

Modeling the Space of Camera Response Functions

Michael D. Grossberg, *Member, IEEE Computer Society*, and Shree K. Nayar

Abstract—Many vision applications require precise measurement of scene radiance. The function relating scene radiance to image intensity of an imaging system is called the camera response. We analyze the properties that all camera responses share. This allows us to find the constraints that any response function must satisfy. These constraints determine the theoretical space of all possible camera responses. We have collected a diverse database of real-world camera response functions (DoRF). Using this database, we show that real-world responses occupy a small part of the theoretical space of all possible responses. We combine the constraints from our theoretical space with the data from DoRF to create a low-parameter empirical model of response (EMoR). This response model allows us to accurately interpolate the complete response function of a camera from a small number of measurements obtained using a standard chart. We also show that the model can be used to accurately estimate the camera response from images of an arbitrary scene taken using different exposures. The DoRF database and the EMoR model can be downloaded at <http://www.cs.columbia.edu/CAVE>.

Index Terms—Radiometric response function, camera response function, calibration, real-world response curves, empirical modeling, high-dynamic range, recovery of radiometry, nonlinear response, gamma correction, photometry, sensor modeling.

1 SCENE RADIANCE TO IMAGE INTENSITY

COMPUTATIONAL vision seeks to determine properties of a scene from images. These properties include 3D geometry, reflectance, and lighting. Algorithms to determine such properties require accurate models of image formation. Models of image formation must account for the characteristics of the imaging system. For example, algorithms that recover geometric properties of a scene must model the system's imaging geometry and account for the system's spatial resolution. Algorithms that require accurate color measurements must account for the system's spectral response.

Many computer vision algorithms require precise measurements of scene radiance to recover desired scene properties. Examples of algorithms that explicitly use scene radiance measurements are color constancy [10], [19], construction of linear high-dynamic range images [3], [21], [26], [6], [22], photometric stereo [2], [28], [31], shape from shading [17], [28], [33], estimation of reflectance and illumination from shape and intensity [20], recovery of BRDF from images [4], [7], [25], [29], [32], and surface reconstruction using Helmholtz stereopsis [34].

The goal of this work is to provide an accurate and convenient model of the mapping from scene radiance to image intensity by an imaging system.¹ In general, this mapping comprises several complex factors, including vignetting and lens fall-off [1]. In a digital camera, some other factors are fixed pattern noise, shot noise, dark current, and read noise [16]. Note, we focus on modeling

the expected image intensity given a scene radiance at a pixel. Thus, we will ignore zero-mean noise for determining the mapping. In the case of a film camera, factors influencing the mapping include the photosensitive response of the film as well as the film developing process [18]. The mapping must also account for the results of digitally scanning the film.

Regardless of the individual factors involved, we can assume the mapping is a composite of just two functions, s and f , as shown in Fig. 1. The function s represents the effect of transmission through the imaging system's optics as well as the fixed pattern noise which is present in digital cameras. This function may vary spatially over the image, but can be assumed to be linear with respect to scene radiance [1], [16]. The function s models the transformation of scene radiance to image irradiance. We also note that to simplify our exposition, we assume fixed integration time.

The photosensitive elements of the image sensor respond to the image plane irradiance E by producing a signal either electronically in a solid state camera, or chemically on film. For an analog video camera, this typically passes through a capture board which then produces a value we call image intensity B . For film, the developed image must be digitally scanned to obtain an image intensity.² We model the entire imaging system response with the function f from Fig. 1. This response is generally a nonlinear function of image irradiance and is called the *camera response function*.

In many imaging devices, the nonlinearity of f is intentional. A nonlinear mapping is a simple means to compress a wide range of irradiance values within a fixed range of measurable image intensity values. Manufacturers produce photographic films with specific nonlinear characteristics. In the case of solid-state cameras, the CCD or CMOS responds linearly to irradiance. Nonlinearities are purposely introduced in the camera's electronics to mimic

1. Part of this work has appeared in [14].

• The authors are with the Computer Science Department, Columbia University, 500 West 120th Street, Room 450, New York, NY 10027. E-mail: {mdog, nayar}@cs.columbia.edu.

Manuscript received 15 Apr. 2003; revised 21 Dec. 2003; accepted 24 Feb. 2004.

Recommended for acceptance by Z. Zhang.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-0035-0403.




Fig. 1. Flow diagram showing the two basic transformations, s and f , that map scene radiance L to image intensity B . The function s models the optics of the imaging system. This function may vary spatially, but is generally linear. The mapping f of image irradiance to image intensity is called the camera response function. It is usually nonlinear, but can be assumed to be spatially uniform.

the nonlinearities of film, to mimic the response of the human visual system, or to create a variety of aesthetic effects. Once we determine the camera response function, we can invert it, making it possible to transform pixel intensity values to image irradiances. Going from image irradiance to scene radiance can then be accomplished by finding s , [1], [16].

Though nonlinear, a camera's response function f is generally uniform across the spatial dimensions of the image. Hence, it is described by a one-variable function of irradiance, $B = f(E)$. Inversion of the camera response function allows the transformation of image intensity to image irradiance. Going from image irradiance to scene radiance can then be accomplished by finding s , which is easy to do once the response f is known [1], [16]. In any case, this work will focus on the first step, which is the recovery of the nonlinear camera response.

A number of algorithms have been introduced in computer vision and computer graphics to estimate the camera response from multiple images of a scene taken with different exposures [6], [23], [22], [26], [30]. All these methods make a priori assumptions about the form of the response function.³ For example, in [22], Mann and Picard describe the recovery of the parameters for a parameterized response model, such as when the response has the form of a gamma curve, $f(E) = \alpha + \beta E^\gamma$. They find the parameters γ , α , and β from multiple registered images of a static scene taken using different exposures. The single image recovery algorithm presented in [9] also assumes the response is a gamma curve. The response functions found in cameras, however, can vary significantly from a gamma curve.

In [23], Mann lists a number of alternative analytical forms for modeling response functions. In contrast, in [6] [24], and [30], no particular form is assumed, instead smoothness constraints are imposed. In what can be considered a compromise between these two extreme assumptions, Mitsunaga and Nayar [26] use the general approximation model of polynomials. They assume a low degree polynomial gives a sufficient approximation to the response and estimate its coefficients.

It is important to note that, while much recent work acknowledges the importance of the camera response, a careful analysis and modeling of the response has yet to be done. We wish to address this void. In doing so, we seek answers to the following fundamental questions:

- **What is the space of possible camera response functions?** We show that all response functions must lie within a convex set that results from the intersection of a hyperplane and a positive cone in

3. It was recently shown in [13] that, to avoid ambiguities, a priori constraints on the response function are imperative when finding the response from multiple images.

function space. This gives us both guidance on the form of our model as well as constraints.

- **Which camera response functions within this space arise in practice?** We compiled a Database of Response Functions (DoRF) of a variety of imaging systems including film, CCD, and solid-state camera components that are currently used. Our goal is to represent the variety of response functions which occur in complete imaging systems. The database currently includes a total of 201 real-world response functions.
- **What is a good model for response functions?** We combine the constraints from our analysis and the data from DoRF to formulate a new Empirical Model of Response (EMoR) which can model a wide gamut of response functions with a very small number of parameters. We show that EMoR outperforms alternative models including previously used ones, in terms of accuracy. We provide a comparison with a log-space version of our model.

We show that EMoR works well by using a number of different evaluation metrics. We demonstrate that EMoR can be used to recover complete response functions from an image of a chart with a few known reflectances. It can also be used to accurately determine a camera's response from a set of images of a scene taken at different exposures. We have made the DoRF database and the EMoR model available at <http://www.cs.columbia.edu/CAVE>.

2 WHAT IS THE THEORETICAL SPACE OF CAMERA RESPONSE FUNCTIONS?

Defining the theoretical space of camera response functions helps us build a model of these functions. It guides us as to the form such a model should take. It also determines the constraints these functions must satisfy. We begin by stating our assumptions.

Our first assumption is that the response function f is the same at each pixel. In theory, an imaging system could have a different camera response function for every pixel. For example, a CCD has a linear variation called fixed pattern noise [15] which changes the response from pixel to pixel. Linear spatial variations can be folded into the function s (see Fig. 1), which includes effects such as lens fall-off [1]. As pointed out in Section 1, after removing such variations, the response is a one-variable function of image irradiance, $f(E) = B$, where B is image intensity.

Our second assumption is that the range of our camera's response goes from B_{MIN} to B_{MAX} . These values are easily computed. For example, in digital cameras B_{MIN} is the mean of the thermal noise. This may be estimated from an image taken with the lens cap on. The number B_{MAX} may be determined by taking an image of a very bright object such as a light source, so that parts of the image are saturated.




Fig. 2. To visualize the theoretical space of camera response functions, we represent the response functions as vectors. Vectors that satisfy $f(1) = 1$ lie on the hyperplane W_1 . Vectors satisfying the monotonic condition lie in the shaded solid polygonal cone Λ . The theoretical space W_{RF} of camera response functions is the darkly shaded intersection $\Lambda \cap W_1$.

The units of response are arbitrary, so we normalize the response so that $B_{MIN} = 0$ and $B_{MAX} = 1$.

Our third assumption is that the normalized response function is monotonic. If the response is not *strictly* monotonic, it is *many to one* and, thus, cannot be linearized. This limits its usefulness in computer vision. In our analysis, it is easier to work with monotonic functions. In practice, the minor discrepancy between monotonic and strictly monotonic is not a problem. Without loss of generality, we assume f monotonically *increases*. If f is monotonically decreasing, such as in the case of a negative image, we replace the given response with the function $1 - f$. This implies that corresponding to B_{MIN} and B_{MAX} are minimum and maximum detectible irradiances, E_{MIN} and E_{MAX} . These parameters may be incorporated into the function s in Fig. 1, since they represent a linear scaling and shift along the irradiance axis. Therefore, we normalize irradiance so that $E_{MIN} = 0$ and $E_{MAX} = 1$.

With these assumptions, we define the space of camera response functions as:

$$W_{RF} := \{f | f(0) = 0, f(1) = 1, \text{ and } f \text{ monotonically increasing}\}.$$

The exact form of the space W_{RF} is easier to understand in terms of vectors. Any function f of irradiance (not necessarily a response) may be thought of as a vector by sampling it at a set of fixed increasing irradiance levels. That is, the function f becomes the finite-dimensional vector⁴ $(B_1, \dots, B_P) = (f(E_1), \dots, f(E_P))$. We set the brightest sampled irradiance to be $E_P = 1$.

Response functions f are normalized such that $B_P = f(1) = 1$. Therefore, all response vectors must lie in the hyperplane W_1 shown in Fig. 2, where the last component B_P is 1. If f and f_0 are any two response vectors in the hyperplane W_1 , the difference $h = f - f_0$ lies in a parallel hyperplane W_0 going through the origin (see Fig. 2). Therefore, any response function can be expressed as $f = f_0 + h$ where f_0 is some base response function and $h \in W_0$.

4. We will treat f interchangeably as a vector and a function. We assume the function f is smooth enough so that we can recover it from a vector of dense samples by interpolation.

Now, the additional constraint that a response function is monotonic, can be interpreted as requiring the function's first derivatives to be positive. Any positive linear combination of two functions with positive derivatives must also have positive derivatives. We know that a set is a cone when it has the property that positive linear combinations of its elements lie within it. Therefore, monotonic functions can be represented by a cone, shown as Λ in Fig. 2.

Combining both of the above constraints, we see that W_{RF} is the intersection (the darkly shaded region in Fig. 2) of the cone Λ with the hyperplane W_1 :

$$W_{RF} = W_1 \cap \Lambda. \quad (1)$$

Note that the convexities of the hyperplane and the cone imply that the intersection of (1) is also convex. If $p, q \in W_{RF}$ and $0 \leq \alpha \leq 1$, then $\alpha p + (1 - \alpha)q \in W_{RF}$. That is, positive weighted sums of response functions are also response functions. As a consequence, the mean of any set of camera response functions is also a valid camera response function. Another consequence of convexity is that we can approximate the set response functions with a series of linear inequalities. By approximating the set in terms of linear inequalities, we can use standard optimization algorithms to, for example, find the closest function in W_{RF} to an arbitrary function f .

3 APPROXIMATION MODELS FOR THE RESPONSE FUNCTION

Even though the theoretical space of response functions W_{RF} is restricted to an intersection of a hyperplane and a cone, it is still infinite-dimensional. However, there are a limited number of processes that are used in films and in solid state detectors to collect light. As a result, many functions within the theoretical space never arise in practice. It therefore makes sense to look for a finitely parameterized subset of W_{RF} which approximates the set of real-world response functions. We describe two approaches which give approximation models. In each approach we choose the number of parameters M of our model. As M becomes large, the model better approximates elements in W_{RF} .

The simplest approach to parameterizing W_{RF} uses (1). We note that $W_1 = f_0 + W_0$. We further observe that any choice of basis h_1, h_2, \dots for the vector space W_0 gives an approximation model. The first M basis elements give the M th order approximation:

$$f_0(E) + \sum_{n=1}^M c_n h_n(E), \quad (2)$$

where c_1, \dots, c_M are the coefficients or parameters of the model.⁵

A somewhat more complex approach generalizes both the gamma functions⁶ used in Mann and Picard [22] and the log-space least-squares solutions used in Debevec and Malik [6]: We parameterize the log of the response functions W_{log} . Functions $g_{log} \in W_{log}$ in this space have $g_{log}(1) = 0$, and $\lim_{E \rightarrow 0} g_{log}(E) = -\infty$. These linear conditions define a vector space. If we choose a basis for this space $h_{log,1}, h_{log,2}, \dots$, then

5. Note that, due to normalization of f in Section 2, the model implicitly has the scale and offset parameters B_{MIN} , B_{MAX} , E_{MIN} , and E_{MAX} .

6. For a normalized response, $\alpha = 0$ and $\beta = 1$, in the notation of Section 1.

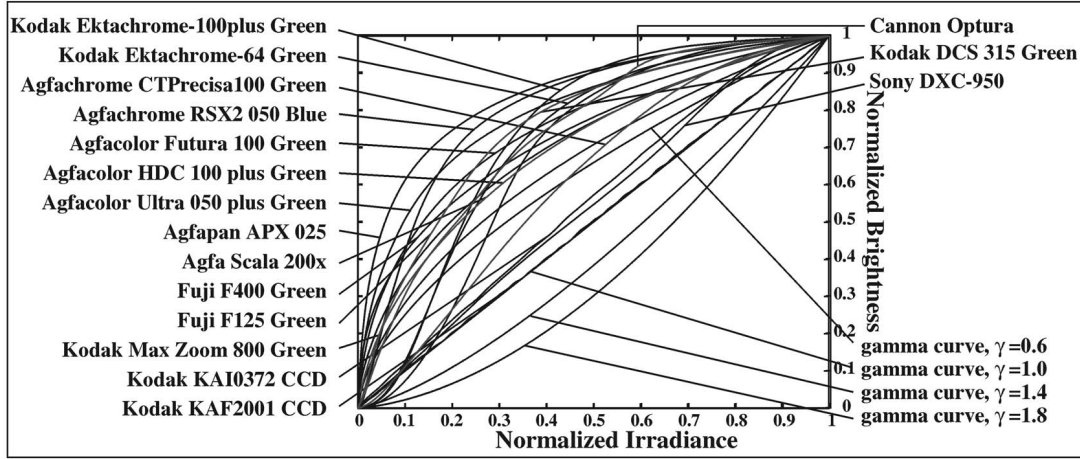


Fig. 3. Examples from our database of 201 real-world response functions (DoRF). The database includes photographic films, digital cameras, CCDs, and synthetic gamma curves. Note that even within a single brand of film, for example, Agfa, there is considerable variation between response curves.

the finite span $\sum_{n=1}^M c_n h_{\log,n}(E)$ in log-space becomes $\prod_{n=1}^M (e^{h_{\log,n}(E)})^{c_n}$ in the space of response functions. This gives an M -parameter approximation space, with c_1, \dots, c_M being the parameters of the model. If the basis is chosen so $h_{\log,1}(E) = E$, then the one parameter approximation is the family of gamma functions.

In both approaches, an approximation is determined by a choice of basis in W_0 or in W_{\log} . For example, in [26], the polynomial basis presented is a special case of the above simpler approach of parameterizing W_0 , to approximate W_{RF} . If we instead use a finite-dimensional discrete approximation of log-space by sampling at integer gray-levels and let the basis $h_{\log,n}$ be delta functions at those gray-levels, this specializes to the method used in [6].

In our notation in (2), the polynomial model is obtained using $f_0(E) := E$ and $h_n(E) := E^{n+1} - E$. One can also obtain a trigonometric approximation model by using $f_0(E) := E$ and the half-sine basis $h_n(E) := \sin(n\pi E)$. Clearly, there are many more choices. Thus, while the description of W_{RF} in Section 2 in terms of W_0 and f_0 has suggested the general form of an approximation model, it has not given us criteria to decide which basis of W_0 to use. The efficiency of any basis depends on how close the responses of actual imaging systems are to the space spanned by the first few basis elements. Hence, a natural approach is to use the response curves of real-world imaging systems to determine the appropriate basis for the approximation model.

4 REAL-WORLD RESPONSE FUNCTIONS

We collected a diverse set of response curves in order to cover the range of curves found in complete imaging systems. Examples of such systems include video cameras with capture cards and digitally scanned photographs. We collected response curves for a wide variety of CCD sensors, digital cameras (detector + electronics), and photographic films. The response curves for photographic film remain important even as the use of film declines. The film response curves have been designed to produce attractive images. Digital cameras often emulate these curves.

Companies such as Kodak, Agfa, and Fuji have published response curves for some of their films on their Web sites. The curves we gathered include representatives from positive and negative films, consumer and professional films, still and motion picture films, in both color as well as black and white. We treated the three response curves for color films as three different responses. We also included curves of the same film type but different ASA speeds. For some of the black and white films (for example, Agfa Scala 200), the response curves for different developing times were available and so included. Examples of film brands we included are Agfacolor Future, Auxochrome RX-II, Fuji F125, Fuji FDIC, Kodak Advanced, Kodak Gold, and Monochrome.

We also obtained response curves for several CCD sensors, in particular, Kodak's KAI and KAF series. In the case of digital cameras, the manufacturers we contacted were unwilling to provide the responses of their cameras. However, Mitsunaga and Nayar have measured the responses of a variety of digital and video cameras, including the Sony DC 950 and the Canon Optura using their algorithm RASCAL [26]. These curves were included. Many camera manufacturers design the response to be a gamma curve. Therefore, we included a few gamma curves, chosen from the range $0.2 \leq \gamma \leq 2.8$, in our database. Currently, the database contains a total of 201 curves, a few of which are shown in Fig. 3.

The companies provided the curves for film and CCDs in an assortment of formats. We first converted all the plotted curves to high-resolution images. We manually removed all extraneous information from these images. After removal of this information, some of the curves had gaps. To interpolate through these gaps, as well as to remove any effects of rasterization, we applied a local linear regression to the curves.

As we discussed in Section 2, we assume that response curves are monotonic. For this reason, the 201 response curves we chose were all monotonic. The few nonmonotonic ones we came across were disregarded. In the case of negative film, we transformed the curves to make them monotonically increasing rather than monotonically decreasing. Many of the film curves we collected were originally published on log-linear or log-log scales. All curves were converted to linear-linear scale in response and irradiance. Those curves that were not originally provided




Fig. 4. (a) The mean of 175 camera responses in DoRF used as the base curve f_0 in the EMoR model. (b) Four eigenvectors (functions) corresponding to the largest four eigenvalues of the covariance matrix for the 175 curves. (c) A plot showing the percentages of the energies captured by V_M , the span of the M principal components. The subspace corresponding to the three largest eigenvalues (an EMoR model with three parameters) captures more than 99.5 percent of the energy.

on a linear-linear scale were no longer uniformly sampled after conversion. We resampled these curves uniformly using linear interpolation to preserve monotonicity. We choose the original sampling densely enough for any interpolation error to be small.

5 AN EMPIRICAL MODEL OF RESPONSE

In this section, we present a new model for the camera response which combines the general form of the approximation model of (2) with the empirical data in the DoRF database described in Section 4. To create as well as test such a model, we segregated the DoRF database into a *training set* of 175 response curves and a *testing set* with 26 curves. We denote the training curves as $\{g_1, \dots, g_N\} \subset W_{\text{RF}}$, where $N = 175$ and W_{RF} is the theoretical space of responses defined in Section 2.

Strictly speaking, W_{RF} is a space of continuous functions, we may extend $\{g_1, \dots, g_N\}$ to continuous functions using interpolation since they are densely sampled. Alternatively, we may think of W_{RF} in terms of its finite-dimensional approximation where it consists of functions sampled uniformly at the sample points of $\{g_1, \dots, g_N\}$.

We would like to find a low-dimensional approximation of W_{RF} , based on our linear model from (2). Our goal is to determine a base function f_0 and a basis $\{h_1, h_2, \dots, h_M\}$ for (2) so that the root mean square approximation error of the model is small for the empirical data from DoRF. We can achieve this by applying Principal Component Analysis (PCA) to the training curves from DoRF. We will refer to this basis as the Empirical Model of Response (EMoR).

Recall from Section 2 that W_{RF} is a convex set. This implies that the mean curve $(1/N) \sum_{n=1}^N g_n$ is also a response function. We choose f_0 in (2) to be the mean curve, which is shown in Fig. 4a. It represents the 0th order approximation to W_{RF} .

Performing PCA on continuous functions, we obtain the basis functions $\{h_1, h_2, \dots\}$ as eigenfunctions of the covariance operator.⁷ The covariance operator is an integral operator C , where $(Cf)(E) = \int_0^1 c(E, x)f(x)dx$, and the kernel of this operator $c(E_1, E_2)$ is defined as

$$c(E_1, E_2) = \sum_{n=1}^N (g_n(E_1) - f_0(E_1))(g_n(E_2) - f_0(E_2)). \quad (3)$$

This integral operator is symmetric and, hence, diagonalizable with a basis of eigenvectors.

To find the basis of (2) using PCA, we work with a finite-dimensional approximation. By densely sampling each response curve f at points $\{E_1, \dots, E_P\}$, we approximate f by the vector $(f(E_1), \dots, f(E_P))$. Using all the response vectors in our training set, the elements of its symmetric covariance matrix C are found as:

$$C_{m,n} = \sum_{p=1}^N (g_p(E_n) - f_0(E_n))(g_p(E_m) - f_0(E_m)).$$

We write V_M for the span of the eigenspaces associated with the largest M eigenvalues of the matrix C . The space V_M is the best M -dimensional approximation to the space W_0 [8]. The curves in Fig. 4b are the eigenvectors for the four largest eigenvalues of the covariance matrix C .

The cumulative energies associated with the eigenvalues increase rapidly, as seen in Fig. 4c. This shows that EMoR represents the space of response functions well. In fact, three eigenvalues explain more than 99.5 percent of the energy. This suggests that even a 3-parameter model should work reasonably well for most response functions found in practice.

Although this analysis gives us a parametric model, it also gives us some insight into the nature of the regularity of response functions. A number of current calibration methods use a global regularity parameter λ to constrain the response rather than explicitly parameterize the response [6], [23], [30]. Our analysis suggests that the regularity of response functions varies since the mean and the principal component curves are quite smooth toward the middle of the domain with higher derivatives occurring near the endpoints. This suggests it would be more appropriate to use a varying regularity constraint to restrict the approximation.

To approximate a new response function f in W_{RF} with an M -parameter EMoR, we project $f - f_0 \in W_0$ into V_M . Let $H := [h_1 \dots h_M]$ be the matrix whose columns are the first M unit eigenvectors. Then, the EMoR approximation \tilde{f} to the response curve f is

$$\tilde{f} = f_0 + Hc, \quad (4)$$

7. We restrict our attention to functions which are square integrable on $[0, 1]$.




Fig. 5. (a) The mean of the camera responses in DoRF in log-space. This is equivalent to the geometric mean of the curves in DoRF. (b) Four principal components in log-space. (c) A plot showing the percentages of the energies in log-space. This shows a three-dimensional subspace captures more than 99.6 percent of the energy in log-space.

where $c = H^T(f - f_0)$ are the model coefficients. We also performed PCA in log-space. In that case, the mean curve in log-space $f_{\log,0}$, corresponds to the geometric mean in W_{RF} , shown in Fig. 5a. The first four principal components $h_{\log,1}, \dots, h_{\log,4}$ in log-space of the DoRF curve are plotted in Fig. 5b as they appear in W_{RF} after being exponentiated. With this model, a response function f is approximated

$$f(E) \approx (e^{f_{\log,0}(E)}) \prod_{n=1}^M (e^{h_{\log,n}(E)})^{c_n}. \quad (5)$$

Since W_{\log} is a vector space, the isolation of the mean $f_{\log,0}(E)$ is not necessary the way it is in the linear case. Isolating the mean does, however, improve the approximation. The log model gives comparable results to the simpler technique of finding a linear basis directly in W_{RF} . Nevertheless, there may be cases where it is better to use a log model. For example, consider images coming from a digital capture card connected to a video camera. Some cards apply a gamma to the camera response. Thus, the response of the complete imaging system is a composite of one response followed by the other. In log-space application of gamma corresponds to multiplication by a constant. Thus, a linear model in log-space is able to represent such composite responses easily. However, we should note that PCA is no longer optimal with regard to least-squares in W_{RF} for the log model. Moreover, the log-model is very sensitive for small irradiance values.

6 IMPOSING MONOTONICITY

The EMoR approximation \tilde{f} from (4) satisfies the constraint that the function lies in the hyperplane $f \in W_1$. Nevertheless, as discussed in Section 2, functions in the theoretical space W_{RF} must also be monotonic. In this section, we will describe finding a *monotonic* EMoR approximation in terms of finding a least squares approximation subject to a set of inequalities.

We first note that the space of monotonic EMoR approximations $\mathcal{L}_M := (f_0 + V_M) \cap W_{RF}$ is convex. This is because it is an intersection of convex sets. Suppose f is the true response curve to be approximated. The convexity of \mathcal{L}_M implies that there is a unique closest point, \tilde{f}^{mon} , in \mathcal{L}_M to f . Here, we measure distance in terms of the norm in L^2 . Thus, \tilde{f}^{mon} is the monotonic least squares EMoR approximation to f .

We note that the function \tilde{f}^{mon} is monotonic if its derivative is positive. If D is the discrete derivative matrix, then the monotonicity constraint becomes a system of inequalities, which in matrix form can be written $D\tilde{f}^{mon} \geq 0$. Let H be the matrix with the first M PCA eigenvectors as columns. Then, \tilde{f}^{mon} will be of the form $\tilde{f}^{mon} = f_0 + H\hat{c}$, where the coefficient vector \hat{c} is determined as $\hat{c} = \arg \min_c \|Hc - f - f_0\|^2$, subject to the constraint

$$DH\hat{c} \geq -f_0. \quad (6)$$

Thus, finding \hat{c} turns into a standard problem of quadratic programming. More details on how we performed this minimization are given in the Appendix. We note that since the function log is monotonic, a function is monotonic in log-space if and only if it is monotonic. Thus, the monotonicity may be also be imposed in log-space using quadratic programming.

As an example, we apply quadratic programming to approximate all the DoRF curves in the 2-parameter approximation space $\mathcal{L}_2 := (f_0 + V_2) \cap W_{RF}$. In this space, it is possible to visualize both the space of response functions W_{RF} and how the DoRF curves appear within it. In Fig. 6, \mathcal{L}_2 is the pie shaped region and the black dots and circles are all the curve approximations to DoRF curves \mathcal{L}_2 . For reference, the approximation to the gamma curves in DoRF are shown as circles.

The space of curves V_2 is two-dimensional and can be parameterized by the derivatives of the curves at 0 and 1. Monotonicity bounds the coefficients and thus the derivatives. From Fig. 6, one can see that the bounds on the derivatives at 0 and 1 are approximately 44 and 5.1, respectively. As pointed out in [26], we can use the derivative of the response function to estimate the signal-to-noise ratio (SNR) of any given gray-level. Note also that there are large sections of the pie-shaped convex region where no curves (dots or circles) appear. Although this tells us nothing about dimensionality, it does indicate that the space occupied by the diverse set of 201 response curves in DoRF is quite localized in \mathcal{L}_2 .

7 EVALUATING THE EMOR MODEL

The various approximation models described in Section 3 can fit increasingly complex response functions at the cost of using many parameters (model coefficients). What distinguishes these models from each other is the rate and manner with which the quality of the approximation varies

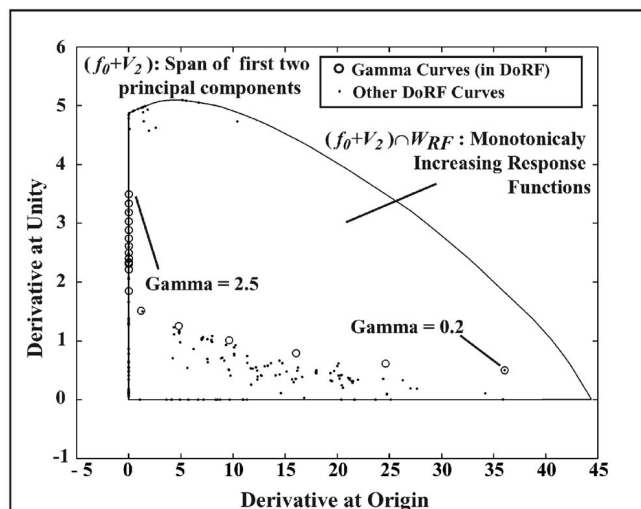


Fig. 6. The span of the eigenvectors associated with the two largest eigenvalues. We parameterize this span by the derivatives of the response curves at 0 and 1. The pie-shaped region (bounded by the solid curve) is the intersection with the cone of monotonic functions. The circles are approximations to the gamma curves and the black dots are approximations to the remaining curves in DoRF. The unoccupied areas in the pie-shaped region represent monotonic functions that typically are not camera response functions in practice.

with the number of parameters. To see this in the case of EMoR, we chose two curves from the DoRF database which were difficult to fit. Fig. 7 shows approximations of these two curves with the number of parameters $M = 1, 3, 5, 7,$ and 9 . Even the low-parameter approximations follow the curves grossly. With five parameters, it is hard to distinguish the response curves from the approximations. The approximations for $M = 11$ are almost identical to the original curves. These worst-case curves show qualitatively how fitting improves with the number of parameters.

We conducted an extensive quantitative evaluation of the EMoR model. Table 1 shows how accuracy increases with dimensionality. We used EMoR to approximate the 175 training curves as well as the 26 testing curves described in Section 4. We used four metrics with each set to evaluate the results. The results are shown in Table 1. The error values based on root-mean-square error (RMSE) appear in rows labeled **RMSE Case**. We compute the **Mean Error** by averaging the RMSE over all curves in each set (training and testing). The largest RMSE over all curves in the set gives the value called **Worst Curve**. The errors in

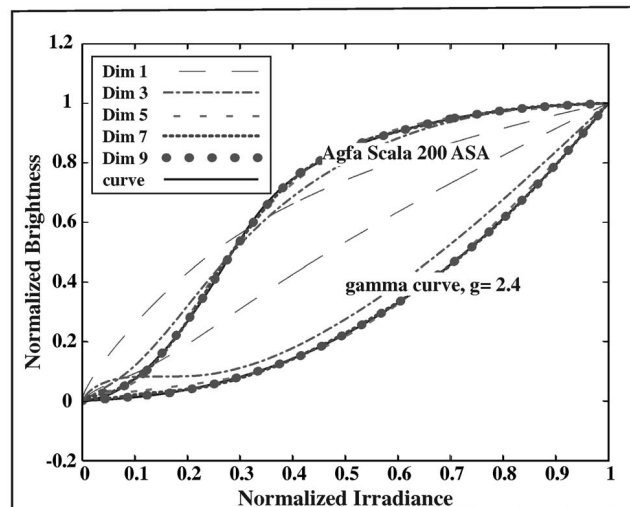


Fig. 7. A qualitative illustration of how the fit of the EMoR model improves with the numbers of model parameters. Here, we show two of the most difficult responses in the DoRF database. For each of these responses, approximation curves with one, three, five, seven, and nine parameters are shown. Even with five parameters, the approximation is quite good. With 11 parameters, there is little difference between the approximate and actual curves.

rows labeled **Disparity Case** are computed from the maximum disparity of the fit and the original curve. The columns of Table 1 correspond to the dimensions (parameters) used for the EMoR model. Note that most curves are well-approximated using an EMoR model with just three parameters.

Taking $-\log_2$ of the entries in Table 1 gives the accuracy in bits. The number of parameters needed for acceptable accuracy depends on the application. For example, suppose we wish to construct a mosaic by blending a set of images taken with an 8-bit camera, and an error of four gray-levels is acceptable. Choosing three parameters gives an RMSE Case/Mean Error accuracy of 6.8. In an application that is more sensitive to errors, such as stereo, choosing six parameters exceeds 9.0 bits of accuracy using the same measure. Some algorithms may require very accurate scene radiance measurements. The Disparity Case/Mean Error metric gives the average of the worst errors. This measure indicates that using 11 parameters ensures 8.3 bits of accuracy.

TABLE 1
An Evaluation of the Performance of the EMoR Model as the Number of Parameters Increase

		Dim	1	2	3	4	5	6	7	8	9	10	11
Training data	RMSE	Mean Error	4.12E-02	1.94E-02	8.87E-03	4.15E-03	2.82E-03	1.91E-03	1.58E-03	1.15E-03	9.10E-04	7.60E-04	6.02E-04
	Case	Worst Curve	2.34E-01	1.37E-01	6.49E-02	2.48E-02	1.76E-02	1.29E-02	1.08E-02	8.08E-03	6.39E-03	4.53E-03	2.94E-03
	Disparity	Mean Error	9.07E-02	4.77E-02	2.60E-02	1.51E-02	1.05E-02	8.09E-03	7.46E-03	4.96E-03	4.30E-03	3.95E-03	3.12E-03
	Case	Worst Curve	3.43E-01	2.41E-01	2.06E-01	1.64E-01	1.31E-01	1.02E-01	9.81E-02	5.37E-02	4.05E-02	3.93E-02	2.56E-02
Testing data	RMSE	Mean Error	4.00E-02	1.73E-02	6.27E-03	2.54E-03	1.77E-03	1.07E-03	9.55E-04	6.37E-04	5.23E-04	4.74E-04	4.00E-04
	Case	Worst Curve	2.04E-01	1.06E-01	3.79E-02	9.47E-03	9.38E-03	3.56E-03	2.78E-03	1.90E-03	1.75E-03	1.62E-03	1.60E-03
	Disparity	Mean Error	8.28E-02	3.62E-02	1.74E-02	1.02E-02	7.58E-03	5.65E-03	5.49E-03	3.14E-03	2.75E-03	2.62E-03	2.04E-03
	Case	Worst Curve	2.88E-01	1.56E-01	6.75E-02	3.19E-02	2.23E-02	2.16E-02	2.15E-02	1.61E-02	1.37E-02	1.36E-02	8.21E-03

The model was used to approximate curves in our training set of 175 curves and testing set of 26 curves from DoRF. In the **RMSE Case**, the **Mean Error** is the RMSE averaged over all the curves in each set. The largest RMSE for the set is listed in the row labeled **Worst Curve**. The **Disparity Case** uses the maximum disparity between approximated and actual curves. In this case, the **Mean Error** and the **Worst Curve** values are the mean disparity and maximum disparity computed over all the curves in the set. Most curves are well-approximated using only three parameters.

TABLE 2
Table Showing RMSE of Various Approximation Models Averaged over the Testing Curves

model \ dim	RMSE case						
	1	2	3	4	5	6	7
gamma	3.46E-02	N. A.	N. A.	N. A.	N. A.	N. A.	N. A.
polynomial	7.37E-02	3.29E-02	1.71E-02	1.06E-02	6.93E-03	4.95E-03	3.65E-03
trigonometric	6.83E-02	3.91E-02	2.58E-02	1.89E-02	1.44E-02	1.16E-02	9.46E-03
EMOR	4.00E-02	1.73E-02	6.27E-03	2.54E-03	1.77E-03	1.07E-03	9.55E-04

To Compute Accuracy in Bits, Take $-\log_2$ of the average RMSE. EMoR clearly outperforms all the other models.

TABLE 3
Table Showing the Maximum Disparity between Curves in the Testing Set and Their Model Approximations

model \ dim	Disparity case						
	1	2	3	4	5	6	7
gamma	2.52E-01	N. A.	N. A.	N. A.	N. A.	N. A.	N. A.
polynomial	4.22E-01	2.90E-01	2.12E-01	1.51E-01	1.16E-01	8.96E-02	7.55E-02
trigonometric	4.12E-01	3.30E-01	2.74E-01	2.32E-01	2.01E-01	1.77E-01	1.57E-01
EMOR	3.04E-01	1.38E-01	8.54E-02	4.90E-02	3.15E-02	2.25E-02	1.86E-02

Again, to compute the accuracy in bits, we take $-\log_2$ of each table entry. Once again, EMoR has the best performance.

TABLE 4

An Evaluation of the Performance of the Log Version of the EMoR Model as the Number of Parameters Increase as in Table 1

		Dim	1	2	3	4	5	6	7	8	9	10	11
Training data	RMSE Case	Mean Error	7.04E-02	1.71E-02	8.83E-03	4.99E-03	3.44E-03	2.80E-03	2.52E-03	1.67E-03	1.36E-03	1.79E-03	9.54E-04
		Worst Curve	2.10E-01	1.27E-01	3.55E-02	3.12E-02	2.13E-02	1.99E-02	1.93E-02	1.23E-02	7.80E-03	6.55E-02	5.67E-03
	Disparity Case	Mean Error	1.11E-01	3.44E-02	1.85E-02	1.11E-02	8.14E-03	6.58E-03	5.77E-03	4.12E-03	3.36E-03	4.14E-03	2.49E-03
		Worst Curve	3.77E-01	2.93E-01	1.46E-01	1.08E-01	8.18E-02	5.20E-02	5.06E-02	3.23E-02	2.04E-02	1.23E-01	1.91E-02
Testing data	RMSE Case	Mean Error	5.92E-02	1.22E-02	6.52E-03	3.27E-03	2.06E-03	1.49E-03	1.36E-03	8.66E-04	7.50E-04	6.28E-04	6.18E-04
		Worst Curve	1.50E-01	6.02E-02	2.84E-02	1.65E-02	7.98E-03	5.94E-03	4.62E-03	2.54E-03	2.24E-03	1.70E-03	1.69E-03
	Disparity Case	Mean Error	9.11E-02	2.19E-02	1.22E-02	6.70E-03	4.40E-03	3.22E-03	2.94E-03	2.15E-03	1.94E-03	1.72E-03	1.71E-03
		Worst Curve	2.45E-01	1.01E-01	5.26E-02	3.17E-02	1.65E-02	1.16E-02	9.41E-03	6.67E-03	6.31E-03	6.41E-03	6.08E-03

Entries in the table show the log-space version of EMoR gives comparable results to the simpler version of EMoR.

We evaluated EMoR by comparing its performance to other approximation models, when monotonicity is imposed in all cases. These include the gamma function, E^γ , as well as the polynomial and trigonometric approximation models described at the end of Section 3. The gamma curve has only one parameter, given our normalizations. Table 2 summarizes our results for the DoRF testing curves. The accuracy of the gamma curve model, using the RMSE averaged across the testing curves, is 4.86 bits. For one-parameter models, the gamma curves are superior to other models. All models, except the gamma curve, may be made more accurate by using more parameters. As the number of parameters increases, the EMoR-based monotonic fit significantly outperforms the other models. Table 3 shows that EMoR is also superior to other models when accuracy is measured in terms of the maximum disparity of the worst curve in the testing set.

We also evaluated the empirical approximation model of (5) using log-space. Table 4 is identical to Table 1, except the DoRF curves were transformed to log-space. The model curves were exponentiated as in (5) before the errors were measured. The errors are comparable but more often larger than those in Table 4. This is an indication that approximating a response in log-space does not minimize the least-square error for curves in W_{RF} . Since PCA in log-space does not minimize the least square error in W_{RF} , we are not even guaranteed that increasing the number of parameters will yield a better least squares answer. Nevertheless, as pointed out in Section 5, the model in log-space is advantageous

when the imaging system's response consists of a camera response composed with a gamma response. With either the linear or log-space EMoR model, we can achieve a good monotonic approximation with a small number of parameters. Our evaluation shows that our empirical models perform better than those based on polynomials, trigonometric functions, or gamma curves.

8 CAMERA RESPONSE FROM SPARSE SAMPLES

The most popular way to estimate a camera's response function is by imaging a color chart of known reflectances, such as the Macbeth chart [5]. The Macbeth chart includes six patches with known reflectances going from white through gray to black. Typically, one applies standard interpolation to these points to obtain a continuous response function. There is no guarantee that the interpolated values correspond to the actual response function of the camera.

The EMoR model enables us to obtain accurate interpolations from chart measurements. Fig. 8 shows interpolation results obtained by fitting four different 3-parameter models (including EMoR) to the sparse response samples obtained from an image of a Macbeth chart taken using a Nikon 990 digital camera.⁸

⁸ The Nikon 990 camera was not part of the training or testing curves in DoRF.

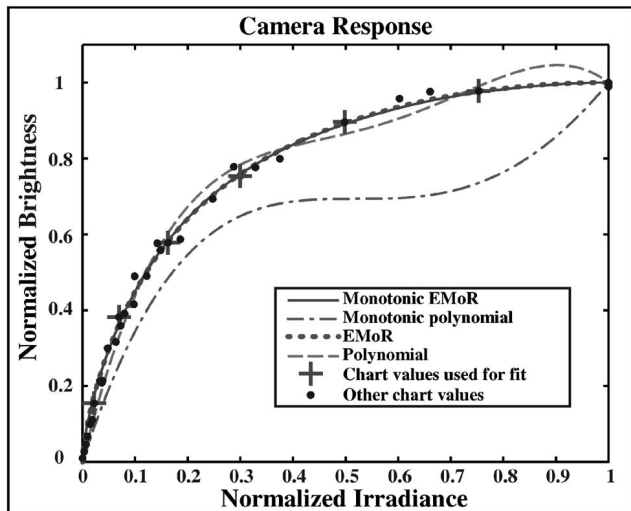


Fig. 8. The response curve of a Nikon 990 camera interpolated from sparse samples, obtained using a Macbeth chart, using 3-parameter EMoR and polynomial models both with and without monotonicity imposed. Images of the chart taken with the same camera at different exposures provide the additional measurements (ground truth) used to estimate the accuracies of the interpolations. The monotonic EMoR has the smallest RMS error.

From the chart and its image, we have six normalized irradiance values⁹ E_1, \dots, E_6 and corresponding intensity values B_1, \dots, B_6 . The three coefficients for the EMoR model are computed using (4), where the matrix H comes from the EMoR basis function h_1, h_2 , and h_3 evaluated at E_1, \dots, E_6 . Similarly, we take the first three basis vectors for the polynomial model evaluated at E_1, \dots, E_6 to obtain H and compute the three coefficients of that model. We also computed the first three coefficients of the EMoR and the polynomial model with monotonicity imposed, using the method described in Section 6.

The interpolations obtained from the different models were evaluated (using RMSE) against many more chart measurements obtained by simply changing the camera's exposure. For the 3-parameter EMoR, the RMSE was 0.11, while for the polynomial model, it was 0.12. Moreover, the polynomial fit is already beginning to exhibit the kind of oscillations one expects from overfitting. These oscillations worsen as the number of parameters increases. This is because we only fit using six data points. When we constrain the fits to be monotonic, the RMSE for the EMoR model is 0.11 and the RMSE for the polynomial model is 0.57. In summary, the EMoR model enables accurate reconstructions of response curves from very few samples.

9 RESPONSE FROM MULTIPLE IMAGES

A number of algorithms recover a camera's response from multiple images of an arbitrary static scene taken using different exposures [6], [23], [22], [26], [30]. These algorithms recover the inverse response function $f^{-1} = g$, where $g(B) = E$. Since an inverse response function has all the

9. We estimated the normalization constant such that $E_{\text{MAX}} = 1$ separately, so in fact, we are using seven samples. Alternately, we can choose an exposure so that E_6 is close to, but not saturated. We then obtain a response where we treat the unnormalized brightness value associated to that patch, as if it were the saturation value.



Fig. 9. Three images of a static scene taken with a Nikon 990 Coolpix camera using exposures e , $2e$, and $4e$ (left to right). These images were used to recover the inverse response of the camera (see Fig. 10).

properties of a response function, we can apply PCA to the inverses of the curves in DoRF to get an EMoR representation of the inverse camera response. We write $g(B) = g_0(B) + \sum_n^M c_n h_n^{\text{inv}}(B)$ in terms of the mean g_0 and the eigenvectors h_n^{inv} of the covariance matrix of the inverse curves.

Now, suppose two images of the same scene are captured with exposures e and $k \times e$, where k is the ratio of exposures. Suppose the images are registered. If a response B_a at a point in one image corresponds to a response B_b in the second image, then their irradiances must satisfy $g(B_a) = k g(B_b)$. Since the equations $g(B_a) - k g(B_b) = 0$ are linear in the coefficients c_n , the coefficients can be found using least-squares techniques, when k is known.

Using this approach, we recovered the response of the Nikon 990 Coolpix camera from the three images¹⁰ of a scene shown in Fig. 9, taken using exposures e , $2e$, and $4e$ (i.e. $k = 2$). The monotonic EMoR fit (using just three parameters) is shown as a solid curve in Fig. 10. For comparison, the polynomial method of Mitsunaga and Nayar [26], and the log method of Debevec and Malik [6] is also shown. Measurements from the Macbeth chart (black dots) are included as ground truth. All the recovered curves are reasonable fits although only the EMoR fit is monotonic. The EMoR fit was found to be closest to the ground truth (chart data).

10 CONCLUSION

Camera response curves represent the critical link between image measurements and scene radiances. Our goal is to provide a simple model which takes us from image intensity to scene radiance for the entire imaging system. To do this, we began with by analyzing the a priori properties of response curves. From these properties, we derived the theoretical space of all camera response functions. We have shown that all responses must lie within a convex set that results from the intersection of a hyperplane and the positive cone of monotonic functions. We used this space to formulate two general approximation models for camera responses. One set of models directly describes response functions with a linear combinations of component functions. This set of models subsumes previously used ones such as the polynomial model, as well as others such as the trigonometric model. A second set of models represents responses in log-space.

To fully exploit our theoretical insights, we created a database, DoRF, of 201 real-world response functions. We used the empirical data from DoRF and the general approximation model from our theoretical analysis to develop a powerful approximation model for responses

10. "Ernie" image used courtesy of Sesame Workshop.

called EMoR. We used several measures to show that the EMoR model performs far better than other models used in the literature.

As a basis for future work, we note that the same methodology can be applied to spectral response curves. Many film and camera manufacturers have published spectral response curves. The constraint for these response curves is positivity rather than monotonicity. Like monotonicity, the positivity constraint can be expressed as a system of inequalities. Therefore, given an empirical linear model, it is possible to use the methods of quadratic programming to determine coefficients for an empirical model for the spectral response that maintains the positivity constraint.

In this work, we showed two example applications of the EMoR model. The first used a few patches on a reflectance chart to fully recover the response curve of the camera. The second used EMoR to recover the camera response from three images of an arbitrary scene taken with different exposures. Our experimental results show that the EMoR model provides an accurate and efficient low-parameter model of real-world camera responses. The DoRF database and the EMoR model can be downloaded at <http://www.cs.columbia.edu/CAVE>.

APPENDIX

FORCING MONOTONICITY BY QUADRATIC PROGRAMMING

In Section 6, we outline how the EMoR model can be used to approximate a response function while ensuring it is monotonic. The constrained approximation may be posed as a standard problem of quadratic programming. To see this we will show that subject to a linear inequality $A_{ineq}s \leq b_{ineq}$ our problem is to minimize a quadratic objective function $G(s)$.

This quadratic objective function comes from recalling that in Section 6, we showed that for the parameters $s = (s_1, \dots, s_M)$, the squared approximation error is given by $\|Es - v\|^2$, where E is the matrix whose columns are basis vectors of W_0 , and where $v = f - f_0$, with f the function we are trying to approximate, and f_0 the 0th order approximation. Therefore, minimizing the squared error is the same as minimizing the quadratic objective function, $G(s) := \frac{1}{2}s^T Hs + d^T s$, where $H := E^T E$ and $d := E^T v$.

We impose the monotonic constraint on $\tilde{f}^{mon} = f_0 + Es$ by saying that the derivative of the approximation is positive. This is a linear condition. Using a dense sampling of the domain, (E_1, \dots, E_P) , we represent a function \tilde{f}^{mon} , by the vector $(\tilde{f}^{mon}(E_1), \dots, \tilde{f}^{mon}(E_P))^T$. We compute the discrete derivative $D\tilde{f}^{mon}$, where D is an $(P-1) \times P$ matrix and with $D_{n,n} := -1/Q_n$, $D_{n,n+1} := 1/Q_n$ and zero elsewhere, where $Q_n = E_{n+1} - E_n$. Then, $D\tilde{f}^{mon} \geq 0$ can be written as $A_{ineq}s \geq b_{ineq}$, where $A_{ineq} := DE$, and $b_{ineq} := -Df_0$. Note, we do not have to constrain the boundary points.¹¹

To perform the optimization used Matlab 6.0 function *quadprog*, which efficiently solves this problem using an active set method described in [11], [12]. The algorithm requires an initial guess for the solution for which we use

11. Recall that $f_0(0) = 0$ and $f_0(1) = 1$. Moreover, all elements of W_0 are functions which vanish at the endpoints. All elements of $f_0 + W_0$ and, thus, $f_0 + V_m$, satisfy the boundary conditions.

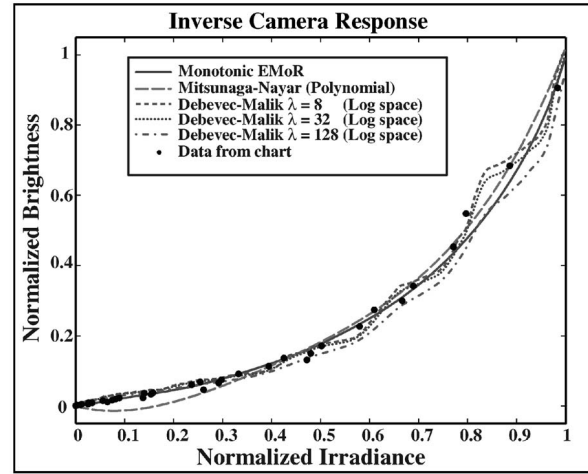


Fig. 10. Inverse response curves recovered from the images in Fig. 9 using the monotonic EMoR model (with three parameters), the Mitsunaga-Nayar polynomial model, and the Debevec-Malik log model (with three smoothing parameters, $\lambda = 8, 32, 128$). The dots correspond to ground truth obtained by calibration using a Macbeth reflectance chart.

the EMoR estimate \tilde{f} given by (4). If \tilde{f} satisfies the constraints, then it is already monotonic and no further optimization is needed. If not, we note that trivially f_0 satisfies the constraints; it is monotonic. Thus, there is some J such that the convex combination $(1 - 2^J)m + 2^J \tilde{f}$ is an initial solution to the constraints. Since, in practice, \tilde{f} is close to monotonic, usually $J = 1$ or $J = 2$.

ACKNOWLEDGMENTS

This work was completed with support from a US National Science Foundation ITR Award (IIS-00-85864) and a grant from the Human ID Program: Flexible Imaging over a Wide Range of Distances Award No. N000-14-00-1-0929. The authors thank Sesame Workshop for granting permission for use of the “Ernie” image.

REFERENCES

- [1] N. Asada, A. Amano, and M. Baba, “Photometric Calibration of Zoom Lens Systems,” *Proc. Int’l Conf. Pattern Recognition*, p. A73.7, 1996.
- [2] R. Basri and D.W. Jacobs, “Photometric Stereo with General, Unknown Lighting,” *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 374-381, 2001.
- [3] P.J. Burt and R.J. Kolczynski, “Enhanced Image Capture through Fusion,” *Proc. Int’l Conf. Computer Vision*, pp. 173-182, 1993.
- [4] B. Cabral, M. Olano, and P. Nemecek, “Reflection Space Image Based Rendering,” *Computer Graphics, Proc. SIGGRAPH*, pp. 165-170, 1999.
- [5] Y.C. Chang and J.F. Reid, “RGB Calibration for Color Image-Analysis in Machine Vision,” *IEEE Trans. Image Processing*, vol. 5, no. 10, pp. 1414-1422, Oct. 1996.
- [6] P.E. Debevec and J. Malik, “Recovering High Dynamic Range Radiance Maps from Photographs,” *Computer Graphics, Proc. SIGGRAPH*, pp. 369-378, 1997.
- [7] P.E. Debevec, “Rendering Synthetic Objects into Real Scenes,” *Computer Graphics, Proc. SIGGRAPH*, pp. 189-198, 1998.
- [8] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, second ed. New York: Wiley, 2000.
- [9] H. Farid, “Blind Inverse Gamma Correction,” *IEEE Trans. Image Processing*, vol. 10, no. 10, pp. 1428-1433, Oct. 2001.
- [10] G.D. Finlayson, S.D. Hordley, and P.M. Hubel, “Color by Correlation: A Simple, Unifying Framework for Color Constancy,” *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 23, no. 11, pp. 1209-1221, Nov. 2001.

- [11] P.E. Gill, W. Murray, M.A. Saunders, and M.H. Wright, "Procedures for Optimization Problems with a Mixture of Bounds and General Linear Constraints," *ACM Trans. Math. Software*, vol. 10, no. 3, pp. 282-298, 1984.
- [12] P. Gill, W. Murray, and M. Wright, *Numerical Linear Algebra and Optimization*. vol. 1, Redwood City, Calif.: Addison-Wesley, 1991.
- [13] M. Grossberg and S. Nayar, "What Can Be Known About the Radiometric Response Function from Images?" *Proc. European Conf. Computer Vision*, pp. 189-205, 2002.
- [14] M. Grossberg and S. Nayar, "What Is the Space of Camera Response Functions?" *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2003.
- [15] G. Healey and R. Kondepudy, "Modeling and Calibrating CCD Cameras for Illumination-Insensitive Machine Vision," *SPIE Proc., Optics, Illumination, and Image Sensing for Machine Vision VI*, vol. 1614, pp. 121-132, 1992.
- [16] G. Healey and R. Kondepudy, "Radiometric CCD Camera Calibration and Noise Estimation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 16, no. 3, pp. 267-276, Mar. 1994.
- [17] B.K.P. Horn and M.J. Brooks, *Shape from Shading*. MIT Press, 1989.
- [18] Eastman Kodak, *Student Filmmaker's Handbook*, 2002.
- [19] E.H. Land and J.J. McCann, "Lightness and Retinex Theory" *J. Optical Soc. of Am.*, vol. 61, no. 1, pp. 1-11, 1971.
- [20] Q.T. Luong, P. Fua, and Y. Leclerc, "The Radiometry of Multiple Images," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 1, pp. 19-33, Jan. 2002.
- [21] B.C. Madden, "Extended Intensity Range Image," Technical Report 366, Grasp Lab, Univ. of Pennsylvania, 1993.
- [22] S. Mann and R. Picard, "Being 'Undigital' with Digital Cameras: Extending Dynamic Range by Combining Differently Exposed Pictures," *Proc. IS&T, Soc. for Imaging Science and Technology 46th Ann. Conf.*, pp. 422-428, 1995.
- [23] S. Mann, "Comparametric Equations with Practical Applications in Quantigraphic Image Processing," *IEEE Trans. Image Processing*, vol. 9, no. 8, pp. 1389-1406, Aug. 2000.
- [24] S. Mann, "Comparametric Imaging: Estimating Both the Unknown Response and the Unknown Set of Exposures in a Plurality of Differently Exposed Images," *Proc. Conf. Computer Vision and Pattern Recognition*, Dec. 2001.
- [25] S. Marschner, S. Westin, E. Lafortune, and K. Torrance, "Image-Based BRDF Measurement," *Applied Optics*, vol. 39, no. 16, 2000.
- [26] T. Mitsunaga and S.K. Nayar, "Radiometric Self Calibration," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, vol. 2, pp. 374-380, June 1999.
- [27] S. Narasimhan and S. Nayar, "Removing Weather Effects from Monochrome Images," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 186-193, 2001.
- [28] S.K. Nayar, K. Ikeuchi, and T. Kanade, "Shape from Interreflections," *Int'l J. Computer Vision*, vol. 6, no. 3, pp. 173-195, Aug. 1991.
- [29] R. Ramamoorthi and P. Hanrahan, "A Signal-Processing Framework for Inverse Rendering," *Computer Graphics, Proc. SIGGRAPH*, pp. 117-128, 2001.
- [30] Y. Tsin, V. Ramesh, and T. Kanade, "Statistical Calibration of the CCD Imaging Process," *Proc. Int'l Conf. Computer Vision*, pp. 480-487, 2001.
- [31] R.J. Woodham, "Photometric Method for Determining Surface Orientation from Multiple Images," *OptEng*, vol. 19, no. 1, pp. 139-144, Jan. 1980.
- [32] Y. Yu, P.E. Debevec, J. Malik, and T. Hawkins, "Inverse Global Illumination: Recovering Reflectance Models of Real Scenes from Photographs," *Computer Graphics, Proc. SIGGRAPH*, pp. 215-224, 1999.
- [33] R. Zhang, P.S. Tsai, J.E. Cryer, and M. Shah, "Shape from Shading: A Survey," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 21, no. 8, pp. 690-706, Aug. 1999.
- [34] T. Zickler, P.N. Bellhumeur, and D.J. Kriegman, "Helmholtz Stereopsis: Exploiting Reciprocity for Surface Reconstruction," *Proc. European Conf. Computer Vision*, 2002.



Michael D. Grossberg received the PhD degree in mathematics from the Massachusetts Institute of Technology in 1991. He is a research scientist with the Columbia Automated Vision Environment (CAVE) at Columbia University. His research in computer vision has included topics in the geometric and photometric modeling of cameras and analyzing features for indexing. Dr. Grossberg was a lecturer in the Computer Science Department at Columbia University. He was also a Ritt assistant professor of mathematics at Columbia University. He has held postdoctoral fellowships at the Max Plank Institute for Mathematics in Bonn, and the Hebrew University in Jerusalem. He has authored and coauthored papers that have appeared in ICCV, ECCV, and CVPR. He has filed several US and international patents for inventions related to computer vision. He is a member of 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 Honorable Mention Award at the 2000 IEEE CVPR Conference, the David Marr Prize at the 1995 ICCV held in Boston, Siemens Outstanding Paper Award at the 1994 IEEE CVPR Conference held in Seattle, 1994 Annual Pattern Recognition Award from the Pattern Recognition Society, 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 holds several US and international patents for inventions related to computer vision and robotics. Dr. Nayar was the recipient of the David and Lucile Packard Fellowship for Science and Engineering in 1992, the National Young Investigator Award from the US National Science Foundation in 1993, and the Excellence in Engineering Teaching Award from the Keck Foundation in 1995.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.