

Illumination Multiplexing within Fundamental Limits

Netanel Ratner Yoav Y. Schechner
Department of Electrical Engineering
Technion - Israel Institute of Technology
Haifa 32000, ISRAEL

ratner@tx.technion.ac.il

yoav@ee.technion.ac.il

Abstract

Taking a sequence of photographs using multiple illumination sources or settings is central to many computer vision and graphics problems. A growing number of recent methods use multiple sources rather than single point sources in each frame of the sequence. Potential benefits include increased signal-to-noise ratio and accommodation of scene dynamic range. However, existing multiplexing schemes, including Hadamard-based codes, are inhibited by fundamental limits set by Poisson distributed photon noise and by sensor saturation. The prior schemes may actually be counterproductive due to these effects. We derive multiplexing codes that are optimal under these fundamental effects. Thus, the novel codes generalize the prior schemes and have a much broader applicability. Our approach is based on formulating the problem as a constrained optimization. We further suggest an algorithm to solve this optimization problem. The superiority and effectiveness of the method is demonstrated in experiments involving object illumination.

1. Illumination Multiplexing

In computer vision research and image-based rendering, objects or people are often acquired under variable lighting directions [4, 6, 13, 14, 17, 18, 19, 24, 25, 27, 31, 32, 37]. Such images are then used for object recognition and identification [4, 13, 21, 23, 31, 32], rendering [7, 13, 19, 25], shape estimation [9, 11, 36, 37] and analysis of specularities, shadows and occlusions [26]. Traditionally, such images were taken by moving a light source around the object, or by sequential operation of individual sources in a constellation. However, recently, there is a growing interest in illumination that is not based on single point sources. Rather, it is based on a sequence of images, in each of which lighting may simultaneously arrive from several directions or sources [6, 13, 15, 20, 27, 28, 29, 30, 34]. Some of the benefits include significant improvement in signal to noise ratio (SNR) [29] (See for example Fig. 1), significant reduc-

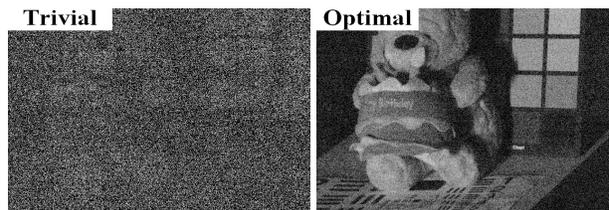


Figure 1. [Left] An image taken under a single light source. [Right] An image of the same scene, decoded from images illuminated by 57 multiplexed sources. It is decoded as if illuminated by the same single source. The multiplexing code is optimal.

tion of dynamic range problems in the presence of saturated pixels, and convenience when photographing people [34]. Other potential advantages are efficiency of the acquisition process [28, 20], and image enhancement by simultaneous use of flashes and ambient lighting [1, 26].

The question is, given all the possibilities of simultaneous operation of sources, what is the optimal way to multiplex the sources in each frame. Ref. [28] suggested that Hadamard-based codes should be used. However, its analysis did not account for a very important problem: image noise depends on the image irradiance itself, which may make Hadamard multiplexing counterproductive, as was later experienced by [34]. Fundamentally, this is due to *photon noise*. It exists in images no matter the quality of the camera, as it stems from the quantum mechanical nature of light. Moreover, no prior study accounted for saturation when seeking optimal lighting. This is despite the acknowledgment that saturation and scene dynamic range are important aspects when using multiple sources [29, 34].

This paper directly seeks multiplexing codes that are optimal under the fundamental limitations of photon noise and saturation, in addition to camera readout noise. This problem and its solution have implications much broader than computer vision and graphics. The reason is that multiplexing of radiation sources is used in many sensing modalities, such as X-ray imaging [10, 33], spectroscopy [10], coded-aperture imaging, and communication in fiber optics.

Hence, the approach presented here has wide applicability. It is based on a constrained optimization formulation. We also describe an algorithm for solving this problem. The resulting novel codes are superior to prior multiplexing codes. We demonstrate this in experiments of object lighting.

2. Theoretical Background

2.1. Multiplexing

Consider a setup where N light sources illuminate an object from various directions. Let $\mathbf{i} = (i_1, i_2, \dots, i_N)^t$ be a set of intensity values of a certain pixel, where each value corresponds to illumination by any *individual* light source in this setup. Here t denotes transposition.

In general, several light sources can be turned on at a time (multiplexing). Define an $N \times N$ multiplex matrix \mathbf{W} , often referred to as a *multiplexing code*. Each element of its m th row represents the power of the corresponding illumination source in the m th measurement. The power is measured relative to its maximum value, where 0 states that the source is completely off and 1 indicates a fully activated source. The measurements acquired at each pixel are denoted by the vector $\mathbf{a} = (a_1, a_2, \dots, a_N)^t$. It is given by

$$\mathbf{a} = \mathbf{W}\mathbf{i} + v \quad , \quad (1)$$

where v is the measurement noise. Any bias to this noise is assumed to be compensated for. The noise v is assumed to be uncorrelated in different pixels, with variance of σ_a^2 .

Once images have been acquired under multiplexed illumination, they can be *demultiplexed* computationally, to derive estimates for the pixel values under single-source illumination $\hat{\mathbf{i}}$. The best linear estimator in the sense of mean square error (MSE) for the single source images is

$$\hat{\mathbf{i}} = \mathbf{W}^{-1}\mathbf{a} \quad . \quad (2)$$

The MSE of this estimator [10, 29] is

$$\text{MSE}_{\hat{\mathbf{i}}} = \frac{\sigma_a^2}{N} \text{trace} \left[(\mathbf{W}^t \mathbf{W})^{-1} \right] \quad . \quad (3)$$

This is the expected noise variance of the recovered images. The lower it is, the better the SNR. The SNR is defined as the ratio between the expected $\hat{\mathbf{i}}$ and $\sqrt{\text{MSE}_{\hat{\mathbf{i}}}}$. Without multiplexing, \mathbf{W} is the identity matrix (trivial sensing: only a single source is on at a time). The improved SNR by multiplexing, relative to the SNR without multiplexing

$$G = \text{SNR}_{\text{Multiplexed}} / \text{SNR}_{\text{Single}} \quad (4)$$

is the *multiplex gain*.

2.2. Noise Mechanisms

To analyze the effect of multiplexing, we should first understand the sources of image noise. In this section we briefly review the *affine noise model*. It exists in high

grade detectors, which have a linear radiometric response. The noise can be divided into two components, signal-dependent and signal-independent. Regardless of the photon flux, signal-independent noise is created by dark current [12, 16, 29], amplifier noise and the quantizer in the camera circuitry [16]. Denote the graylevel variance of the signal-independent noise by κ_{gray}^2 .

Fundamental signal-dependent noise is related to two random effects. The photon flux and the uncertainty of the electron-photon conversion process which occurs in the detector. Overall, the random number $n_{\text{electr}}^{\text{photo}}$ of photo-generated electrons is Poisson distributed [2, 3, 12]. In this distribution, the variance of $n_{\text{electr}}^{\text{photo}}$ is

$$\text{VAR}(n_{\text{electr}}^{\text{photo}}) = \mathcal{E}(n_{\text{electr}}^{\text{photo}}) \quad , \quad (5)$$

where \mathcal{E} denotes expectation. This variance linearly increases with the measured electric signal $n_{\text{electr}}^{\text{photo}}$. This is *photon noise*. The number of detected electrons $n_{\text{electr}}^{\text{photo}}$ is proportional to the gray-level of the acquired pixel value a

$$a = n_{\text{electr}}^{\text{photo}} / Q_{\text{electr}} \quad . \quad (6)$$

Here Q_{electr} is the number of photo-generated electrons required to change a unit gray-level. Typically $Q_{\text{electr}} \gg 1$. Combining Eqs. (5,6) yields a variance in gray levels

$$\mathcal{E}(n_{\text{electr}}^{\text{photo}}) / Q_{\text{electr}}^2 = a / Q_{\text{electr}} \quad . \quad (7)$$

Compounded with signal-independent noise, the total noise variance of the measured gray level [12, 29] is

$$\sigma_a^2 = \kappa_{\text{gray}}^2 + a / Q_{\text{electr}} \quad . \quad (8)$$

Now, consider a diffuse object and sources that illuminate the object from similar directions. In this case, each light source yields a similar object radiance, hence, a similar level of noise. In each measurement, let C sources be activated, each at maximum power. We rephrase Eq. (8) as

$$\sigma_a^2 = \kappa_{\text{gray}}^2 + C\eta^2 \quad . \quad (9)$$

Here η^2 is the photon noise variance, induced by object irradiance from a *single* source turned on completely. Eq. (9) is an affine function of the number of active sources C .

As an example, Fig. 2 plots the average noise variance in raw images acquired by a PtGrey Dragonfly camera. In each measurement, C light sources were activated. The dynamic range of the 16-bit raw data was $a \in [0, 65535]$ graylevels, while $\sigma_a \in [70, 220]$ graylevels. Fitting a straight line to this plot yields κ_{gray}^2 and η^2 .

2.3. Photon Noise and Multiplexing

A well known multiplexing code is based on Hadamard codes. Its multiplex matrix is known as the \mathbf{S} -matrix [8, 10, 22, 29, 33, 34]. It was used in Refs. [29, 34] to

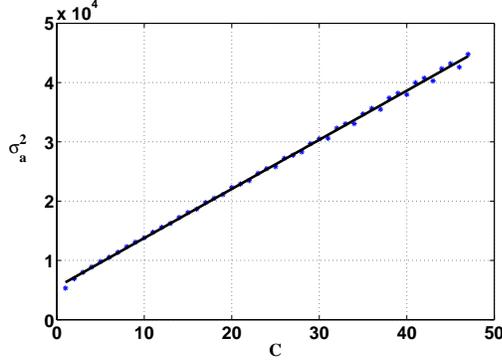


Figure 2. Noise calibration for $N = 47$. In our system, the noise variance linearly increases with the number of activated sources C , in agreement with the affine noise model (9).

multiplex illumination sources. Here $C = (N + 1)/2$. The MSE obtained using this code is

$$\text{MSE}_{i,\text{Had}} = \frac{4N}{(N+1)^2} \sigma_a^2. \quad (10)$$

Using Eq. (9) with $C = (N + 1)/2$, Eq. (10) yields

$$\text{MSE}_{i,\text{Had}} = \frac{4N}{(N+1)^2} \kappa_{\text{gray}}^2 + \frac{2N}{N+1} \eta^2. \quad (11)$$

In the special case where the photon noise is negligible, *i.e.* $\kappa_{\text{gray}}^2 \gg C\eta^2$, Eq. (10) becomes:

$$\text{MSE}_{i,\text{Had}} = \frac{4N}{(N+1)^2} \kappa_{\text{gray}}^2 \quad (12)$$

and the corresponding SNR gain [29, 35] is

$$G_{\text{Had}} = \frac{\text{SNR}_{\text{Hadamard}}}{\text{SNR}_{\text{Single}}} = \frac{N+1}{2\sqrt{N}}. \quad (13)$$

Hence, in such a scenario, Hadamard multiplexing is highly beneficial. Ref. [10] shows that then, the \mathbf{S} -matrix is optimal, minimizing Eq. (3).

On the other hand, when photon noise dominates, then $C\eta^2 \gg \kappa_{\text{gray}}^2$. In this case, Eq. (11) indicates that the demultiplexed images $\{\hat{i}\}$ are *more* noisy than those obtained by simple single-source acquisition [10, 34]. The noise variance doubles by this process, if $N \gg 1$. The reason is that increasing the signal by multiplexing light sources increases the photon noise as well.

Ref. [35] looked into the problem of multiplexing under photon noise. It formulated a general expression for the multiplex gain under the affine model of Eq. (5):

$$G = G_0 \sqrt{\frac{1 + \chi^2}{1 + C\chi^2}}, \quad (14)$$

where

$$\chi = \eta / \kappa_{\text{gray}}. \quad (15)$$

Here,

$$G_0 = \sqrt{N / \text{trace} \left[(\mathbf{W}^t \mathbf{W})^{-1} \right]} \quad (16)$$

is the multiplex gain when photon noise is not considered. Hence, for a given characteristic χ of the noise, G in Eq. (14) is maximized by reducing C while increasing G_0 . Ref. [35] proposed multiplexing codes, which optimize G out of the set of cyclic binary matrices \mathbf{W} , hence they are not general multiplexing matrices. Moreover, these codes, termed *perfect sequences*, exist only for a very limited set of N and noise parameters. For most values of χ and N , perfect sequences do not exist.

3. Optimal Saturated Multiplexing

We begin the discussion by considering saturation. While an object may be moderately bright when illuminated by a single source, it can become saturated if illuminated by numerous light sources. When this is the case, multiplexing too many sources, *e.g.* using the \mathbf{S} -matrix is impractical. While exposure time may be reduced to counter saturation, Refs. [28, 29] proved that such a step should be avoided: it is better to decrease the number of illumination sources C activated in each measurement. This raises the need for new multiplexing codes, that comply with a constraint on C .

Assume that the saturation phenomenon is insensitive to the specific identities of the illuminating sources. Saturation is assumed to occur when the total illumination radiance exceeds a threshold, C_{sat} . If all light sources yield a similar object radiance, then C_{sat} expresses units of light sources, and is analogous to C in Sec. 2.3.

Saturation is avoided if

$$\sum_{s=1}^N w_{m,s} \leq C_{\text{sat}} \quad \forall m \in \{1, 2, \dots, N\}. \quad (17)$$

Recall that all sources can be activated with some portion of their maximum intensity *i.e.*

$$0 \leq w_{m,s} \leq 1 \quad \forall m, s \in \{1, 2, \dots, N\}. \quad (18)$$

We use Eq. (16) to formulate a maximization problem on the multiplex gain, G_0 . In this section, we do not consider photon noise. Hence, a signal-dependency of the noise is not used here. Maximizing G_0 is equivalent to minimizing its reciprocal square *i.e.*

$$\begin{aligned} \arg \max_{\mathbf{W}} G_0 &\equiv \arg \min_{\mathbf{W}} \frac{1}{G_0^2} = \\ &\arg \min_{\mathbf{W}} \frac{1}{N} \text{trace} \left[(\mathbf{W}^t \mathbf{W})^{-1} \right]. \end{aligned} \quad (19)$$

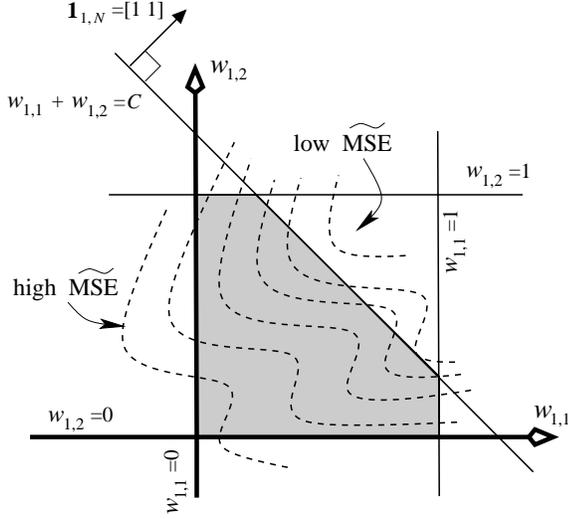


Figure 3. A 2D illustration of the optimization task. The shaded area is the domain in which \mathbf{w}_1 satisfies the constraints.

The constraints for our problem are taken from Eqs. (17,18). Thus, the optimization problem is

$$\min_{\mathbf{W}} \frac{1}{N} \text{trace} \left[(\mathbf{W}^t \mathbf{W})^{-1} \right] \quad (20)$$

$$\text{s.t. } \mathbf{1}_{1,N} \cdot \mathbf{w}_m - C_{\text{sat}} \leq 0 \quad \forall m \in \{1, \dots, N\} \quad (21)$$

$$-w_{m,s} \leq 0 \quad \forall m, s \in \{1, \dots, N\} \quad (22)$$

$$w_{m,s} - 1 \leq 0 \quad \forall m, s \in \{1, \dots, N\}. \quad (23)$$

Here $\mathbf{1}_{1,N}$ is a row vector, all of whose elements are 1 and \mathbf{w}_m is the m 'th row of \mathbf{W} . See Fig. 3 for an illustration of this optimization task.

This problem is simple if $C_{\text{sat}} > (N + 1)/2$. In this case, codes based on the \mathbf{S} -matrix are optimal. The reason is that saturation is not met in Hadamard multiplexing when $C_{\text{sat}} > (N + 1)/2$. Hence, the optimality [10] of Hadamard codes holds in this case.

We thus focus on $C_{\text{sat}} \leq (N + 1)/2$. Simulations we performed found local minima in (20). The best minimum occurred when (21) was active. This may be intuitively explained by arguing that one prefers to exploit the maximum radiance for every measurement.¹ We therefore replace Eq. (21) by the equality constraint

$$\mathbf{1}_{1,N} \cdot \mathbf{w}_m = C_{\text{sat}} \quad \forall m \in \{1, 2, \dots, N\}. \quad (24)$$

While using C_{sat} in Eq. (24) facilitates optimization under saturation, for the remainder of the work we favor the use of C instead. This is done to allow a subsequent generalization of the formulation to photon noise. Note that Eq. (24) means that \mathbf{w}_m must lie on a hyperplane (see Fig. 3), whose unit normal vector is $(1/\sqrt{N})\mathbf{1}_{1,N}^t$.

¹This argument holds if the noise is signal independent. The more general case is discussed in Sec. 4.

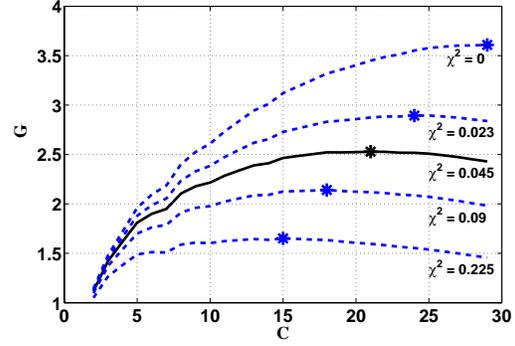


Figure 4. The expected multiplex gain for various values of χ^2 . Here $N = 57$, thus $C \in \{2, \dots, 29\}$. The solid line corresponds to our experimental system. The values of C_{opt} are marked by asterisks. They shift as the photon noise increases relative to κ_{gray}^2 .

4. Optimal Photon Limited Lighting

Sec. 3 considered only saturation. We now extend the approach presented in Sec. 3 to cope with photon noise. Solving the optimization problem in Eq. (20) subject to the constraints (22,23,24) results in an illumination matrix $\mathbf{W}(C)$, that is optimal, for a given C . In other words, we determine the values in each row \mathbf{w}_m of $\mathbf{W}(C)$, such that $\mathbf{1}_{1,N} \cdot \mathbf{w}_m$ is exactly C , while $\mathbf{W}(C)$ has the highest gain, G_0 , under a signal-independent noise model. Eq. (14) then converts G_0 to the multiplex gain under the general affine noise model.²

Recall that χ^2 can be obtained from calibration, as described in Sec. 2.2 and Eq. (15). Based on χ^2 and $G_0(\mathbf{W}(C))$, Eq. (14) yields the multiplexing gain $G(C)$. Now, let a range of values of C be scanned. For each C , we obtain $\mathbf{W}(C)$, as well as $G(C)$. This function $G(C)$ is plotted in Fig. 4, for actual results obtained by our algorithm. Out of the plot, the value of C that maximizes G is selected. In other words, this scan finds the *number of activated sources per measurement* that maximizes the gain, accounting for photon noise (via Eq. 14), and the system characteristic χ . To recap,

1. Calibrate the system to find χ^2 .
2. Scan the range of C values from 1 to C_{sat} . For each³ value of C , perform the subsequent steps 3 and 4.
3. Find the matrix $\mathbf{W}(C)$ that optimizes Eq. (20) subject to Eqs. (22,23,24).
4. Calculate the expected multiplex gain $G(C)$ using Eqs. (14,16).

²Note that there is no point in checking cases where $C \geq \frac{N+1}{2}$. They are certainly suboptimal, for a given N , as we now explain. Recall that for signal-independent noise and no saturation, G is optimized by the \mathbf{S} -matrix. From (14) it can also be seen that if G_0 is optimized, there is no point in increasing C , as it will only degrade G .

³There is no necessity for exhaustive search of $G(C)$. Since $G(C)$ is well behaved, one can incorporate efficient optimization procedures.

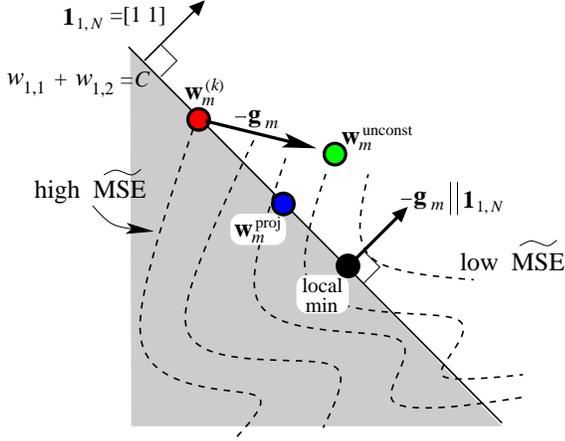


Figure 5. From \mathbf{w}_m (red dot), gradient descend yields $\mathbf{w}_m^{\text{unconst}}$ (green dot). The latter is projected into the constraint (24), yielding $\mathbf{w}_m^{\text{proj}}$ (blue). When the gradient \mathbf{g}_m is normal to the constraint surface, then $\mathbf{w}_m^{\text{proj}} = \mathbf{w}_m$ (black dot). Hence, the optimization gets stuck at a local minimum.

5. Let $C_{\text{opt}} = \arg \max_C G(C)$, as in Fig. 4.
6. The desired multiplexing code is $\mathbf{W}(C_{\text{opt}})$.

5. Minimization Procedure

We now describe a numerical scheme for solving the system given in Eqs. (20,22,23,24). It consists of a *core*, which is based on a projected gradient method [5]. It also consists of a higher-level procedure, designed to escape local minima. Define

$$\widetilde{\text{MSE}} = \frac{1}{\sigma_a^2} \text{MSE}_i(\mathbf{W}) = \frac{1}{N} \text{trace} \left[(\mathbf{W}^t \mathbf{W})^{-1} \right]. \quad (25)$$

We iteratively minimize $\widetilde{\text{MSE}}$ as a function of \mathbf{W} . The minimization core is based on projected gradient descend. In each basic step, \mathbf{W} is updated by the gradient

$$\Gamma \triangleq \frac{d\widetilde{\text{MSE}}}{d\mathbf{W}} = -\frac{2}{N} (\mathbf{W}^t \mathbf{W} \mathbf{W}^t)^{-1}. \quad (26)$$

The updated \mathbf{W} is then projected onto constraints (18) and (24), one at a time. This is illustrated in Fig. 5. Further details are given in App. A.

The $\widetilde{\text{MSE}}$ in Eq. (20) is a multimodal function of \mathbf{W} . Therefore, the core generally converges to a local minimum, rather than a global one. To escape local minima, we embed the core in a higher level process. When the core converges to a local minimum, \mathbf{W} is modified, as we describe below. Then, the core is re-initialized with the modified \mathbf{W} .

The minimization core gets stuck in a local minimum because specific rows of \mathbf{W} are prevented from undergoing any modification. This prevention is caused by the *constraints*. To understand this, note that Eq. (26) is never

nulled.⁴ Hence, following the Karush-Kuhn-Tucker theorem [5], all of the extrema of $\widetilde{\text{MSE}}$ are obtained when *constraints are active*. For this reason, local minima are caused by matrix rows which reside on constraints, as illustrated in Fig. 5. On the other hand, other rows of \mathbf{W} are free to change. We therefore seek to identify rows that stagnate the minimization core.

The m 'th row of \mathbf{W} is \mathbf{w}_m . Its corresponding row in the gradient matrix Γ is \mathbf{g}_m . When \mathbf{g}_m is parallel to $\mathbf{1}_{1,N}$, it means that this row of the gradient is orthogonal to the constraint surface (24), as illustrated in Fig. 5. If this is the case, then \mathbf{w}_m is equivalent to its projection, stagnating the minimization core. Hence, a sufficient condition of row m of \mathbf{W} to stagnate is that $\mathbf{g}_m \parallel \mathbf{1}_{1,N}$.

While this condition is sufficient, it is not a necessary one. We now describe a wider class of stagnating rows. Suppose that \mathbf{w}_m has elements s for which $w_{m,s} = 1$ or 0 and that $\mathbf{w}_m - \mathbf{g}_m$ shifts them beyond the bounds of Eqs. (22,23). Denote the set of indices of these elements by S^{overflow} . Now, define a row vector $\mathbf{g}_m^{\text{eff}} \in \mathbb{R}^{N-|S^{\text{overflow}}|}$. It is extracted from \mathbf{g}_m . It is defined as $\mathbf{g}_m^{\text{eff}} \triangleq \mathbf{g}_{m,s \notin S^{\text{overflow}}}$. Hence, it consists only of those elements s in \mathbf{g}_m whose indices are *not* in S^{overflow} . It can be shown that

$$\mathbf{g}_m^{\text{eff}} \parallel \mathbf{1}_{1,N-|S^{\text{overflow}}|} \quad (27)$$

is a necessary condition for stagnation of row m .

An algorithm is intended to detect a local minimum of the core, and then escape it:

1. Execute the *minimization core* given in App. A once. Use its output multiplexing code and corresponding $\widetilde{\text{MSE}}$ to initialize \mathbf{W}^0 and $\widetilde{\text{MSE}}^{\text{min}}$.
2. Iterate the subsequent steps 3,4,5 until the number of allowed iterations is exhausted. The iteration index is l .
3. For all $m \in \{1, \dots, N\}$, if Eq. (27) holds, then row m is detected as stagnated. Replace it by a random row vector. This new row complies with (18,24) and is formulated as described in App. B.
4. Execute the *minimization core* again. Initialize it by $\mathbf{W}^{(l-1)}$. Its output is $\mathbf{W}^{(l)}$, as well as $\widetilde{\text{MSE}}^{(l)}$ and its corresponding gradient $\Gamma^{(l)}$.
5. If $\widetilde{\text{MSE}}^{(l)} < \widetilde{\text{MSE}}^{\text{min}}$, then $\widetilde{\text{MSE}}^{\text{min}} := \widetilde{\text{MSE}}^{(l)}$.

6. Experiments

We demonstrate the new multiplexing codes by applying them to lighting. An EPSON EMP-7800 projector created patterns of light patches on a white diffuse wall, as in Ref. [29]. Light reflected by these patches acted as distinct

⁴A valid inverse of a matrix A can never be nulled. If it could, it would have yielded a contradiction: $\mathbf{A}^{-1}\mathbf{A} = \mathbf{I}_{N \times N} = \mathbf{0}_{N \times N}\mathbf{A}$, where $\mathbf{0}_{N \times N}$ is an $N \times N$ null matrix.



Figure 6. Multiplexing codes produced by our algorithm. [Left] $\{N, C\} = \{57, 24\}$. [Right] $\{N, C\} = \{47, 12\}$. Here, black pixels denote $w_{m,s} = 0$. White denotes $w_{m,s} = 1$. The intermediate values are in gray.

sources irradiating the viewed objects. The exposure time of the Dragonfly camera was $63msec$, corresponding to a $15Hz$ frame rate. It eliminates radiance fluctuations of the projector [29], which have a period of $7msec$.

6.1. Calibration

For noise calibration, images of the object were taken, by simply turning on C of the N illumination sources. For each value of C , a sequence of 10 frames was taken. From this sequence, the noise variance $\sigma_a^2(x, y, C)$ was estimated per pixel (x, y) . Then, spatial mean yielded $\sigma_a^2(C)$. This process was repeated for a range of C values. The resulting $\sigma_a^2(C)$ generally agreed with the affine noise model, as in Fig. 2. From the plot of $\sigma_a^2(C)$, the parameters κ_{gray} and η were extracted. Consequently, Eq. (15) yielded χ . For example, in an experiment using $N = 57$ light patches (sources), we obtained $\kappa_{gray} = 42.4$ graylevels, $\eta = 9.0$ [graylevels/light-source] thus $\chi^2 = 0.045$.

6.2. Constructing Multiplexing Codes

Following the calibration, multiplexing codes were tailored. Our algorithm can deal with an arbitrary value of N or χ , even if no Hadamard code or a matrix suggested by [35] exists for these parameters. In this domain, lack of competing codes in the literature has been the rule, rather than exception. However, to make a comparison when possible, we deliberately selected, in the following experiments, special cases having values of N and χ , for which competing codes exist.

For $N = 57$, we obtained $\mathbf{W}(C_{opt} = 24)$, shown in Fig. 6. Ref. [35] yielded $C = 20$ as the value that should be used, considering our calibrated value of χ . However, Ref. [35] does not provide a way to obtain a multiplexing code having $\{N, C\} = \{57, 20\}$. Rather, it only offers $\{N, C\} = \{57, 8\}$: it is not optimal, but it is the only competing code in this specific case.

Experiments were also conducted to compare performance vs. Hadamard codes (\mathbf{S} -matrices), for $N = 47$ and $N = 11$. The respective values of C_{opt} in our setup were $C_{opt} = 12$ and 5 . The matrix corresponding to $\{N, C\} = \{47, 12\}$ is shown in Fig. 6. In all case, we also compared the performance to that obtained by multiplexing using the identity matrix (trivial acquisition).

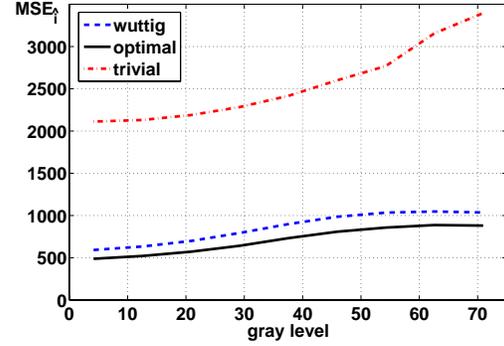


Figure 7. MSEs obtained by decoding of illumination multiplexed images ($N = 57$). Our optimal code outperforms both the code of Ref. [35] and trivial illumination.

6.3. Measurements

We used each set of codes to illuminate a scene while acquiring image sets. From each set of acquired images, we reconstructed the scene as if illuminated by individual illumination sources. This procedure was repeated 10 times, to facilitate empirical estimation of MSE_i .

For $N = 57$, an example of a demultiplexed image is shown in Fig. 1. The corresponding MSE_i is plotted in Fig. 7. A demultiplexing example for $N = 47$ is plotted in Fig. 8. The estimated MSE_i for both $N = 11$ and $N = 47$ are shown in Fig. 9. In these rare cases, where competing codes exist, the best multiplexing scheme (lowest output noise) is the one created by our method.

The experiment using $N = 47$ demonstrates the implication of saturation on the applicability of Hadamard codes. In this case, single-source illumination created a bright spot in a small part of the raw image (the soda can in Fig. 8). Although most of the image is dark (graylevels up to 1000 in the 16-bit data), the highlight exhausted the dynamic range. In this case, saturation bans Hadamard code from being used. Nevertheless, we did use Hadamard codes, and the plots in Fig. 9 use only the unsaturated pixels. Needless to say, in pixels that were saturated by Hadamard-coded illumination, the data was useless. However, in these pixels saturation was avoided by our codes.

7. Discussion

Our approach provides optimal multiplexing codes for every desired number of light sources N and radiance inhibition (saturation, photon noise). It does so for cases that are much more general than those reported in the literature, covering cases for which no codes are known. By accounting for fundamental physical limits of image acquisition, we achieve results that are superior to other multiplexing codes, even when such codes exist. Our work may apply to many applications that use multiplexing, other than object

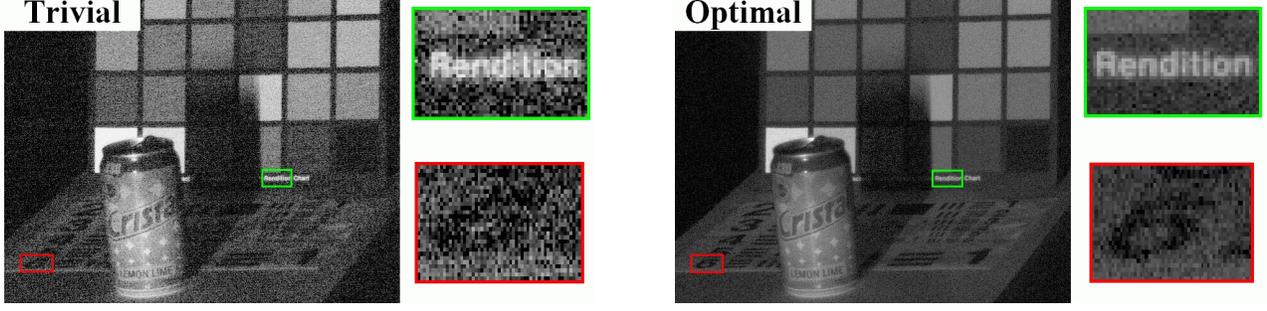


Figure 8. [Left] An image taken under a single light source. [Right] An image of the same scene, decoded from $\{N = 47, C = 12\}$ multiplexed frames. It is decoded as if illuminated by the same single source. The multiplexing code is optimal. The marked rectangles are magnified to the right of each image.

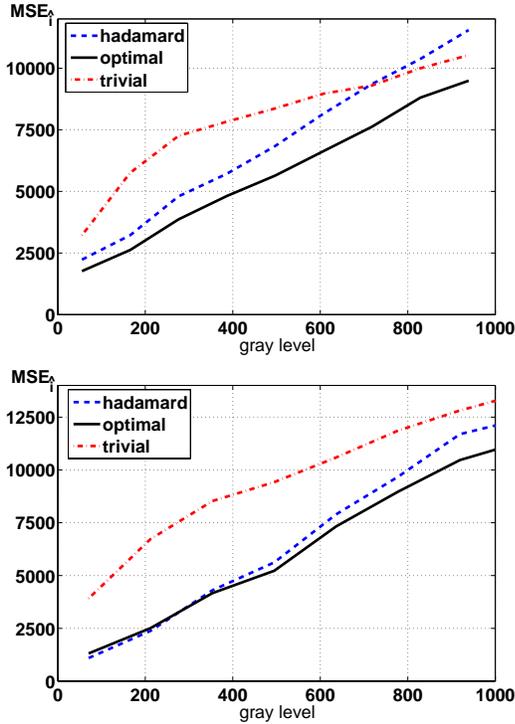
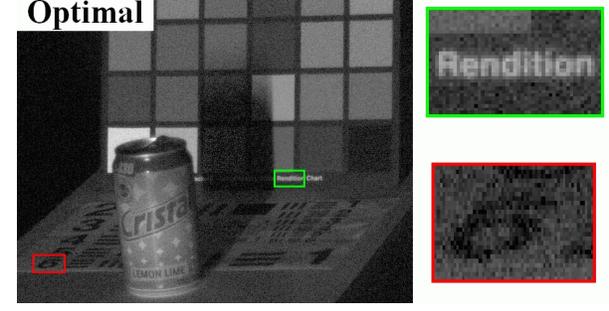


Figure 9. MSEs of the images decoded from illumination multiplexed frames. [Top] $N = 11$. Hadamard multiplexing becomes counter productive for high gray levels. Our multiplexing code is better than the Hadamard code and the identity (trivial) matrix. [Bottom] $N = 47$. Plotting only pixels unsaturated by Hadamard.

lighting (Xray, spectroscopy, coded aperture imaging etc.).

Ref. [29] showed that compensating for nonlinearity in γ -corrected cameras induces radiance noise that is similar to the affine noise model. Hence, the formalism used here may apply to such cameras. Moreover, σ_a can also have a multiplicative component that stems from fluctuations in the light sources being multiplexed, as analyzed in Ref. [29]. It is thus worth accounting for this effect as well.



A. Minimization Core

We iterate on k , minimizing $\widetilde{\text{MSE}}$ as a function of \mathbf{W} .

1. Given N and C , create an initial matrix \mathbf{W}_0 (See App. B). The initial matrix \mathbf{W}_0 complies with constraints (22,23,24).
2. Repeat the subsequent stages 3 and 4, until $|\widetilde{\text{MSE}}_k - \widetilde{\text{MSE}}_{k+1}| < \epsilon$, where ϵ is a pre-determined small threshold.
3. Calculate the gradient, Γ_k (Eq. 26). Then, calculate an updated matrix $\mathbf{W}_{k+1}^{\text{unconst}} \triangleq \mathbf{W}_k - \Gamma_k$, as in standard gradient descent (Fig. 5). We take care of the step size in stage 4.
4. Project $\mathbf{W}_{k+1}^{\text{unconst}}$ in the following way:
 - (i) Project $\mathbf{W}_{k+1}^{\text{unconst}}$ onto the hyperplane used in (24) as in Fig. 5. It is easy to show that for each row $\mathbf{w}_m^{\text{unconst}}$, its projection is

$$\mathbf{w}_m^{\text{proj}} = \mathbf{w}_m^{\text{unconst}} + \left[C - \sum_{s=1}^N w_{m,s}^{\text{unconst}} \right] \frac{\mathbf{1}_{1,N}^t}{N}. \quad (28)$$

- (ii) Denote $\mathbf{d}_k \triangleq \mathbf{W}_{k+1}^{\text{proj}} - \mathbf{W}_k$. Then, update $\mathbf{W}_{k+1}^{\text{unbounded}} = \mathbf{W}_k - \beta \mathbf{d}_k$, where β is a parameter controlling the step size.
- (iii) Project $\mathbf{W}_{k+1}^{\text{unbounded}}$ onto constraints (22,23) to create \mathbf{W}_{k+1} . This is done by truncating the elements of $\mathbf{W}_{k+1}^{\text{unbounded}}$ to $[0, 1]$.

B. Initialization of the Minimization Core

We now describe the initialization procedure for the initial matrix \mathbf{W}_0 . We randomly generate $w_{m,s} \in [0.1, 0.9] \forall m, s$. This avoids activation of constraints (22,23) in this step. We then normalize each row, m of \mathbf{W}_0 such that its sum is C . Any element violating (18) is regenerated and the normalization process is then repeated until satisfaction of Eqs. (18, 24).

Acknowledgments

We thank Einav Namer for her help in experiments. Yoav Schechner is a Landau Fellow - supported by the Taub Foundation, and an Alon Fellow. The work was supported by the Israeli Ministry of Science, Culture and Sport (Grant 3-3426). It was conducted in the Ollendorff Minerva Center. Minerva is funded through the BMBF.

References

- [1] A. Agrawal, R. Raskar, S. K. Nayar, and Y. Li. Removing photography artifacts using gradient projection and flash-exposure sampling. *ACM TOG*, 24:828–835, 2005.
- [2] F. Alter, Y. Matsushita, and X. Tang. An intensity similarity measure in low-light conditions. In *Proc. ECCV Vol. 4*, pages 267–280, 2006.
- [3] H. H. Barrett and W. Swindell. *Radiological Imaging*, volume 1. Academic, New York, 1981.
- [4] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. *IEEE Trans. PAMI*, 25:218–233, 2003.
- [5] D. P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1999.
- [6] O. G. Cula, K. J. Dana, D. K. Pai, and D. Wang. Polarization multiplexing for bidirectional imaging. In *Proc. IEEE CVPR Vol. 2*, pages 1116–1123, 2005.
- [7] P. Debevec. Image-based lighting. *IEEE Computer Graphics and Applications*, 22:26–34, 2002.
- [8] E. E. Fenimore and T. M. Cannon. Coded aperture imaging with uniformly redundant arrays. *Applied Optics*, 17:337–347, 1978.
- [9] D. Goldman, B. Curless, A. Hertzmann, and S. M. Seitz. Shape and spatially-varying BRDFs from photometric stereo. In *Proc. IEEE ICCV*, pages 341–348, 2005.
- [10] M. Harwit and N. J. A. Sloane. *Hadamard Transform Optics*. Academic, New York, 1979.
- [11] G. Healey and T. O. Binford. Local shape from specularities. In *Proc. IEEE ICCV*, pages 151–160, 1987.
- [12] S. Ioué and K. R. Spring. *Video Microscopy*, 2nd ed., 1997. ch. 6,7,8, Plenum Press, New York.
- [13] K. C. Lee, J. Ho, and D. J. Kriegman. Acquiring linear subspaces for face recognition under variable lighting. *IEEE Trans. PAMI*, 27:684–698, 2005.
- [14] H. Lensch, J. Kautz, M. Gosele, W. Heidrich, and H. Seidel. Image-based reconstruction of spatial appearance and geometric detail. *ACM TOG*, 22:234–257, 2003.
- [15] M. Levoy, B. Chen, V. Vaish, M. Horowitz, I. McDowall, and M. Bolas. Synthetic aperture confocal imaging. *ACM TOG*, 23:825–834, 2004.
- [16] C. Liu, W. T. Freeman, R. Szeliski, and S. B. Kang. Noise estimation from a single image. In *Proc. IEEE CVPR Vol. 1*, pages 901–908, 2006.
- [17] Q. T. Luong, P. Fua, and Y. Leclerc. Recovery of reflectances and varying illuminants from multiple views. In *Proc. ECCV Vol. 3*, pages 163–179, 2002.
- [18] S. R. Marschner, S. H. Westin, E. P. F. Lafortune, and K. E. Torrance. Image-based bidirectional reflectance distribution function measurement. *Applied Optics*, 39:2592–2600, 2000.
- [19] W. Matusik, H. Pfister, A. Ngan, P. Beardsley, R. Ziegler, and L. McMillan. Image-based 3D photography using opacity hulls. *ACM TOG*, 21:427–437, 2002.
- [20] F. Moreno-Noguer, S. K. Nayar, and P. N. Belhumeur. Optimal illumination for image and video relighting. In *Proc. IEE Conf. Visual Media Production*, pages 201–210, 2005.
- [21] Y. Moses, S. Ullman, and S. Edelman. Generalization to novel images in upright and inverted faces. *Perception*, 25:443–461, 1996.
- [22] G. Nitzsche and R. Riesenberger. Noise, fluctuation and HADAMARD-transform-spectrometry. In *Proc. SPIE*, volume 5111, pages 273–282, 2003.
- [23] M. Osadchy and D. Keren. Efficient detection under varying illumination conditions and image plane rotations. *Computer Vision & Image Understanding*, 93:245–259, 2004.
- [24] J. A. Paterson, D. Claus, and A. W. Fitzgibbon. BRDF and geometry capture from extended inhomogeneous samples using flash photography. In *Proc. Eurographics*, 2005.
- [25] R. Ramamoorthi and P. Hanrahan. Frequency space environment map rendering. *ACM TOG*, 21:517–526, 2002.
- [26] R. Raskar, K. H. Tan, R. Feris, J. Yu, and M. Turk. Nonphotorealistic camera: depth edge detection and stylized rendering using multiframe imaging. *ACM TOG*, 23:679–688, 2004.
- [27] I. Sato, T. Okabe, Y. Sato, and K. Ikeuchi. Using extended light sources for modeling object appearance under varying illumination. *Proc. IEEE ICCV*, pages 325–332, 2005.
- [28] Y. Y. Schechner, S. K. Nayar, and P. N. Belhumeur. A theory of multiplexed illumination. In *Proc. IEEE ICCV Vol. 2*, pages 808–815, 2003.
- [29] Y. Y. Schechner, S. K. Nayar, and P. N. Belhumeur. Multiplexing for optimal lighting. *IEEE Trans. PAMI*, to be published, 2007.
- [30] P. Sen, B. Chen, G. Garg, S. R. Marschner, M. Horowitz, M. Levoy, and H. P. A. Lensch. Dual photography. *ACM TOG*, 24:745–755, 2005.
- [31] A. Shashua. On photometric issues in 3D visual recognition from a single 2D image. *Int. J. Comp. Vis.*, 21:99–122, 1997.
- [32] T. Sim, S. Baker, and M. Bsat. The CMU pose, illumination and expression database. *IEEE Trans. PAMI*, 25:1615–1618, 2003.
- [33] G. K. Skinner. X-ray imaging with coded masks. *Scientific American*, 259:84–89, 1988.
- [34] A. Wenger, A. Gardner, C. Tchou, J. Unger, T. Hawkins, and P. Debevec. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM TOG*, 24:756–764, 2005.
- [35] A. Wuttig. Optimal transformations for optical multiplex measurements in the presence of photon noise. *Applied Optics*, 44:2710–2719, 2005.
- [36] A. Yuille, J. M. Coughlan, and S. Konishi. The KGBR viewpoint-lighting ambiguity. *JOSA A*, 20:24–31, 2003.
- [37] T. Zickler. Reciprocal image features for uncalibrated helmholtz stereopsis. In *Proc. IEEE CVPR Vol. 2*, pages 1801–1808, 2006.