappa_{\mathrm{gray}}^2 + \frac{2N}{N+1}\frac{i_s(x,y)}{g_{\mathrm{electr}}}. \tag{61}$$

If photon noise is negligible, i.e., $\kappa_{\mathrm{gray}}^2 \gg a(x,y)/g_{\mathrm{electr}}$, (60) and (61) are consistent with (11) and (12). However, when photon noise is overwhelming $i_s(x,y)/g_{\mathrm{electr}} \gg \kappa_{\mathrm{gray}}^2$, (60) and (61) indicate that a demultiplexed image is more noisy, with its noise variance approximately doubled relative to a single-source image. This special case is consistent with [15].

In general, multiplexing of N sources benefits the SNR if

$$\mathrm{MSE}_{\hat{i},\mathrm{Hadamard}} \leq \sigma_{i,\mathrm{single}}^2. \tag{62}$$

Inserting the explicit expressions from (60) and (61) into (62) and performing several trivial operations yields the condition

$$i_s(x,y) \leq i_{\mathrm{limit}}^{\mathrm{photon}} \equiv g_{\mathrm{electr}}\kappa_{\mathrm{gray}}^2\frac{N-1}{N+1}. \tag{63}$$

How significant is this limitation? Recall that, in no case, should the camera be saturated. Thus, following (13), the useable single-source intensity range is

$$i_s(x,y) \leq i_{\mathrm{limit}}^{\mathrm{saturation}} \equiv \frac{2v}{N+1}, \tag{64}$$

where we assume $t = t_0$. Define the ratio

$$p_{\mathrm{limit}}^{\mathrm{relative}} \equiv \frac{i_{\mathrm{limit}}^{\mathrm{photon}}}{i_{\mathrm{limit}}^{\mathrm{saturation}}} = \frac{g_{\mathrm{electr}}\kappa_{\mathrm{gray}}^2}{2v}(N-1). \tag{65}$$

We now reach our conclusion: If $p_{\mathrm{limit}}^{\mathrm{relative}} > 1$, then photon noise does *not* limit the useful intensity range in which Hadamard multiplexing offers an SNR advantage. Rather, this range is limited by saturation, as in Section 5. On the other hand, for $p_{\mathrm{limit}}^{\mathrm{relative}} < 1$, Hadamard multiplexing offers an SNR advantage for a portion $p_{\mathrm{limit}}^{\mathrm{relative}}$ of the single-source range.[12]

Let us look more closely at (63). We would benefit from a higher bound $i_{\mathrm{limit}}^{\mathrm{photon}}$ since it allows a larger absolute range of intensities to gain from Hadamard multiplexing. The bound $i_{\mathrm{limit}}^{\mathrm{photon}}$ monotonically increases with N. Its value corresponding to $N = 3$ doubles when N is large. Hence, a larger range of single-source intensities i_s can benefit from multiplexing, if *all* the available sources are involved, until saturation. This aspect is even clearer from (65): The *relative* range $p_{\mathrm{limit}}^{\mathrm{relative}}$ of intensities benefiting from multiplexing grows linearly with N. Hence, Hadamard multiplex is most advantageous when the lighting setup includes a large number of dim sources. Thus, for the projector-based system, the method works best with a large number of small segments, hence assisting high-directional resolution of lighting.

12. Some cameras have an antiblooming mechanism. It protects pixels from electron overflow of adjacent saturated pixels, hence keeping saturation local. In such a camera, the radiometric response is distorted in the top 10-20 percent of its output [21]. Effectively, this reduces v in (65), increasing $p_{\mathrm{limit}}^{\mathrm{relative}}$.

8.1.1 Examples of Specific Cameras

As examples of the consequences of (61) and (65), we estimate these expressions for several cameras on the current market. These estimates are based mostly on data derived from the manufacturers of the cameras or their detector chips. Some camera specifications quote g_{electr}, D and ρ_{read} in e^- units. Otherwise, these parameters should be inferred.[13] Often, the well-depth [19], denoted here as d_{well} is specified. This is the number of accumulated electrons at a pixel that corresponds to v. Hence, $g_{\mathrm{electr}} = d_{\mathrm{well}}/v$. The read noise ρ_{read} is derived as the ratio of d_{well} and the DR [10], [19], [22] specification.[14] We collected specification data of various cameras,[15] and, in Table 1, we standardize the units.

Equation (59) derives κ_{gray} based on ρ_{read}, ρ_{digit}, D, and exposure time t. The latter is set as the inverse of the camera's frames-per-second (FPS) specification. The derived parameters g_{electr} and κ_{gray} are then given in Table 1 as well. For the Dragonfly camera, the quoted parameters are not based on camera specifications (unavailable). Rather, they were directly calibrated by us in an experiment, as we describe in Section 10. These numbers can be used to predict the cameras' performance in various multiplexing schemes. In particular, we used this data in (60), (61), (64), and (65) to estimate $\sigma_{i,\mathrm{single}}^2$, $\mathrm{MSE}_{\hat{i},\mathrm{Hadamard}}$, and $p_{\mathrm{limit}}^{\mathrm{relative}}$ for $N = 255$ and $i_s \in [0, i_{\mathrm{limit}}^{\mathrm{saturation}}]$. The resulting MA is plotted in Fig. 6.

9 γ-Correction

Some cameras introduce a nonlinear transformation [19], namely, γ-correction to their images, prior to the final quantization of the data. This operation does *not* increase the DR of the detector [19], which is set by d_{well}, but actually compresses the DR. We now analyze this effect. A γ-correction transforms a linear measurement $a \in [0, v]$ to an acquired value

$$\tilde{a} = \chi a^{\gamma} \quad \text{where} \quad \chi = \tilde{v}/(v^{\gamma}), \tag{66}$$

while $\tilde{a} \in [0, \tilde{v}]$. Here, γ is a small number. Typically, $\gamma \approx 0.5$. The constant χ matches the input to output dynamic ranges. Transforming a 12-bit input to an 8-bit output using $\gamma = 0.5$ requires $\chi = 4$. This output quantization introduces an uncertainty q to $\tilde{a}$.

Linearizing $\tilde{a}$, we recover the image irradiance

$$\hat{a} = (\tilde{a}/\chi)^{(1/\gamma)}. \tag{67}$$

The noise variance of $\hat{a}$ is affected by that of a and by q. Approximating these noise contributions as independent and making use of (58), (66), and (67),

$$\sigma_{\hat{a}} = \sigma_a^2 + \left(\frac{\partial \hat{a}}{\partial \tilde{a}}q\right)^2 = \kappa_{\mathrm{gray}}^2 + \frac{a}{g_{\mathrm{electr}}} + \frac{a^{(2-2\gamma)}q^2}{\gamma^2\chi^2}. \tag{68}$$

13. The dark current D is sometimes given in mV/sec, and is converted to e^-/sec using a conversion factor of $\mu\mathrm{V}/e^-$ given in the specifications.

14. Specifications sometimes mistakenly denote the DR as SNR. The DR can be equivalent to SNR only in the absence of photon noise.

15. Typically, camera specifications correspond to the basic camera amplifier mode. A higher gain changes the noise and the saturation level, in e^- units.

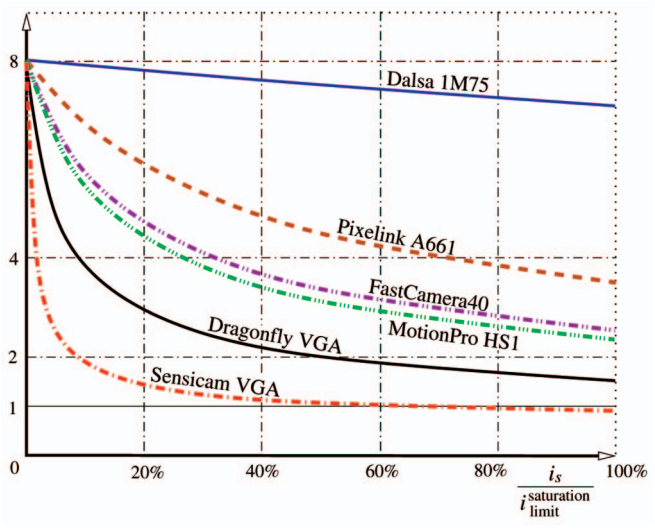

Fig. 6. MA for $N = 255$, accounting for the various noise componets. The plots are based on the data in Table 1. The intensity domain is scaled by each camera's saturation limit, under $N = 255$. The MA decreases with the single-source intensity i_s. For most of the plotted cameras, multiplexing enhances the SNR for all usable intensities. For the Sensicam, this multiplexing is counterproductive $(\text{MA} < 1)$ for the upper 34 percent of its useable intensities.

For a uniform quantizer, the noise STD is $q = \sqrt{1/12}$. The case of $\gamma = 0.5$ and $\chi = 4$ described above yields

$$\sigma^2_{\hat{a}} = \kappa^2_{\text{gray}} + a\left(\frac{1}{g_{\text{electr}}} + \frac{1}{48}\right). \tag{69}$$

Compare this result with (58). The effect of γ-correction is similar to that of photon statistics: noise is amplified with the detected image irradiance, undermining the MA. This problem is exacerbated when using a smaller γ, e.g., $\gamma = 0.45$ in conjunction to linearization by $1/\gamma \approx 2.2$ in (68). To conclude, for successful multiplexing and demultiplexing, it is advisable to avoid γ-correction camera modes.

10 Experiments

We used a setup as described in Section 4. The projector created $n = 255$ patterns on a wall, each containing 255 segments. The wall then illuminated the objects. In addition, we also acquired images under the corresponding individual sources (segments), using the same setup parameters (projector brightness, exposure time, lens aperture, and camera gain). The objects were also imaged when a "dark pattern" was projected. This "darkness frame" was subtracted from all frames acquired under active lighting. Based on the multiplexed raw frames, single-source images are decoded (demultiplexed).

First, we performed experiments using an analog Sony NTSC monochrome camera having a linear radiometric response. The projected patterns were as shown in Fig. 5. Two of the images acquired under multiplexed illumination are displayed in Fig. 7a. Two of the decoded images are displayed in Fig. 7b. The corresponding images taken under a single-source are displayed in Fig. 7c. The single-source

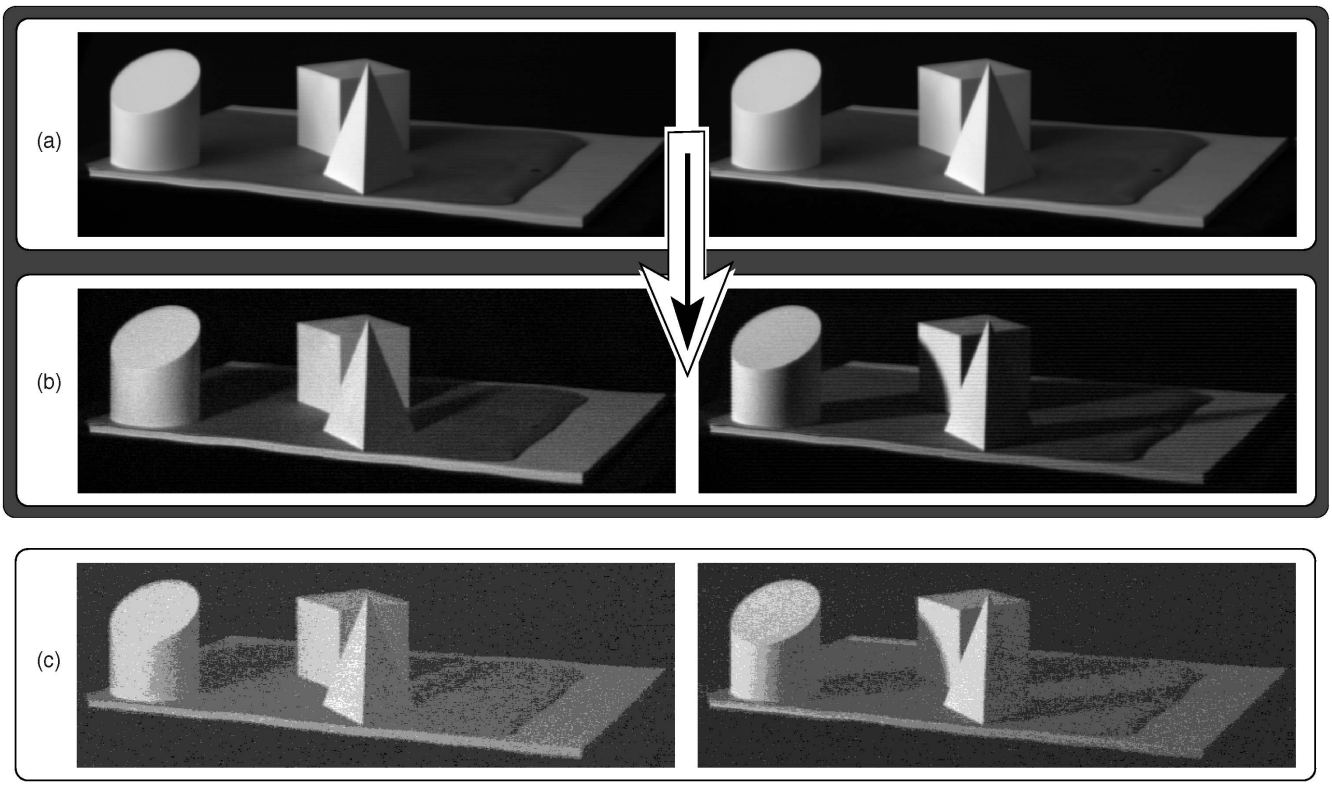

Fig. 7. Experimental results. All images are contrast stretched for display purposes. (a) Frames are acquired with multiplexed illumination. (b) Decoded images. (c) Corresponding images acquired by single-source illumination. The single-source images have a significantly lower SNR than their corresponding decoded images and low gray-level information.

Fig. 8. Experimental results of imaging shiny objects.

images are very dark, and are contrast stretched for display purposes. The correspondence of the decoded images to the single-source images is seen in the cast shadows. Yet, the decoded images have a much better quality. Fig. 8 shows results of an additional experiment done with specular objects. To uncover the object details, we saturated the specular highlights in the display. Yet, none of the raw captured images has saturated pixels. We have also performed a color experiment (see [39]): In conjunction to multiplexing the illumination *direction*, we have also multiplexed the illumination *color*. We projected cyan, magenta, and yellow patterns, each capturing $\approx 2/3$ of the color bandwidth. After decoding the illumination direction and color, we obtain true-color images as if illuminated by a single white source.

The analog camera suffers from a high read noise ρ. We thus experimented with a digital PointGrey Dragonfly camera. It is dominated by photon noise across much of its DR. We verified this in a dedicated experiment. A stable DC halogen light source illuminated a target having a wide range of reflectance values, spanning the camera DR. Then, a sequence of 100 frames was taken with $t = 28$ msec, an amplifier setting of 3.5 dB, and no γ. Based on the sequence, the statistics of each pixel value were derived. Fig. 9 plots the average variance of a pixel value, as a function of its quantized mean gray level. Near darkness, the noise STD is ≈ 0.4 gray levels, which is about twice the quantization noise. This noise raises to ≈ 2 gray levels at very bright pixels (saturation was avoided). This plot fits the noise model in (58) very well, with parameters $\kappa_{\text{gray}} = 0.4$ and $g_{\text{electr}} = 53e^{-}$.

Before using this camera under illumination, we ensured that fluctuations of light power are effectively eliminated. To do so, we studied *beforehand* the temporal behavior of the bulb of our Epson EMP-7800 projector, using a light meter (Newport 1815c) connected to an oscilloscope. We found that the radiance of our projector exhibits a *periodic* variation[16] having a period of 7 msec. Other than this, the measured bulb power appeared stable for long periods. Hence, if the exposure time is a multiple of 7 msec, the temporal variations of the projector should be eliminated. Indeed, we use $t = 28$ msec.

The experiment performed with the Dragonfly is similar to the one done with the analog camera. Here, $i_s \in [0, 97]$ gray levels, where $v = 255$. The brightest pixels are specular

16. This may originate from the bulb's electric power supply in the projector.

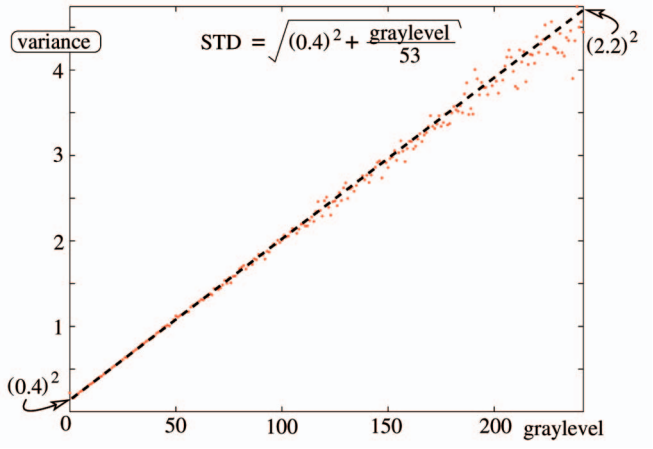

Fig. 9. The noise variance of the Dragonfly camera increases with the measured signal, as modeled by (58) and (60). For the settings we used for amplifier gain and exposure time, the noise STD is ≈ 0.4 gray levels at darkness, increasing to ≈ 2 gray levels near saturation.

points at nearly half of the image DR. At multiplexed illumination, $a_m \in [3, 215]$. Saturation does not occur despite a 128-fold increase of illumination intensity, as the brightness of specular points only doubles, as expected from Section 6. Moreover, the darkest pixels yield a readable $a_m \geq 3$ value, despite being nulled (read as $i_s = 0$) by quantization under single-source illumination. This enables subsequent observation of details. The results are shown in Fig. 10. For a fixed camera gain (3.5 dB), multiplexing is superior across the image, particularly in specular objects, which are too dark with single source imaging. For comparison, we also increased the camera amplifier gain to a setting of 21.6 dB. This is a linear amplification factor of $\times 8.04$, relative to the 3.5 dB setting, matching the $\text{MA} = 8.02$ gained by multiplexing $n = 255$ sources. This amplification improves single-source imaging in most of the FOV, but the quality is inferior to that obtained by light multiplexing with a 3.5 dB camera gain. Moreover, in 21.6 dB gain, the highlights are saturated (erroneous readout value of 255), contrary to the multiplexed images.

The quality improvement obtained in these experiments may not apply to other cameras. If a high-bit camera has a low read noise as the Sensicam (See Table 1 and Fig. 6), then multiplexing may not be as beneficial to the SNR. On the other hand, fast cameras as the MotionPro may benefit from multiplexing even more than we have since they typically have a higher read noise (despite their higher bit quantizers).

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 enhances realism in image-based rendering of specular reflections. The cup shown in Fig. 8 is rendered in Fig. 11 as if it is illuminated by a window, out of which is a natural outdoor scene. The scene is specularly reflected from various places on the cup.

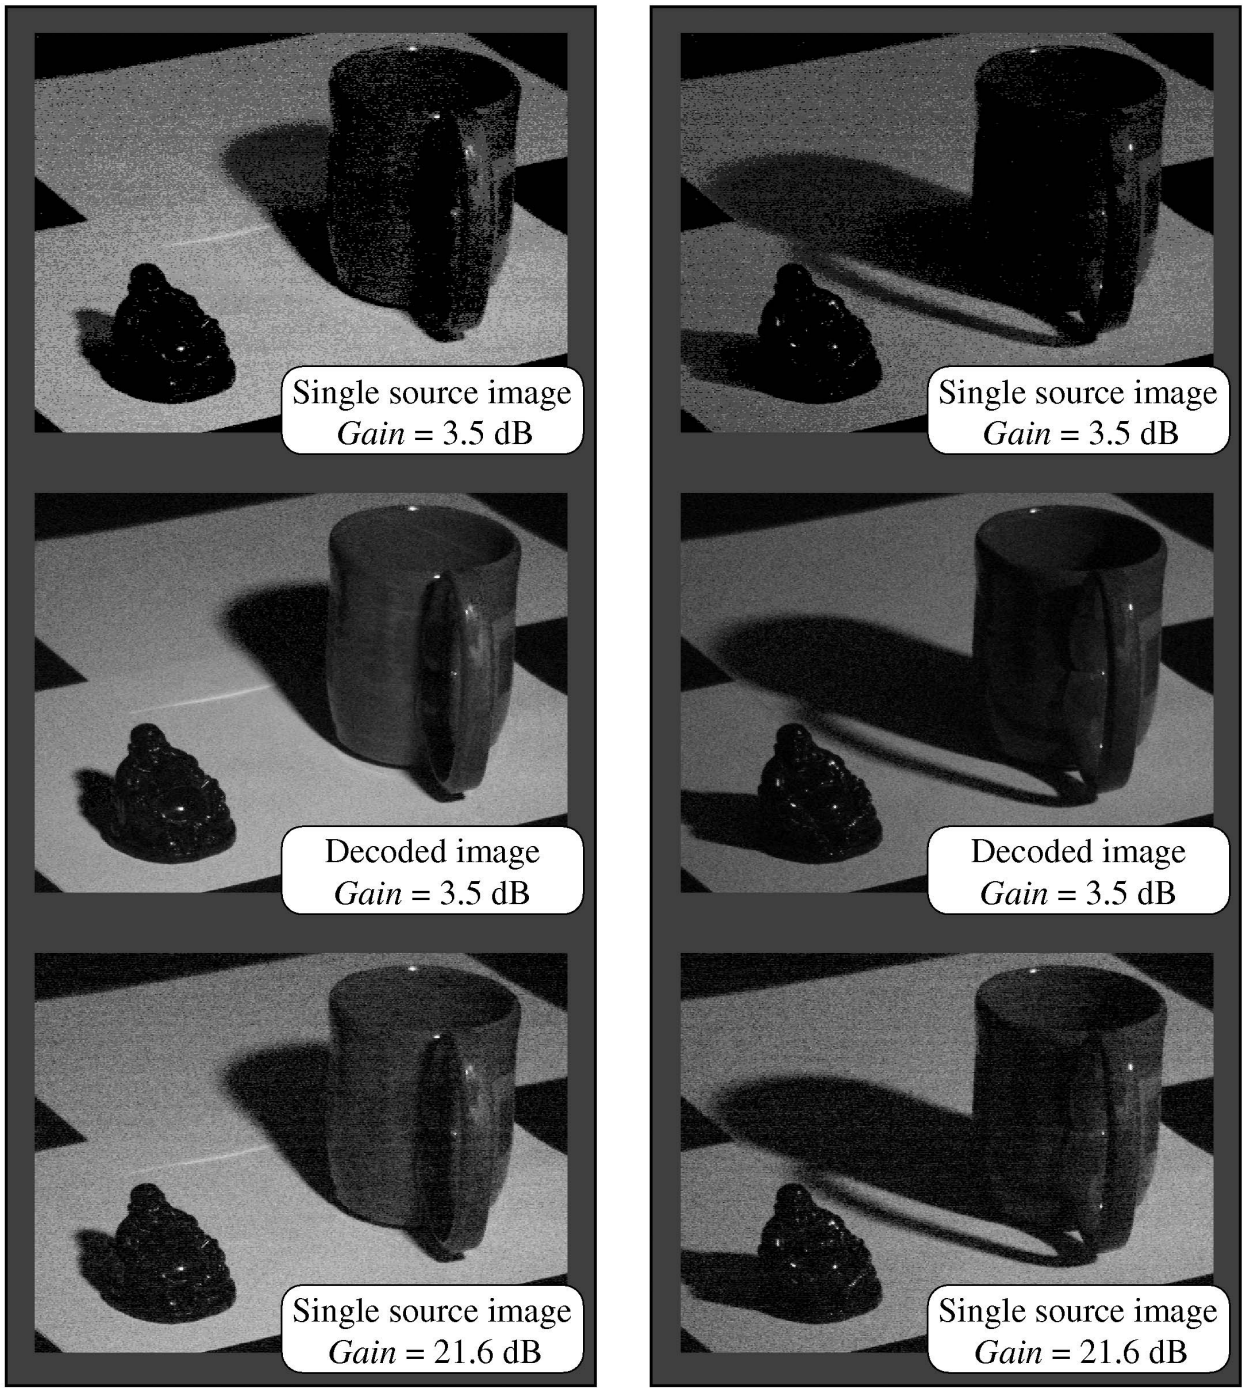

Fig. 10. Results of an experiment using a DragonFly camera and $n = 255$. The inherent noise characteristics of this camera significantly depend on the measured signal, as seen in Fig. 9. Still, (b) the multiplexing approach enhances the results relative to (a) single source illumination. (c) The results are better than increasing the internal camera gain by the MA.

11 Operation Guidelines

To recap, we summarize guidelines for Hadamard illumination multiplexing.

1. Set the exposure time t and camera gain. Avoid γ-correction in the camera, if possible.
2. Ensure that the radiance of the illuminants is stable, when integrated over t.
3. Take a "darkness image" of the scene, at the settings of Step 1. Use this image later, to compensate for stray light and dark current bias.
4. Calibrate the noise specification σ_a^2 of the camera, at the settings of Step 1.
5. Define the number of sources n you wish to operate.
6. Take a preliminary photograph with multiplexed sources.
7. If Step 6 reveals saturation in diffuse areas, decrease the number of multiplexed sources to N. Repeat Steps 6 and 7 until saturation is avoided in diffuse areas. Do *not* decrease the exposure time, aperture or source brightness for this purpose.

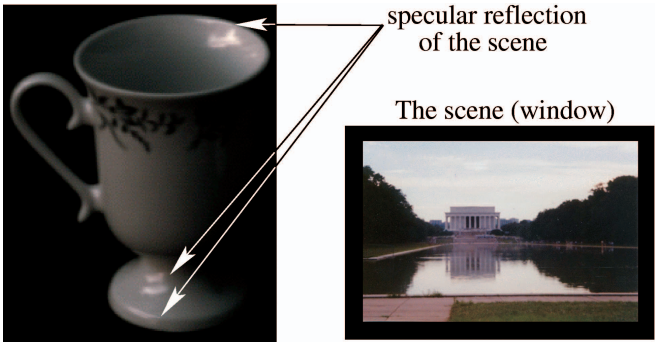

Fig. 11. Rendering the object based on the decoded images, as if illuminated by a wide angle scene. Specular reflections are densely rendered thanks to the large number of sources.

8. Based on N, determine $i_{\text{limit}}^{\text{saturation}}$ by (64) and $\text{MSE}_{\hat{i},\text{Hadamard}}$ by (61).
9. Based on N and the camera noise parameters,[17] derive $i_{\text{limit}}^{\text{photon}}$ using (63).
10. Use multiplexed illumination with all N sources.[18] If $i_{\text{limit}}^{\text{photon}} > i_{\text{limit}}^{\text{saturation}}$, the process will increase the SNR for all pixels. Otherwise, it would benefit $i_s < i_{\text{limit}}^{\text{photon}}$, while it would lead to some loss of SNR in brighter objects. The latter are better imaged using $N = 1$.

12 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 beneficial to these techniques. There are still open questions for further research. In particular, part of our limitation analysis assumed that for diffuse objects the value of i_s is insensitive to s. This assumption does not hold when the illumination direction range is very wide due to foreshortening of object irradiance [45]. It is worth extending our analysis of limitations and tradeoffs to such cases.

This paper analyzed the consequences to Hadamard multiplexing by effects as illumination instabilities, unequal sources (channels), and photon noise. This analysis is not limited to illumination, but is useful to other fields which employ multiplexing, such as spectral sensing [9], [12], [15], [44], [46], X-ray and γ-ray imaging (e.g., in medical imaging). Moreover, these effects may be accounted for earlier in multiplexing design. This may result in new multiplexing codes. We currently pursue this direction. An additional extension is to define objective functions other than SNR, that multiplexing needs to optimize or address. It is interesting to note that light demultiplexing can be based on images taken outdoors under natural illumination [32].

17. We stress that the linear model in (58) should apply to scientific grade cameras. If the electronic circuity creates noise σ_a^2 that is not well modeled by (58), this step should be replaced: calculate whether (62) holds for each gray-level $i_s < i_{\text{limit}}^{\text{saturation}}$.

18. This simple process applies to cameras whose noise σ_a^2 is well modeled by (58). If the noise behaves significantly different, then N may need to be determined by a dedicated analysis.

Appendix

To make the paper self-contained, we describe a method for creating an $\mathbf{S}$ matrix. It is derived from a Hadamard matrix $\mathbf{H}$. There are several recipes [15] for creating such matrices. A simple method involves the following steps:

1. Define the 1×1 matrix $\mathbf{H}_1 = 1$. Then, recursively, define the $2l \times 2l$ matrix
$$\mathbf{H}_{2l} = \begin{bmatrix} \mathbf{H}_l & \mathbf{H}_l \\ \mathbf{H}_l & -\mathbf{H}_l \end{bmatrix}. \tag{70}$$
2. Define the number of rows of $\mathbf{H}$ as $n+1$. In Step 1, all the elements of the first row and column of $\mathbf{H}_{n+1}$ are 1s. Cropping this row and column, we are left with a matrix of size $n \times n$. Similarly to the notation in [15], we denote this matrix as $\mathbf{G}_n$, where
$$\mathbf{H}_{n+1} = \begin{bmatrix} 1 & \mathbf{1}_{1\times n} \\ \mathbf{1}_{n\times 1} & \mathbf{G}_n \end{bmatrix} \tag{71}$$
and $\mathbf{1}$ is a matrix whose elements are all 1s and whose size is given by its subscript.
3. The $n \times n$ matrix $\mathbf{S}$ is given by $\mathbf{S} = (\mathbf{1}_{n\times n} - \mathbf{G}_n)/2$.

This procedure creates $\mathbf{H}$ matrices for which $n+1 = 2^Q$, where Q is an integer [20]. Another recipe creates $\mathbf{H}$ matrices for which $n+1$ is a multiple of 4. This can be easily obtained by the Matlab software, using the command *hadamard*. In any case, as proven in [15], we may always manipulate $\mathbf{H}$ to have all the elements of its first row and column be 1s. Therefore, we may proceed to Steps 2 and 3 in the procedure detailed above, to create $\mathbf{S}$. This is the method we have implemented. There are methods [15] that create *cyclic* $\mathbf{S}$ matrices. There, each row is a cyclic permutation of its proceeding one, as in (10). Therefore, $\mathbf{S}$ in not unique. Rather, for a given n, there is a class of $\mathbf{S}$ matrices, all of which are equivalent with respect to the analysis in this paper.

Acknowledgments

The authors are grateful to the anonymous reviewers for their deep thought and very useful comments. They would like to thank Harish Peri, Einav Namer, Eyal Gordon, and Netanel Ratner for their help with experiments. Yoav Schechner is a Landau Fellow—supported by the Taub Foundation and an Alon Fellow. The work was supported by the US-Israel Binational Science Foundation (BSF). It was conducted in the Ollendorff Minerva Center. Minerva is funded through the BMBF. This work was also supported by US National Science Foundation grants ITR IIS-00-85864, PECASE IIS-9703134, EIA-02-24431, IIS-03-08185 and KDI-99-80058.

References

[1] H.H. Barrett and W. Swindell, *Radiological Imaging*, vol. 1, pp. 82-98, 285-290. Academic Press, 1981.
[2] R. Basri and D. Jacobs, "Lambertian Reflectance and Linear Subspaces," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 25, pp. 218-233, 2003.
[3] O.G. Cula, K.J. Dana, D.K. Pai, and D. Wang, "Polarization Multiplexing for Bidirectional Imaging," *Proc. Computer Vision and Pattern Recognition*, vol. 2, pp. 1116-1123, 2005.

[4] K.J. Dana and J. Wang, "Device for Convenient Measurement of Spatially Varying Bidirectional Reflectance," *J. Optical Soc. Am. A,* vol. 21, pp. 1-12, 2004.
[5] P. Debevec, T. Hawkins, C. Tchou, H.P. Duiker, W. Sarokin, and M. Sagar, "Acquiring the Reflectance Field of a Human Face," *Proc. ACM SIGGRAPH,* pp. 145-156, 2000.
[6] P.E. Debevec and J. Malik, "Recovering High Dynamic Range Radiance Maps from Photographs," *Proc. ACM SIGGRAPH,* pp. 369-378, 1997.
[7] R. Epstein, P.W. Hallinan, and A.L. Yuille, "5 ± 2 Eigenimages Suffice: An Empirical Investigation of Low-Dimensional Lighting Models," *Proc. Physics-Based Modeling in Computer Vision,* pp. 108-116, 1995.
[8] H. Farid and E.H. Adelson, "Separating Reflections and Lighting Using Independent Components Analysis," *Proc. Computer Vision and Pattern Recognition,* vol. I, pp. 262-267, 1999.
[9] W.G. Fateley, R.M. Hammaker, R.A. DeVerse, R.R. Coifman, and F.B. Geshwind, "The Other Spectroscopy: Demonstration of a New De-Dispersion Imaging Spectrograph," *Vibrational Spectroscopy,* vol. 29, pp. 163-170, 2002.
[10] T.J. Fellers and M.W. Davidson, "CCD Noise Sources and Signal-to-Noise Ratio," *Optical Microscopy Primer,* (Molecular Expressions, Florida State Univ.), http://micro.magnet.fsu.edu/primer/, 2004.
[11] D. Goldman, B. Curless, A. Hertzmann, and S.M. Seitz, "Shape and Spatially-Varying BRDFs from Photometric Stereo," *Proc. IEEE Int'l Conf. Computer Vision,* pp. 341-348, 2005.
[12] Q.S. Hanley, P.J. Verveer, and T.M. Jovin, "Spectral Imaging in a Programmable Array Microscope by Hadamard Transform Fluorescence Spectroscopy," *Applied Spectroscopy,* vol. 53, pp. 1-10, 1999.
[13] Hamamatsu Photnics K.K., "Xenon Flash Lamps,"Catalog TLSX1008E04, (Electron Tube Center, Japan), 1998.
[14] Hamamatsu Photnics K.K., "Super-Quiet Xenon Lamps,"Catalog TLSX1002E06, (Electron Tube Center, Japan), 2000.
[15] M. Harwit and N.J.A. Sloane, *Hadamard Transform Optics.* Academic Press, 1979.
[16] M. Hatzitheodorou, "Shape from Shadows, A Hilbert Space Setting," *J. Complexity,* vol. 14, pp. 63-84, 1998.
[17] G. Healey and T.O. Binford, "Local Shape from Specularity," *Proc. Int'l Conf. Computer Vision,* pp. 151-160, 1987.
[18] B.P.K. Horn, *Robot Vision,* chapter 10. MIT Press, 1986.
[19] S. Ioué and K.R. Spring, *Video Microscopy,* second ed., chapters 6, 7, and 8. Plenum Press, 1997.
[20] A.K. Jain, *Fundamentals of Digital Image Processing,* pp. 155-159. Prentice-Hall, 1989.
[21] Eastman Kodak Co., "Interline Image Sensor Photodiode Charge Capacity and Antiblooming," Application Note DS 01-001, Revision 2, 2000.
[22] Eastman Kodak Co., "CCD Image Sensor Noise Sources,"Application Note MTD/PS-0233, Revision 2, 2003.
[23] M. Koudelka, P. Belhumeur, S. Magda, and D. Kriegman, "Image-Based Modeling and Rendering of Surfaces with Arbitrary BRDF's," *Proc. IEEE Computer Vision and Pattern Recognition,* pp. 568-575, 2001.
[24] K.C. Lee, J. Ho, and D. Kriegman, "Acquiring Linear Subspaces for Face Regognition under Variable Lighting," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 27, pp. 684-698, 2005.
[25] H. Lensch, J. Kautz, M. Gosele, W. Heidrich, and H. Seidel, "Image-Based Reconstruction of Spatial Appearance and Geometric Detail," *ACM Trans. Graphics,* vol. 22, pp. 234-257, 2003.
[26] M. Levoy, B. Chen, V. Vaish, M. Horowitz, I. McDowall, and M. Bolas, "Synthetic Aperture Confocal Imaging," *ACM Trans. Graphics,* vol. 23, pp. 825-834, 2004.
[27] Q.T. Luong, P. Fua, and Y. Leclerc, "Recovery of Reflectances and Varying Illuminants from Multiple Views," *Proc. European Conf. Computer Vision,* pp. 163-179, 2002.
[28] T. Malzbender, D. Gelb, and H. Wolters, "Polynomial Texture Maps," *Proc. ACM SIGGRAPH,* pp. 519-528, 2001.
[29] S. Mann and R.W. Picard, "On Being "Undigital" with Digital Cameras: Extending Dynamic Range by Combining Differently Exposed Pictures," *Proc. IS&T Ann. Conf.,* pp. 422-428, 1995.
[30] S.R. Marschner, S.H. Westin, E.P.F. Lafortune, and K.E. Torrance, "Image-Based Bidirectional Reflectance Distribution Function Measurement," *Applied Optics,* vol. 39, pp. 2592-2600, 2000.
[31] V. Masselus, P. Dutré, and F. Anrys, "The Free Form Light Stage," *Proc. ACM Eurographics Workshop Rendering,* pp. 247-256, 2002.
[32] W. Matusik, M. Loper, and H. Pfister, "Progressively-Refined Reflectance Functions from Natural Illumination," *Proc. Eurographics Symp. Rendering,* 2004.
[33] W. Matusik, H. Pfister, A. Ngan, P. Beardsley, R. Ziegler, and L. McMillan, "Image-Based 3D Photography Using Opacity Hulls," *ACM Trans. Graphics,* vol. 21, pp. 427-437, 2002.
[34] Y. Moses, S. Ullman, and S. Edelman, "Generalization to Novel Images in Upright and Inverted Faces," *Perception,* vol. 25, pp. 443-461, 1996.
[35] G. Nitzsche and R. Riesenberg, "Noise, Fluctuation, and Hadamard-Transform Spectrometry," *Proc. SPIE,* vol. 5111, pp. 273-282, 2003.
[36] M. Osadchy and D. Keren, "Efficient Detection under Varying Illumination Conditions and Image Plane Rotations," *Computer Vision & Image Understanding,* vol. 93, pp. 245-259, 2004.
[37] R. Ramamoorthi and P. Hanrahan, "Frequency Space Environment Map Rendering," *ACM Trans. Graphics,* vol. 21, pp. 517-526, 2002.
[38] I. Sato, T. Okabe, Y. Sato, and K. Ikeuchi, "Using Extended Light Sources for Modeling Object Apperance under Varying Illumination," *Proc. Int'l Conf. Computer Vision,* vol. I, pp. 325-332, 2005.
[39] Y.Y. Schechner, S.K. Nayar, and P.N. Belhumeur, "A Theory of Multiplexed Illumination," *Proc. Int'l Conf. Computer Vision,* vol. 2, pp. 808-815, 2003.
[40] Y.Y. Schechner, S.K. Nayar, P.N. Belhumeur, and H.S. Peri, "Imaging in Multiplexed Illumination," *Proc. SPIE,* vol. 5529, pp. 198-205, 2004.
[41] P. Sen, B. Chen, G. Garg, S.R. Marschner, M. Horowitz, M. Levoy, and H.P.A. Lensch, "Dual Photography," *ACM Trans. Graphics,* vol. 24, pp. 745-755, 2005.
[42] A. Shashua, "On Photometric Issues in 3D Visual Recognition from a Single 2D Image," *Int'l J. Computer Vision,* vol. 21, pp. 99-122, 1997.
[43] T. Sim, S. Baker, and M. Bsat, "The CMU Pose, Illumination, and Expression Database," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 25, pp. 1615-1618, 2003.
[44] J.F. Turner II and P.J. Treado, "Near-Infrared Acousto-Optic Tunable Filter Hadamard Transform Spectroscopy," *Applied Spectroscopy,* vol. 50, no. 2, pp. 277-284, 1996.
[45] A. Wenger, A. Gardner, C. Tchou, J. Unger, T. Hawkins, and P. Debevec, "Performance Relighting and Reflectance Transformation with Time-Multiplexed Illumination," *ACM Trans. Graphics,* vol. 24, pp. 756-764, 2005.
[46] A. Wuttig, R. Riesenberg, and G. Nitzsche, "Subpixel Analysis of a Double Array Grating Spectrometer," *Proc. SPIE,* vol. 4480, pp. 334-344, 2002.
[47] A. Yuille, J.M. Coughlan, and S. Konishi, "The KGBR Viewpoint-Lighting Ambiguity," *J. Optical Soc. Am. A,* vol. 20, pp. 24-31, 2003.

Yoav Y. Schechner received the BA and MSc degrees in physics and PhD in electrical engineering from the Technion-Israel Institute of Technology in 1990, 1996, and 2000, respectively. From 2000 to 2002, he was a research scientist in the Computer Science Department at Columbia University. Since 2002, he has been a faculty member in the Department of Electrical Engineering at the Technion, where he heads the Hybrid Imaging Lab. His research is focused on computer vision, the use of optics and physics in imaging and computer vision, and on multimodal sensing. He was the recipient of the Wolf Foundation Award for Graduate Students in 1994, the Guttwirth Special Distinction Fellowship in 1995, the Israeli Ministry of Science (Eshkol) Distinction Fellowship and the Ollendorff Award in 1998, the Scwartz Foundation Award in 1999, and the Morin Fellowship in 2000-2002. He is now a Landau Fellow—supported by the Taub Foundation, and an Alon Fellow. He has received the Klein Research Award in 2006. He is a member of the IEEE and the IEEE Computer Society.

Shree K. Nayar received the PhD degree in electrical and Computer engineering from the Robotics Institute at Carnegie Mellon University in 1990. He is the T.C. Chang Professor of Computer Science at Columbia University. He currently heads the Columbia Automated Vision Environment (CAVE), which is dedicated to the development of advanced computer vision systems. His research is focused on three areas: the creation of cameras that produce new forms of visual information, the modeling of the interaction of light with materials, and the design of algorithms that recognize objects from images. His work is motivated by applications in the fields of computer graphics, human-machine interfaces, and robotics. Dr. Nayar has authored and coauthored papers that have received the Best Paper Award at the 2004 CVPR Conference held in Washington, DC, the Best Paper Honorable Mention Award at the 2000 IEEE CVPR Conference held in Hilton Head, the David Marr Prize at the 1995 ICCV held in Boston, the Siemens Outstanding Paper Award at the 1994 IEEE CVPR Conference held in Seattle, the 1994 Annual Pattern Recognition Award from the Pattern Recognition Society, the Best Industry Related Paper Award at the 1994 ICPR held in Jerusalem, and the David Marr Prize at the 1990 ICCV held in Osaka. He was the recipient of the Columbia Great Teacher Award in 2006, the Excellence in Engineering Teaching Award from the Keck Foundation in 1995, the NTT Distinguished Scientific Achievement Award from NTT Corporation, Japan, in 1994, the National Young Investigator Award from the US National Science Foundation in 1993, and the David and Lucile Packard Fellowship for Science and Engineering in 1992. He holds several US and international patents for inventions related to computer vision and robotics. He is a member of the IEEE.

Peter N. Belhumeur received the ScB degree in information sciences from Brown University in 1985. He received the PhD degree in engineering sciences from Harvard University under the direction of David Mumford in 1993. He was a postdoctoral fellow at the University of Cambridge's Isaac Newton Institute for Mathematical Sciences in 1994. He was made assistant, associate, and professor of electrical engineering at Yale University in 1994, 1998, and 2001, respectively. He joined Columbia University in 2002 where he is currently a professor in the Department of Computer Science and he is the director of the Laboratory for the Study of Visual Appearance. His main research focus is on illumination, reflectance, and shape and their relation to visual appearance. Within these areas, he concentrates on two subproblems: the representation and recognition of objects under variable illumination and the estimation of the geometry of objects from low-level cues like image brightness, binocular stereopsis, and motion. Applications include face and object recognition, image-based rendering, computer graphics, content-based image and video compression, and human computer interfaces. He is a recipient of the Presidential Early Career Award for Scientists and Engineers (PECASE) and the US National Science Foundation Career Award. He won both the Siemens Best Paper Award at the 1996 IEEE Conference on Computer Vision and Pattern Recognition and the Olympus Prize at the 1998 European Conference of Computer Vision. His research in computational vision is funded by NSF (PECASE), NIH, NSF KDI-LIS, NSF ITR, ARO, and DARPA.