Determining the Camera Response from Images: What Is Knowable?

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

Abstract—An image acquired by a camera consists of measured intensity values which are related to scene radiance by a function called the camera response function. Knowledge of this response is necessary for computer vision algorithms which depend on scene radiance. One way the response has been determined is by establishing a mapping of intensity values between images taken with different exposures. We call this mapping the intensity mapping function. In this paper, we address two basic questions. What information from a pair of images taken at different exposures is needed to determine the intensity mapping function? Given this function, can the response of the camera and the exposures of the images be determined? We completely determine the ambiguities associated with the recovery of the response and the ratios of the exposures. We show all methods that have been used to recover the response break these ambiguities by making assumptions on the exposures or on the form of the response. We also show when the ratio of exposures can be recovered directly from the intensity mapping, without recovering the response. We show that the intensity mapping between images is determined solely by the intensity histograms of the images. We describe how this allows determination of the intensity mapping between images without registration. This makes it possible to determine the intensity mapping in sequences with some motion of both the camera and objects in the scene.

Index Terms—Calibration, histogram, response function, ambiguities, illumination, radiometry, comparagram, dynamic range, intensity mapping, histogram specification, comparametric.

---◆---

1 DETERMINING THE CAMERA RESPONSE FROM DIFFERENT EXPOSURES

THE image a camera acquires consists of a collection of measurements we refer to as intensity values.[1] At a single point in the image, an intensity value is related to the scene radiance by a nonlinear function called the camera response function. We will assume that the response is the same for each point in the image. A typical camera response has a variation across the image which is linear in scene radiance. Once the response is found, this variation may be calibrated separately and removed [1]. By determining the response, or rather the inverse of the response, we can obtain scene radiance from image intensity.

Obtaining scene radiance is required for numerous applications in machine vision that determine properties of the scene from the physics of the interaction of light with scene objects. For example, Belhumeur and Kriegman require scene radiance from images to determine the illumination space of an object [3]. Marschner estimates the BRDF from images [20]. Shape from shading algorithms use scene radiance to determine the orientation of surface normals with respect to the illumination [27]. Photometric

stereo algorithms use changes in scene radiance between images to determine the scene structure and the illumination [22]. Color constancy algorithms use estimates of scene radiance to separate the illumination from the reflectance of the scene [14]. High dynamic range illumination maps using accurate wide field of view measurements of scene radiance permit the realistic merging of synthetic and real objects in an image [7]. Estimating the spectral response and the polarization of light from objects, using generalized mosaicing, requires obtaining the scene radiance from images [24]. In all these cases, one must find scene radiance from measured intensity values by determining the camera response of the imaging system.

The camera response function can be obtained by taking an image of a uniformly illuminated chart with patches of known reflectances, such as the Macbeth chart, as done in [5]. Nevertheless, placing a chart in the scene can be inconvenient or difficult in the field. For example, when images are taken with a camera attached to a remote mobile device. Additionally, changes in temperature alter the response function requiring frequent recalibration.

The problems associated with using charts have led researchers to develop methods to determine a camera's response from arbitrary scenes. Farid showed if one assumes the response has the form of a gamma curve, one can estimate the parameters of the curve by making assumptions on the statistics of scene radiance [9]. Rather than making assumptions about the statistics of scene radiance, a number of researchers obtained the response function from a sequence of images of the same scene taken at different exposures. By comparing corresponding intensity values between images in the sequence, Mann and

---

1. By intensity value, we simply mean the number the camera reports at a pixel. A typical color camera will report one number at each pixel for each channel (e.g., red, green, and blue). We treat these channels as giving three monochrome images.

● *The authors are with the Computer Science Department, Columbia University, New York, NY 10027.*
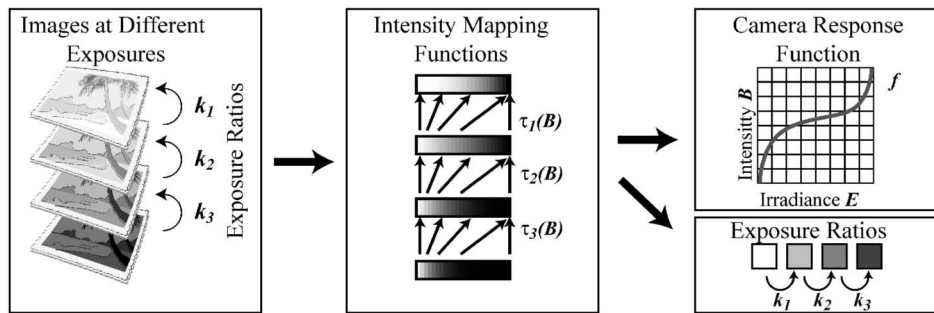*E-mail: {mdog, nayar}@cs.columbia.edu.*

Fig. 1. A diagram showing the two parts of recovering the response function and exposure ratios from a sequence of images taken at different exposures. All the information in the images relevant for chartless recovery is contained in the intensity mapping functions. These functions describe how an intensity in one image maps to an intensity in another image. Thus, the problem of recovering the camera response function falls into two parts: recovery of the intensity mapping functions from images and recovery of the camera response function and the exposure ratios from the intensity mapping functions.

Picard demonstrated constraints on the response of a camera in the pioneering paper [18]. These constraints allowed them to determine the response at a sequence of points. By using a parameterized model of this response, such as a gamma curve, they determined the response assuming known exposure ratios between the images.[2] Debevec and Malik also assumed the ratio of exposures is known, but they found the log of the inverse camera response without using a parametric form for the response [6]. Instead, they imposed a smoothness constraint on the response.

In between having assumed a specific parametric form and having assumed no parametric form, Mitsunaga and Nayar assumed that the inverse response function can be closely approximated by a polynomial [21]. They were then able to obtain the coefficients of the polynomial and the exposure ratios, from rough estimates of those ratios, by alternating between recovering the response and the exposure ratios. Tsin et al. [26] and, separately, Mann [17] determined both the response and exposure ratios by combining the iterative[3] approach from [21], with the nonparametric recovery in [6].

All chartless[4] methods base the recovery of the response function on the constraint that between a pair of images of a static scene, all irradiances change by the same exposure ratio [6], [17], [18], [21], [26]. The essential information from the images that all methods use is how intensity values in one image map to intensity values in another. We describe this mapping with a function called the intensity mapping function.[5] Fig. 1 shows that the recovery of the response decomposes into two parts: the recovery of the intensity mapping function from images and the recovery of the combination of the response and exposure ratios from the intensity mapping function.

2. In later work, Mann has suggested alternate parameterized models [16], as well as the use of nonparametric regression [17].

3. Iteration is unnecessary if the exposure ratios are known, as in earlier methods.

4. When we refer to "chartless" methods, we do not include Farid's method [9]. While it does not depend on knowing the values of scene reflectances, it makes assumptions on the higher order statistics of scene radiances.

5. Most methods use all pairs of intensity values at corresponding points in the images. Mann shows that the intensity mapping function, which he calls the comparametric function, summarizes this information [16].

In this paper, we address both parts of the recovery problem. We will first completely describe the limitations of the second part of this problem, that is, obtaining the response and exposure ratios from the intensity mapping function, by answering the following questions:

1. Given the exposure ratio of an image pair, is there a unique response function for each intensity mapping function?
2. Is it possible to recover the response function and the exposure ratios simultaneously and uniquely?
3. Can we recover the exposure ratio without recovering the response?

We show that, given fixed exposure ratios, different camera response functions give rise to the same intensity mapping function. We present a self-similar ambiguity which arises when recovering the response from an intensity mapping function. We show that either a careful choice of exposures or a priori assumptions on the response are necessary to break this ambiguity. Beyond this ambiguity, we describe an ambiguity to recovering the response function and exposure ratios together and prove it is the only ambiguity. Only by making assumptions on the response function can we expect a unique solution. We will show when it is possible to solve for the exposure ratio without determining the response. Having established the limits of determining the response function and exposure ratios from the intensity mapping function, we address the first part of the recovery problem:

4. What information is necessary to recover the intensity mapping function from images?

Previous work compared registered images which required a static scene and assumed restricted camera motion. We answer the above question by proving a theorem which relates the intensity mapping function to the histograms of the images. This implies that, in situations where the distribution of scene radiances remains almost constant between images, we can still recover the intensity mapping function even for images with scene and camera motion that cannot be registered. We verify the theorem by recovering the response function from a sequence of images with camera motion, motion of objects in the scene, and a combination of both. This type of

camera motion (where parallax occurs) and scene motion have not been handled before.

2 THE FUNDAMENTAL CONSTRAINT FOR CHARTLESS RECOVERY

To determine scene radiance L from an image, we first note that the scene radiance is proportional to image irradiance E [23]. The irradiance and radiance are related by $E = ePL$, where P is a factor due to the optics of the system and e is the exposure. For a simple system, $P = \cos^4 \alpha / c^2$, where α is the angle subtended by the principle ray from the optical axis and c is the focal length.[6] The exposure is given by $e = (\pi d^2)t$, where d is the size of the aperture and t is the time for which the photo-detector is exposed to the light. Even though e contains the integration time t, we can think of E as the image plane irradiance.

A function f, called *the camera response function*, relates the actual measured intensity value $B = f(E)$ at a photo-sensitive element to the image irradiance.[7] Imaging system designers often intentionally create a nonlinear response, for example, to compress the dynamic range. Since measured intensity indicates relative irradiance, we can assume that the response monotonically increases.[8] We normalize the domain of the response, which is the irradiance, to go from 0 to 1. We also normalize[9] the range of f, which is the intensity, so that $f(1) = 1$ and $f(0) = 0$. Up to this normalization, we can determine f if we take an image of a uniformly illuminated chart with known reflectance patches. Without a chart, we must find the constraints that permit us to extract f from images without assuming the knowledge of scene reflectances.

As a special case of how we obtain constraints, suppose we take two images of the same scene with different exposures e_1 and e_2. If the first image has image irradiance E_1 at a point and the corresponding point in the second image has the irradiance E_2, then $E_1/e_1 = E_2/e_2$. The exposure ratio $k := e_2/e_1$ expresses the relationship between irradiances in the two images, $E_2 = kE_1$.

Typically, the response is used to convert intensity to irradiance using the assumption that the response is monotonic and thus invertible. Therefore, we actually need to recover the inverse response $g := f^{-1}$ rather than the response itself. If we let $f(E_1) = B_1$ and $f(E_2) = B_2$, we have the equation

$$g(B_2) = kg(B_1) \qquad (1)$$

6. Details of P for a simple perspective camera can be found in Horn [12]. Whereas Mitsunaga and Nayar [21] discuss selection of pixels in the image where P is nearly constant, we will assume P is constant throughout the part of the image we analyze.

7. Treating intensity as a function of irradiance E is a simplification since irradiance is measured across all frequencies of light. Intensity is a function of irradiance weighted across frequencies by a distribution (in digital cameras called the quantum efficiency). Mann refers to this weighted irradiance as the quantimetric exposure [17].

8. If the response f is monotonically decreasing, as in negative films, we replace f with $1 - f$.

9. The minimum intensity value in a digital imaging system is often effectively greater than zero due to nonzero mean thermal noise called dark current. By taking an image with the lens covered, this effect may be estimated and subtracted.

in terms of the inverse response. All chartless methods base the recovery of the response and the exposure ratio on this constraint equation. In each pair of images, each corresponding pair of pixel intensity values gives one constraint with a fixed exposure ratio. Assuming the exposure ratio is known, and the inverse response is a polynomial, then (1) becomes linear in the coefficients of the polynomial. Mitsunaga and Nayar solve for these coefficients [21]. Debevec and Malik [6] and Mann [17] take the log of both sides of (1). Rather than start with a parameterized model of $\log g$, they evaluate it at a set of discrete intensity values and, therefore, treat it as a vector. By imposing a regularity condition on the discrete second derivatives of $\log g$, they are able to obtain a solution.

When we know the response but not the exposure ratio, we can solve (1) for the exposure ratio. Mitsunaga and Nayar [21] and Mann [17] use an iterative scheme in which they first solve for the inverse response with an initial guess for the exposure ratios. Updating their estimates, they iteratively solve for the exposure ratios and the inverse response.

3 INTENSITY MAPPING FUNCTIONS

The pairs of intensity measurements B_1 and B_2 at corresponding points in two different images of the same scene constitute all the information available from which to recover the response function in chartless recovery. Mann [17] pointed out that all this information is contained in a two variable cross-histogram he calls the *comparagram*. If (B_1, B_2) are any two pairs of intensity values, then the comparagram $J(B_1, B_2)$ is the number of pixels which have intensity value B_1 in the first image and B_2 at the corresponding point in the second image.

For real images, the correspondence between intensity values in one image and intensity values in another cannot be described by a function. The intensity values are not related by a single function because of factors such as noise, quantization of the intensity values, spatial quantization, and saturated pixels. Ignoring these factors for a moment, from (1), we see that a function τ should ideally relate the intensity values in the images

$$B_2 = \tau(B_1) := g^{-1}(kg(B_1)), \qquad (2)$$

which we call the *intensity mapping function*. This function describes how to map intensity values in one image into the second image. We can estimate the intensity mapping function from the comparagram. Once we estimate the intensity mapping τ, we have a modified version of (1), given by

$$g(\tau(B)) = kg(B). \qquad (3)$$

This equation has an advantage over (1): Because the function τ contains all the information about the intensity mapping between images, we can study the mathematical problem of existence and uniqueness of solutions to (3). To study solutions to (3), we must first derive some properties of intensity mapping functions. For example, we will show that τ is monotonic so τ^{-1} exists. Define $\tau^0(B) := B$, $\tau^n(B) := \tau(\tau^{n-1}(B))$, and $\tau^{-n}(B) := \tau^{-1}(\tau^{1-n}(B))$.

Theorem 1. Intensity Mapping Function Properties. *Let g be a smooth monotonically increasing function with smooth inverse. Suppose that $g(0) = 0$, $g(1) = 1$, and $k > 1$, then the function $\tau(B) := g^{-1}(kg(B))$ has the following properties:*

1. *$\tau(0) = 0$,*
2. *τ monotonically increases,*
3. *$B \leq \tau(B)$, and*
4. *$\lim_{n \to \infty} \tau^{-n}(B) = 0$.*

See Appendix A for the proof.

Assuming that $k > 1$ just means that we order our images so that exposure increases. To order the images, note that, when $k > 1$, then $g^{-1}(kE) \geq g^{-1}(E)$ since g^{-1} monotonically increases. In other words, every intensity value in one image maps to a larger value in the other image. Therefore, the intensity mapping function τ goes from the image with darker average pixel value, to the image with lighter average pixel value.[10]

4 SELF-SIMILAR AMBIGUITY

Each pair of images determines an intensity mapping function. In this section, we show that, even when the exposure ratio between the images is known, (3) has many solutions. Hence, the inverse camera response is not uniquely determined by the image pair. To understand why there are many solutions, we first note that, if the exposure ratio k is greater than 1, the intensity mapping function τ expands the B-axis as we see from Theorem 1.1. Thus, (3) can be interpreted as asserting that stretching the intensity axis using the intensity mapping and applying the inverse response g is the same as multiplying g by the exposure ratio between the images. The left side of the equation relates the value of g at points in $(\tau^{-1}(1), 1]$, to the value of g at points in $(\tau^{-2}(1), \tau^{-1}(1)]$ on the right side. Nothing, however, relates the value of g at points in $(\tau^{-1}(1), 1]$ to *each other*. This means as long as $g(1) = 1$, $g(\tau^{-1}(1)) = 1/k$ and g is continuous and monotonic, then g can have arbitrary values on $(\tau^{-1}(1), 1)$ and still be a solution to (3). We call this ambiguity to (3) the self-similar ambiguity.[11]

More formally, we state this by saying that we can build a solution g starting with any function $s(B)$ on $[\tau^{-1}(1), 1]$:

Theorem 2. The Self-Similar Ambiguity. *Suppose that τ satisfies the properties listed in Theorem 1. Suppose $s(B)$ is any continuous, monotonic function on the interval $[\tau^{-1}(1), 1]$ such that $s(1) = 1$, $s(\tau^{-1}(1)) = 1/k$. Then, s extends to a unique, continuous, and monotonic function g on $[0, 1]$ such that $g(B) = s(B)$ for $B \in [\tau^{-1}(1), 1]$ and g satisfies $g(\tau(B)) = kg(B)$, with $g(0) = 0$ and $g(1) = 1$. (See Appendix B for proof.)*

We can more easily understand this self-similar ambiguity by evaluating s at some sample points. Suppose we take three points $b_1, b_2, b_3 \in (\tau^{-1}(1), 1]$, as shown in Fig. 2a.
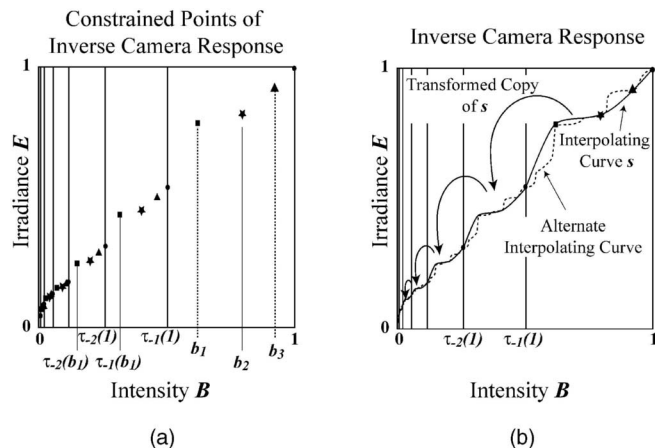

Fig. 2. (a) Graph showing the constraints of (3). The equation does not constrain the relative values of the inverse response within the interval $(\tau^{-1}(1), 1)$ such as at points b_1, b_2, and b_3. Nevertheless, using (3), we can generate a sequence of points from each such point, for example b_1. (b) Two arbitrary curves s on $(\tau^{-1}(1), 1]$ extended uniquely to $[0, 1]$. This means neither (3) nor its log tell us anything about the values of g in $(\tau^{-1}(1), 1]$. Unique solutions of g can only come from prior assumptions on the form or smoothness of g.

We can choose the values of s and, hence, g at these points essentially arbitrarily. The only restriction a priori is that s be monotonic and, thus, $s(\tau^{-1}(1)) = 1/k \leq s(b_n) \leq 1$. Each point $b_1 \in (\tau^{-1}(1), 1]$ gives rise to a sequence of points $b_1 \geq \tau^{-1}(b_1) \geq \tau^{-2}(b_1) \geq \ldots$. Equation (3) determines this sequence from b_1. It places no restrictions on the relationship between the values of s at points in $(\tau^{-1}(1), 1]$. In Fig. 2b, we see alternative solutions to (3).

The self-similar ambiguity presents a problem relating different sequences of points, each obtained by choosing a point b_1 and a value $s(b_1)$ in $(\tau^{-1}(1), 1]$. The choice of initial value $s(b_1)$ determines each sequence, however, only up to a single multiple. Thus, choosing the exposure ratio closer to 1 reduces this ambiguity since $1 - \tau^{-1}(1)$ is smaller.[12] Choosing the exposure ratio too close to 1, however, creates other problems. When the exposure ratio is close to one, the images will be very similar. Thus, the intensity mapping function will be close to the identity, as is clear from (2). Recovery of the inverse response g from this mapping will be very sensitive to noise. The recovery itself assumes that the scene does not change for different exposures. If the images are taken sequentially, even small changes in illumination may create problems in recovering the inverse response.

By assuming the inverse response has a particular form or by imposing a regularity constraint, we can break the self-similar ambiguity. For example, suppose we take a sequence of points in the plane

$$(1, 1), (\tau^{-1}(1), 1/k), (\tau^{-2}(1), 1/k^2), \ldots, (\tau^{-N}(1), 1/k^N).$$

If we assume that the inverse response is a low degree polynomial, then we can solve for the least squares polynomial fit for the points. This breaks the ambiguity

---

10. Note that we do not have to register images to compare the average pixel values.

11. Mann [17] has noted the relationship of his version of (3) to dynamical systems and that the equations allow "ripples" in the solutions.

12. Most consumer cameras do not permit fine adjustment of exposures. Nevertheless, it is often possible to achieve exposure ratios close to one by varying, for example, the aperture and integration time simultaneously.
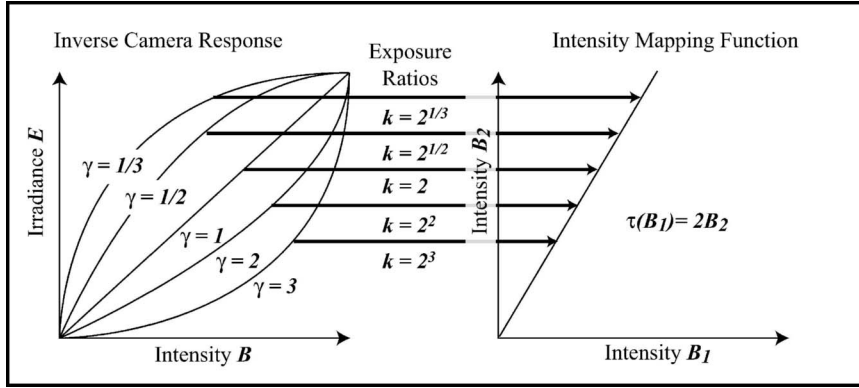
Fig. 3. Shown above is a graph of different response functions giving rise to the same intensity mapping function. The intensity mapping function $\tau(B) = 2B$ results from an inverse response function B^γ and exposure ratio between images of 2^γ, independent of γ. This shows we can only recover the inverse response of the camera and the exposure ratio simultaneously by making assumptions on the inverse response and the exposure ratio that break the ambiguity.

by allowing the polynomial to interpolate between $\tau^{-1}(1)$ and 1 based on the best fit to the rest of the points.

Clearly, multiple exposures cannot help if the same exposure ratios are used, as in [6], [18], [21]. Dror pointed out that, for three exposures, with k_1 the exposure ratio of the first and second exposures, with k_2 the exposure ratio of the second and third exposures, and with $log_{k_1} k_2$ irrational, then, in theory, the self-similar ambiguity disappears [8]. More formally, if we let the intensity mapping functions for the two pairs written as τ_{k_1} and τ_{k_2}, then

$$\tau_{k_1}^m(\tau_{k_2}^{-n}(1)) = g^{-1}(k_1{}^m/k_2{}^n g(1)) = g^{-1}(k_1{}^m/k_2{}^n). \quad (4)$$

It can be shown that $k_1{}^m/k_2{}^n$ is dense in $[0, 1]$ if and only if $log_{k_1} k_2$ is irrational (see Appendix C). Therefore, the inverse response is constrained everywhere. In practice, the numerical irrationality of $log_{k_1} k_2$ is meaningless for a physical quantity such as the ratio of exposures. Nevertheless, some choices of $log_{k_1} k_2$ are particularly poor. For example, when $k_1 \approx k_2$, as is commonly the case [21], [18], [6], multiple exposures offer no advantage over the two-image case, in terms of the self-similar ambiguity. It should also be noted that repeated application of τ_{k_1} and τ_{k_2} will accumulate errors; thus, choosing at least one of the exposure ratios closer to one, for example, $\sqrt{2}$, is advisable.

Implications. Since the extension from s to g is unique, recovering the inverse camera response by solving (3) when the exposure ratio is known is unique up to the self-similar ambiguity. This ambiguity decreases when the exposure ratio is close to 1, keeping in mind that an exposure ratio too close to 1 creates other problems. By taking at least three exposures with exposure ratios, k_1, k_2 with $log_{k_1} k_2$ irrational, the self-similar ambiguity can be broken. Furthermore, this ambiguity can be broken even for two images or for multiple images taken with the same exposure ratio by making a priori assumptions on the inverse response function g.

5 THE EXPONENTIAL AMBIGUITY

In Section 4, we analyzed the recovery of the inverse response when the exposure ratios are known by solving (3). Can we solve (3) to determine the inverse camera

response and the exposure ratios simultaneously? As an example, suppose we have two imaging systems: One system has the inverse response $g_1(B) = B^\gamma$, an example of a "gamma" curve, and another system has a linear inverse response $g_2(B) = B$. Each system takes two images of the same scene. Both systems have identical initial exposures. We change the exposure by a factor of $k = 2^\gamma$ for the first system and $k = 2$ for the second system. The intensity mapping function of the first system is $\tau(B) = g_1^{-1}(2^\gamma g_1(B)) = (2^\gamma B^\gamma)^{-\gamma} = 2B$. The second system also has intensity mapping function $\tau(B) = 2B$. The two systems produce different pairs of images but an identical intensity mapping function for each pair of images, as illustrated in Fig. 3. It is therefore *impossible* to recover the exposure ratio and inverse response simultaneously from the intensity mapping, without making a priori assumptions on the inverse response and exposure ratio.

This exponential ambiguity in solutions for the exposure ratio k and inverse response g is general and not limited to gamma curves. For example, if g and k are arbitrary solutions to $\tau(B) = g^{-1}(kg(B))$, then so are g^γ and k^γ. In other words, if we are given two sets of images of the same scene with identical initial exposure, one from an imaging system with an inverse response g and exposure ratio between the images of k and a second with inverse response function g^γ and exposure ratio k^γ, they have identical intensity mapping functions. The following theorem shows that there are no other ambiguities in (3):

Theorem 3. Exponential Ambiguity. *Suppose we have inverse response functions g_1, g_2 and exposure ratios k_1, k_2, so that*

$$\begin{aligned} g_1(\tau^{-1}(B)) &= k_1^{-1} g_1(B) \\ g_2(\tau^{-1}(B)) &= k_2^{-1} g_2(B). \end{aligned} \quad (5)$$

Define $\beta(B) := g_2(g_1^{-1}(B))$, which is an ambiguity in the solutions g, k to (3), then $\beta(B) = KB^\gamma$ and $k_1^\gamma = k_2$, for some constants γ and K. (See Appendix D for proof.)

Mitsunaga and Nayar [21] discussed this ambiguity in their method for simultaneous recovery of the inverse response and the exposure ratio. They broke the ambiguity by assuming that the inverse response is polynomial. This restriction limits the possible values of γ. They also

assumed they had rough estimates of the exposure ratios. They showed that the multiple solutions are far enough apart so that starting from the rough estimates of the exposure ratios, the correct solution can be found with an iterative method. Tsin et al. break the ambiguity by imposing constraints on the errors in their estimates of the exposure ratios, as well as smoothness and monotonicity of the functions [26]. Without assumptions on the response functions, no method can resolve this ambiguity.

It is also worth noting that using multiple images and, thus, multiple intensity mappings does not break this ambiguity. Suppose we start with three images at different exposures with exposure ratio k_1 between the first and second images and k_2 between the second and third images. Thus, between the first and third images the exposure ratio is $(k_1 k_2)$. Further, suppose the inverse response function is g. We have already shown that using the exposure ratio k_1^γ between the first and second images with camera response g^γ will give the same intensity mapping as when $\gamma = 1$. We may apply the same logic to the second and third images to show that the intensity mapping does not change. What about the first and third images? The exposure ratio between these images is the product of the exposure ratios k_1^γ and k_2^γ. Since, however, $k_1^\gamma k_2^\gamma = (k_1 k_2)^\gamma$, the intensity mapping between the first and third images does not change either. The same holds for an arbitrary number of images. Thus, no algorithm applied to these images can recover exposure ratios and inverse response functions simultaneously without assumptions on the response function.

Implications. Recovery of the exposure ratios and the response function is only possible by making assumptions on the form of the response function or by starting with rough estimates on the exposure ratios as in [21]. We should be wary of applying any algorithm for recovering the response and exposure ratios in situations where we know nothing about either.

6 RECOVERY OF THE EXPOSURE RATIO

We have just shown that it is impossible, without further assumption on the response, to recover the response and exposure ratio together. If we know the response, we can simply recover the exposure k from (3). We now show that, in theory, it is possible to recover the exposure ratio when we don't know the response. If we differentiate both sides of (3) and solve for k, we get

$$k = \tau'(B)g'(\tau(B))/g'(B). \qquad (6)$$

Evaluating this at $B = 0$ using $\tau(0) = 0$ and assuming that the derivative of the inverse response is not zero at zero, $g'(0) \neq 0$, we have $k = \tau'(0)$. This tells us that we can, in theory, directly estimate the exposure ratio from τ, as illustrated in Fig. 4. This may seem to contradict the exponential ambiguity of Section 5. It does not, however, because $(g(B)^\gamma)' = \gamma g(B)^{\gamma-1}g'(B)$, so, since $g(0) = 0$ if $\gamma > 1$, then g^γ has a zero derivative at $B = 0$, and, if $\gamma < 1$, then g^γ has an infinite derivative at $B = 0$. In either case, we cannot cancel $g'(0)$.
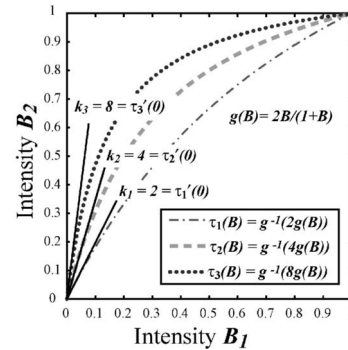

Fig. 4. Graph showing that the exposures ratios k are equal to the slope of the intensity mapping function at the origin, $\tau'(0)$. As an example, we choose the inverse camera response $g(B) = 2B/(B + 1)$ The curves are intensity mapping functions for exposure ratios 2, 4, and 8. This shows that, if the intensity mapping function can be estimated near $B = 0$, the exposure ratios may be recovered without recovering the response function.

Unfortunately, in practice, estimating the exposure ratio using this observation is often impractical. Many cameras have a response which is a gamma curve or have a response which is zero for a significant range of irradiances. In such cases, the camera response violates the hypothesis on $g'(0)$. Moreover, estimating $\tau'(0)$ is difficult since SNR is often very low for small intensity values. Nevertheless, we now show that τ' close to 0 often serves as a rough approximation. Using the Taylor approximation of g' in (6), we get

$$k \approx \left(1 + \frac{g''(B)}{g'(B)}(\tau(B) - B)\right)\tau'(B). \qquad (7)$$

For small B, $\tau(B) - B$ is often small, so, if $g'(B)$ is nonzero, the derivative of the intensity mapping function is a rough estimate of the exposure ratio.

Implications. It is theoretically possible to recover the exposure ratio from τ as long as $g'(0) \neq 0$. In practice, $g(B)$ must be well-behaved with a limited amount of noise near $B = 0$ for this recovery to be possible.

7 INTENSITY MAPPING FUNCTION FROM HISTOGRAMS

As discussed in Section 2, all the information we can recover about the exposure ratios and the response function of the camera comes from the correspondence of intensity values between images. When the images are registered, the pairs of corresponding intensity values may be aggregated as a comparagram, J (see Fig. 5a). The intensity mapping function may be estimated from the comparagram by regression.

Obtaining the comparagram from a pair of images requires the images to be registered. When a scene remains static and the camera motion can be considered a homography, a number of methods have been presented for simultaneously registering images spatially, while recovering the intensity mapping function [2], [4], [15], [19]. Nevertheless, pixel-to-pixel registration between the images can be problematical. Due to occlusion, general object and camera motion can make registration impossible. Even for static scenes and static cameras, pixel-to-pixel estimates of
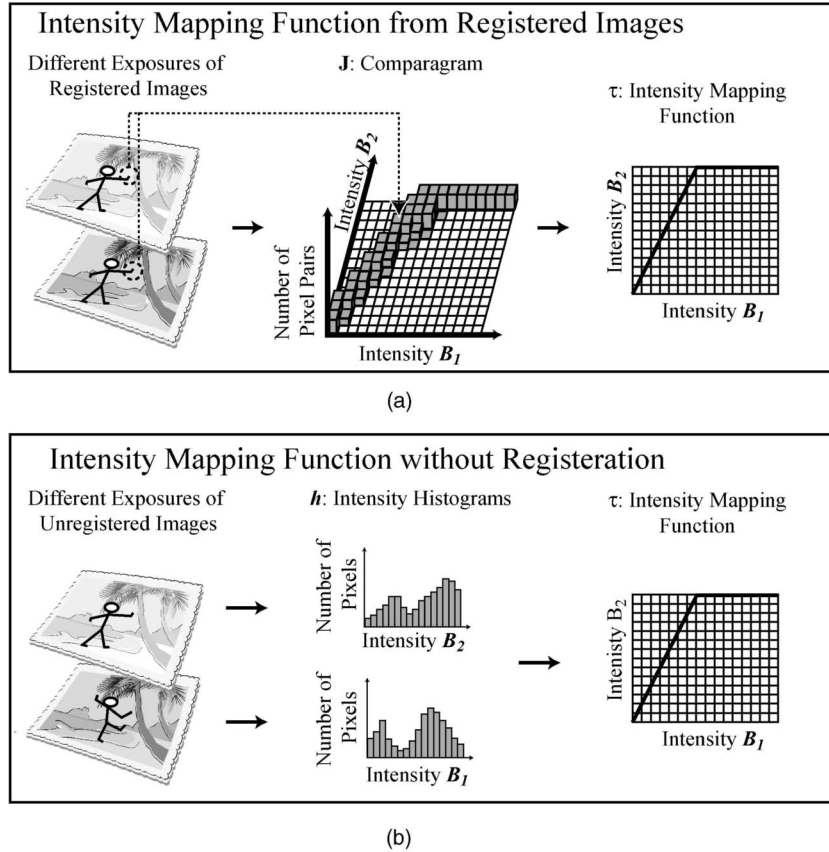
Fig. 5. (a) Diagram showing the recovery of the intensity mapping function from registered images at different exposures. The comparagram counts the number of pixels which have intensity B_1 in one image and intensity B_2 at the corresponding pixel in the second image. From this, we can find a best fit for the intensity mapping function. (b) Diagram showing the recovery of intensity mapping function from differently exposed images without registration. We first compute histograms of the images. Histogram specification gives the intensity mapping function between the images.

the comparagram can be unreliable for edges and textures due to aliasing. While some of these concerns may be addressed by disregarding outliers, it raises the question: What information is required from an image pair to determine the intensity mapping function?

The intensity mapping function is a characteristic of the camera and the exposures and, thus, is not scene dependant. This suggests that it should be possible to determine the intensity mapping by analyzing the change in histograms of the images, as in Fig. 5b. We note that, if the histograms of either image have empty bins, then the images give no information about the intensity mapping function for the corresponding intensities.[13] As an extreme case if one image is completely saturated, the image pair does not have any information about the the intensity mapping. Assuming the image histograms do not have empty bins we have the following theorem:

Theorem 4. Image Information Determining the Intensity Mapping. *Given the histogram of one image, the histogram of the second image is necessary and sufficient to determine the intensity mapping function.*

---

13. This is an issue for all chartless recovery methods. Since we do not control the scene radiances, we must take enough exposures for the scene radiances to provide information across the entire range of intensities. Nevertheless, a parameterized model for the inverse response makes the recovery more robust to missing intensity values and saturation.

To prove this, we start by normalizing all images to have unit area. The total area of the image with intensity values between intensity 0 and intensity B is given by a monotonic function we call $H(B)$. The continuous histogram of intensity values can be defined as the unique function h such that

$$H(B) = \int_0^B h(u)\,du. \qquad (8)$$

This also means that $H(B)$ is the cumulative histogram. Consider two images of a static scene with no camera motion. Ignoring saturation and quantization for the moment, each intensity B_2 in the second image maps to an intensity B_1 in the first image, $B_1 = \tau(B_2)$. The set of image points in the first image with intensity less than B_1, must be the same as the set in the second image with intensity less than B_2 since they correspond to the same set of scene points. Hence, these sets must have equal area, so $H_1(\tau(B_2)) = H_2(B_2)$. This implies that the intensity mapping function and the first histogram determine the second histogram. Replacing $B_1 = u$ and solving for τ, we have

$$\tau(u) = H_2^{-1}(H_1(u)). \qquad (9)$$

This shows that the two histograms determine the intensity mapping function, completing the proof of Theorem 4.

(a)   (b)   (c)   (d)   (e)
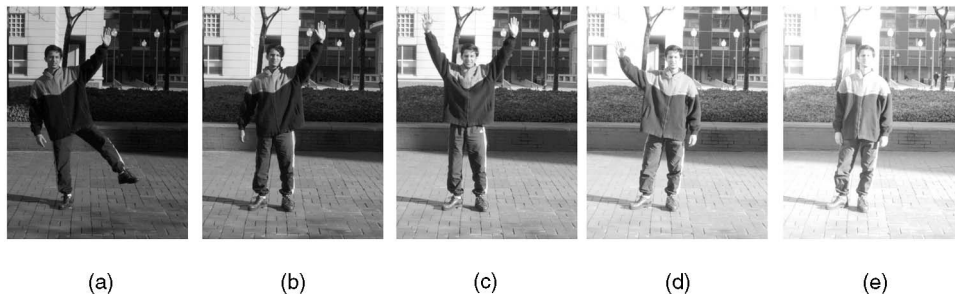
Fig. 6. Sequence of images (a)-(e) with motion of objects in the scene. Here, the man changes the position of one limb between each image. In image (c), both his arms are raised, occluding scene points in the upper part of the background, whereas, in (d), his left arm is lowered, both revealing some pixels in the background and occluding others in the lower part of the image. Even this kind of motion, present in all the image pairs, makes registration of the images impossible. Since the histogram of scene radiance values remains roughly the same, the intensity mapping may be obtained using histogram specification.

We can interpret this theorem in terms of histogram modeling [13]. Histogram modeling changes the histogram of the image by remapping the intensities. One example of histogram modeling is histogram equalization. Assume we have normalized our histograms h_1 and h_2 so that $H_1(1) = H_2(1) = 1$. When intensities in the image are remapped via the function $H_1(u)$, the resulting image has a histogram that is uniform. Another example of histogram modeling is histogram specification, where we attempt to specify the desired histogram for an image. For example, if we want one image to have the same histogram as a second image, we use the function $H_2^{-1}(H_1(u))$ to remap the intensities.

A method obtaining the intensity mapping from the histograms and, thus, omitting the step of registering the images, achieves computational savings. More importantly, recovery of the intensity mapping from the histograms does not require the scene to remain static or camera motion to be a homography. When camera motion involves parallax simple registration is often not possible. If a camera moves around an object, such as a plant, the images may share relatively few scene points due to occlusion. Nevertheless, due to the fact that the images are of the same object from a different viewing direction, it will often have a similar distribution of scene radiances.

Often, scene motion will not change the histogram of scene radiances significantly. For example, consider a natural landscape with trees and grass on a windy day. While scene points may move around spatially in the image, as long as the distribution of scene radiances remains roughly constant, the intensity mapping can be recovered from the image histograms. By not requiring registration, it is possible to extend recovery of the intensity mapping function to a much wider class of images than those permitting registration.

Implication. We have shown that the histograms of an image pair are necessary and sufficient to determine the intensity mapping function. Histogram specification determines the intensity mapping function whenever we expect the histograms of scene radiance to remain approximately constant between images taken under different exposures. This means it is possible to recover the intensity mapping functions in the presences of some scene or camera motion, where registration would be difficult or impossible. It also makes it possible to avoid registration for static scenes,

reducing computational effort and eliminating any errors coming from the registration process.

8   EXPERIMENTAL VERIFICATION

A practical and robust algorithm to recover the intensity mapping function from images using histograms must address a number of issues. We must evaluate the sensitivity of the histogram to missing bins, the sensitivity to changes in scene radiance and to the algorithm's complexity. In this section, we only verify that recovery of the intensity mapping function from images taken at different exposures is possible even in the presence of scene and camera motion.[14] We will verify our recovery of the intensity mapping by determining the inverse response function of the camera. We also show that, under the assumption that the inverse response has nonzero derivative at zero, we can obtain rough estimates of the exposure ratios from our recovered intensity mappings.

Figs. 6a, 6b, 6c, 6d, and 6e show a sequence of five images taken at different exposures with the Nikon 990 Coolpix camera. In order to vary the exposure, we changed the integration times which were: (a) $1/500$ seconds, (b) $1/250$ seconds, (c) $1/125$ seconds, (d) $1/60$ seconds, and (e) $1/30$ seconds. Larger apertures can introduce a spatial variation of image irradiance so we used a small aperture setting of $F = 7.1$.

In between each pair of adjacent images in Fig. 6, the man in the image raises or lowers one limb. This scene motion makes registration difficult since some scene points are occluded in the image while others move. For example, in the image in Fig. 6d, the man's arm covers the building, while, in the image in Fig. 6e, his lower right arm covers the grass. Moreover, the grass and leaves in the scene move due to wind. Registration of such images would require complex tracking and elimination of points that cannot be registered. Nevertheless, this kind of motion does not effect the overall histogram significantly. Hence, we can recover the intensity mapping function simply from histogram specification.

---

14. For our examples, we have changed the exposures manually. In [15], Mann exploits the automatic gain control (AGC) of a camera to obtain different exposures. Since the AGC changes the exposure as a reaction to a change in the distribution of scene radiance, it is difficult to rely on the AGC in our case.
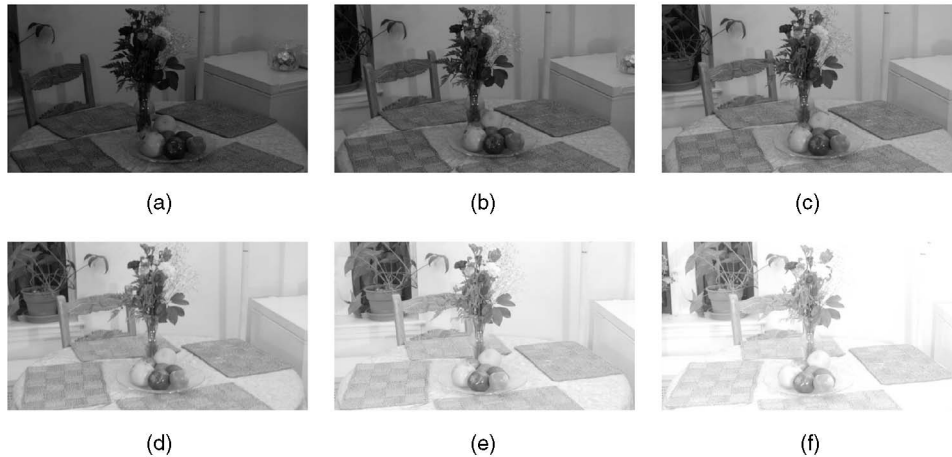
Fig. 7. A sequence of images (a)-(f) of a static scene with camera motion. Each image was taken with the camera directed at the flower vase in the center of the table. The camera's position was moved $5°$ about the axis of the vase. In image (b), a flower pot in the background moved into the field of view on the left. In image (c), a glass bowl of silver wrapped chocolates left the field of view to the right. As in Fig. 6, the parallax and occlusions are problematic for image registration.

An example of a scene with static objects and camera motion is shown in Figs. 7a, 7b, 7c, 7d, 7e, and 7f. The camera here is translated along a circle and rotated while keeping the optical axis intersecting the axis of the flower vase. Each image represents motion along the circle of roughly $5°$. The scene is relatively close and this camera motion introduces parallax and occlusions in the images. For example, the flower pot on the far left of Fig. 7c is outside of the field of view in Fig. 7a, while the bowl of chocolates wrapped in foil is visible in Fig. 7a, but out of the field of view in Fig. 7c. The flowers occlude part of the chair in Fig. 7f and part of the wall in Fig. 7a.

The motion used to create the images in Figs. 7a, 7b, 7c, 7d, 7e, and 7f cannot be modeled by a homography, as in [2], [4], [15], [19], making it impossible to establish correspondences without determining the scene structure. Nevertheless, the overall histogram of scene radiances does not change dramatically between the images as is often the case when changing the viewing direction by a modest amount for the same objects. The images were taken with the aperture fixed at F4.1 and the integration times were (a) 1/30 second, (b) 1/15 second, (c) 1/8 second, (d) 1/4 second, (e) 1/2 second, and (f) 1 second. The images were cropped to minimize spatial variations in exposure.

Finally, Figs. 8a, 8b, 8c, 8d, 8e, and 8f show a sequence of images with both camera motion and complex motion of objects in the scene. In this image sequence, the camera was translated along a line parallel to the face of the building in the image. The camera was rotated to keep the optical axis pointed at the large seated metal statue in the center of the steps. The rotation was approximately $5°$ between each image. The time between when the images were taken was around 3-4 minutes. In that time, some people seated on the stairs left while others arrived. Other people passed through the field of view of the camera. The camera's aperture was fixed at F = 6.2 and the shutter speeds for the images were 1/1,000 sec for (a), 1/500 sec for (b), 1/250 second for (c), 1/125 second for (d), 1/60 second for (e), and 1/30 second for (f). Since the overall histogram of scene

radiances remains roughly constant, the changes of exposure allow us to obtain the intensity mapping using histogram specification.

We computed the cumulative histograms for all the images in Fig. 6. We inverted the cumulative histograms of the images using linear interpolation. For each pair of consecutive images, we computed the intensity mapping using histogram specification. Fig. 9 shows four intensity mapping functions from the intensity values given by the average of the three color channels of the image pairs (a)-(b), (b)-(c), (c)-(d), and (d)-(e), from Fig. 6. The exposure ratios, computed from the integration times reported by the camera for these images, were 2, 2, 2.08, and 2. As pointed out in Section 6, we were able to obtain rough estimates of these ratios from the inspection of the slopes of the intensity mapping near 0. From the graphs, we estimated these slopes by fitting a degree 4 polynomial passing through 0 to the intensity mapping function values for intensity values less than 64. The estimated slopes were 2.4 for image pairs (a)-(b), 2.9 for (b)-(c), 2.6 for (c)-(d), and 1.2 for (d)-(e). These estimates are very rough because it is precisely this part of the intensity mapping function which is most sensitive to noise.

In order to recover the inverse response function of the camera from the intensity mapping, we must make assumptions that break the ambiguities of chartless recovery. We will assume that the inverse response is a low-degree polynomial,[15] as in Mitsunaga and Nayar [21]. We modify their method to recover the camera's response function from the intensity mapping functions, rather than the images. We did not attempt to use our rough exposure estimates as inputs to their iterative scheme for recovering the exposure ratios between pairs of images. For each pair of images, we generated a set of pairs of intensity values B_1 and $B_2 = \tau(B_1)$. We combined the pairs $(n/255, \tau(n/255))$ for 256 intensity values $0 \leq n \leq 255$, with the pairs

---

15. The assumption that the response can be approximated by a polynomial is based on the Taylor polynomials (see [25]). This method is more stable for small intensity values than methods that use the log of (3), such as used in [6].

(a) (b) (c)

(d) (e) (f)

Fig. 8. A sequence of images (a)-(f) of a scene with both object and camera motion. The camera was translated and rotated while facing the large seated statue in the center of the steps in front of the building. The total rotation was approximately $5°$ between each image. The total translation from the image in (a) to the image in (f) was approximately 80 feet. Since each image was taken several minutes apart, people sitting on the steps present in one image are often not in the previous or following image. Other people walk through the field of view of the camera. Despite the scene and camera motion, the overall histogram of scene radiances remains roughly constant permitting recovery of the intensity mapping using histogram specification.

$(\tau^{-1}(n/255), n/255)$. Each pair (B_1, B_2) gave a constraint from (1).

Not all the pairs (B_1, B_2) generated from τ should be weighted equally. Our certainty about our estimation of τ at various intensity values depends on how many pixels in the images have those intensity values. Thus, we weight the pair $(B_1, \tau(B_1))$ by the number of pixels with value equal to B_1, which is $C = h_1(B_1)$. To weigh the least squares problem, we multiplied the constraint (3) by $\sqrt{C}$. Similarly,

we weighed the constraints for the pairs $(\tau^{-1}(B_2), B_2)$ with $C = \sqrt{h_2(B_2)}$. Putting all the constraints together, we assumed that g is a sixth order polynomial and solved a linear system for the coefficients of the polynomial g.

In Fig. 10, we show the recovery of the RGB response curves from the image sequence in Fig. 6. The result is compared with reference points obtained from images with the same camera using a Macbeth chart. Several images of the chart were taken and the data merged using the camera's reported exposure values. Since the global irradiance scale of the response function is unrecoverable, we chose a single best fit scale factor which allows the recovered curve to pass near Macbeth chart points in the middle of the graph. We see that we have excellent overall agreement with the chart recovered response samples. This shows that, from the intensity mapping functions obtained using histogram specification, we can recover the response curves in the presence of modest scene motion.

Fig. 11 shows the recovery of the RGB response curves from the image sequence of Fig. 7. We recovered the curves in Fig. 11 as we did for Fig. 10. The recovery of the red and green channel are in excellent agreement with the Macbeth data. We do not recover the blue channel curve accurately for large blue values. This is because the image does not have enough pixels with information in that part of the blue histogram. This problem faces any chartless recovery method. Chartless recovery is based on the assumption that the scene contains a range of irradiances so that, when combined with the changes of exposure, information about the irradiances associated with each intensity is obtained. If the scene is poor in some range of irradiance, we have no information there. The recovery method used simply extrapolates and, thus, is likely to give poor agreement.
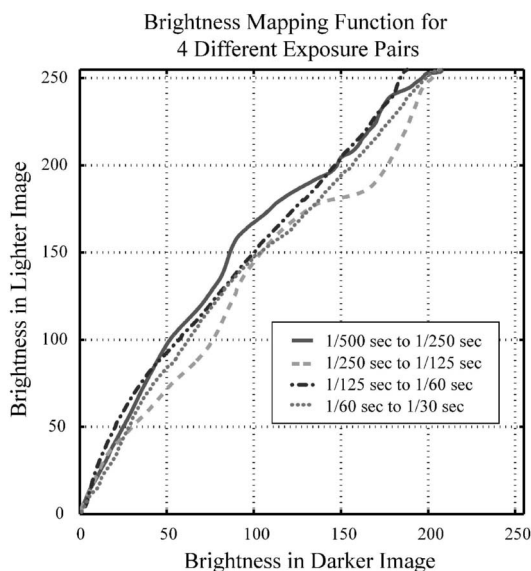

Fig. 9. Intensity mapping functions for the average of the color channels between image pairs from Fig. 6. Very rough estimates of the slopes of these curves near the origin are 2.4 for the image pair Fig. 6a and 6b, 2.9 for Fig. 6b and 6c, 2.6 for Fig. 6c and 6d, and 1.2 for Fig. 6d and 6e. The ground truth ratios of the shutter speeds are 2, 2, 2.1, and 2, respectively.
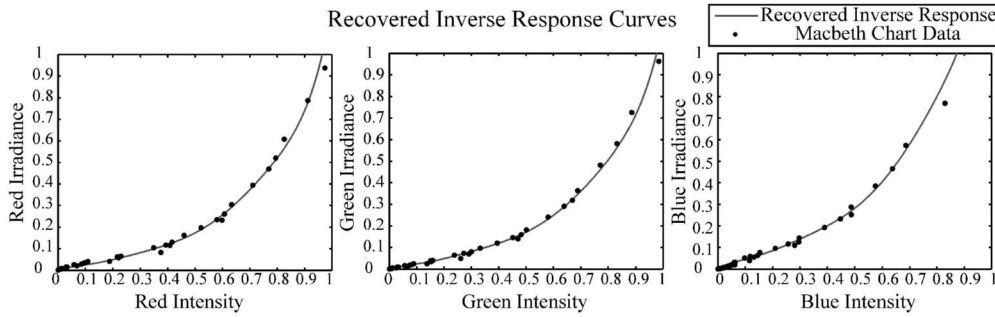
Fig. 10. The recovered red, green, and blue responses computed from Fig. 6 compared with ground truth from the Macbeth chart. We extracted intensity mapping curves from image pairs in Fig. 6 using histogram specification. We broke the ambiguities of recovery by assuming the inverse response to be a polynomial of order 6 as in [4]. The excellent agreement with the Macbeth chart shows that we can extract the intensity mapping functions and recover the inverse camera response even with some scene motion.
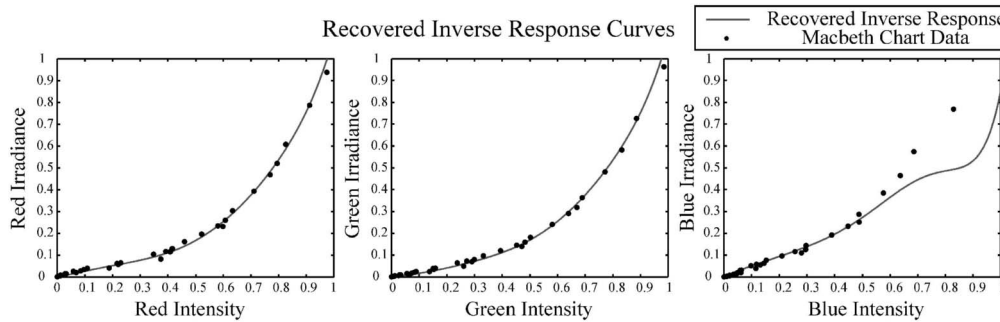

Fig. 11. The recovered red, green, and blue camera responses computed from Fig. 7 compared with ground truth from the Macbeth chart. We recovered these curves as we did for Fig. 7. Note that Fig. 7 is a warm image with red, brown, green, and yellow tones dominating. Here, we see excellent agreement with the Macbeth chart for the red and green responses. We failed to recover the blue response for larger pixel values because there was little information for those colors in Fig. 7. All chartless recovery methods depend on sufficient information across the range of intensity in each channel.
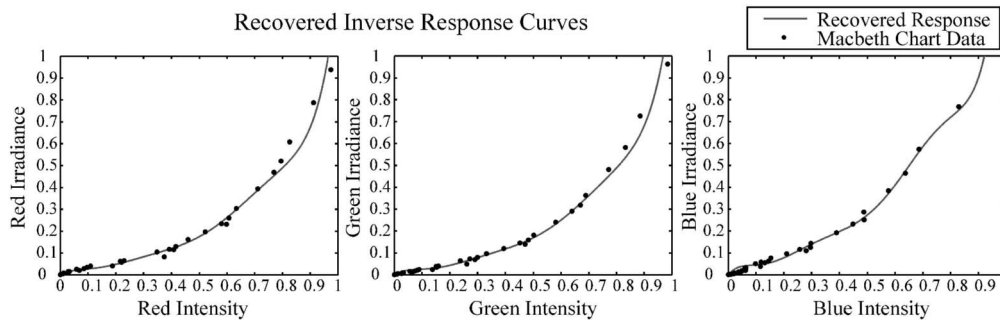

Fig. 12. The recovered red, green, and blue camera responses computed from Fig. 8 compared with ground truth from the Macbeth chart. We recovered these curves as we did for Figs. 10 and 11. Despite the motion of the camera and motion in the scene, we still recover the curves in good agreement with the Macbeth charts in each channel.

In Fig. 12, we see the recovered RGB curves from the images in Fig. 8. In all channels, we have good agreement with the Macbeth data. This demonstrates that image histograms can be used to recover the camera response even with camera motion and scene motion as long as the basic assumption that the histogram of irradiances does not change significantly.

9 SUMMARY

We have proven that the constraints of chartless recovery are not sufficient to allow unambiguous recovery of the camera response function. We have derived a self-similar ambiguity in recovering the camera response function from two images when the exposure ratios are known. We have shown that chartless recovery methods break this ambiguity by assuming a smoothness constraint or using a parameterized form of the response function. We have described how this ambiguity may be reduced or eliminated by choosing appropriate exposure ratios for multiple images. We have proven that assumptions on the response function and the exposure ratios are necessary to recover them simultaneously due to an exponential ambiguity. We have proven these are the only ambiguities. We have shown when the exposure ratios can be found directly from the intensity mapping function. We have shown that the intensity mapping function can be determined by histogram specification. We have proven that chartless recovery can be

accomplished without registration. Finally, we verified that the image histograms alone determine the intensity mapping function. Thus, these histograms can be used to recover a camera response curve from image sequences with both motion of the camera and objects in the scene whenever the histograms of scene radiances are roughly constant.

APPENDIX A

PROOF OF THE PROPERTIES OF THE INTENSITY MAPPING FUNCTION

Recall that we have normalized intensity and irradiance so $g(0) = 0$ and $g(1) = 1$.

1. Evaluating at zero, we find

$$\tau(0) = g^{-1}(kg(0)) = g^{-1}(0) = 0.$$

2. Since g is smooth and monotonically increasing $g' \geq 0$. From (3), we have $\tau'(B) = k \, g'(B)/g'(\tau(B))$. Thus, $\tau'(B) \geq 0$ so τ is monotonically increasing.

3. Since g is monotonically increasing, if $B_1 \leq B_2$, then $g(B_1) \leq g(B_2)$. Since $k > 1$, then $g(B) \leq kg(B) = g(\tau(B))$. Since g^{-1} is also monotonically increasing, $B = g^{-1}(g(B)) \leq g^{-1}(g(\tau(B))) = \tau(B)$.

4. Consider the sequence of decreasing points $B > \tau^{-1}(B) > \tau^{-2}(B) > \dots$. We know that these points are bounded from below by 0. Thus, the sequence must converge to a limit point B^*. At this point, $\tau(B^*) = B^*$. This means $g(B^*) = g(\tau(B^*)) = kg(B^*)$. Since $k > 1$, it must be that $g(B^*) = 0$, thus $B^* = 0$.

APPENDIX B

PROOF OF THE SELF-SIMILAR AMBIGUITY PROPOSITION

Consider the decreasing sequence of points $1 \geq \tau^{-1}(1) \geq \tau^{-2}(1) \geq \dots$. For any point $B \in [0,1]$, since $\lim_{n \to \infty} \tau^{-n}(1) = 0$, there is some nonnegative integer $r(B)$ such that $B \in (\tau^{-r(B)-1}(1), \tau^{-r(B)}(1)]$. Note that $r(\tau(B)) = r(B) - 1$. Define the function

$$g(B) = \frac{1}{k^{r(B)}} s(\tau^{r(B)}(B)) \text{ for } B > 0, g(0) = 0. \qquad (10)$$

Now, observe that

$$
\begin{aligned}
g(\tau(B)) &= \frac{1}{k^{r(\tau(B))}} s(\tau^{r(\tau(B))}(\tau(B))) \\
&= \frac{1}{k^{r(B)-1}} s(\tau^{r(B)-1}(\tau(B))) \\
&= \frac{k}{k^{r(B)}} s(\tau^{r(B)}(B)) = kg(B).
\end{aligned}
\qquad (11)
$$

Thus, we see that g satisfies $g(\tau(B)) = kg(B)$. Since s continuous and monotonic, then so $g(B)$ is inside the union of the disjoint intervals $(\tau^{-n-1}(1), \tau^{-n}(1))$. Since $s(1) = 1$, $s(\tau^{-1}(1)) = 1/k$ and h is monotonic, so g is from (10). Because $s(B)$ is continuous both at $h(1)$ and $h(\tau^{-1}(1))$, so is g at $\tau^{-n}(1)$. We have $\lim_{n \to \infty}(g(\tau^{-n}(1)) = \lim_{n \to \infty} 1/k^n = 0$. Thus, g is continuous at 0.

APPENDIX C

PROOF THAT SELF-SIMILAR AMBIGUITY CAN BE BROKEN WITH THREE EXPOSURES

Equation (4) shows that the self-similar ambiguity can be broken using three images with exposure ratios k_1, and k_2 if k_1^m/k_2^n is dense on $(0,1)$. Here, we give a proof for Dror's observation that k_1^m/k_2^n is dense on on $(0,1)$ if and only if $\log_{k_1} k_2^n$ is irrational [8].

Kronecker's Theorem states that the set $\{z - \lfloor nz \rfloor\}_{n=1}^{\infty}$ is dense on the interval $(0,1)$ if and only if z is irrational (see [11]). This implies that $\{m + nz | (m,n) \in \mathbb{Z}\}$ is dense on the whole real line. Hence, it is also dense on $(-\infty, 0)$. Thus, exponentiating with any positive number k_1, we have that $k_1^m k_1^{-nz}$ is dense on the interval $(0,1)$. If we set $z = \log_{k_1} k_2$, then $k_1^m k_1^{-nz} = k_1^m k_2^{-n}$; therefore, k_1^m/k_2^n is dense on the interval $(0,1)$ if and only if $log_{k_1} k_2$ is irrational.

APPENDIX D

PROOF OF THE EXPONENTIAL AMBIGUITY

If g_1, g_2 are monotonic functions, then so is β. Note that $\beta(g_2(B)) = g_1(B)$. Since $\beta(k_1^{-n} g_1(B)) = \beta(g_1(\tau^{-n}(B))) = g_2(\tau^{-n}(B))) = k_2^{-n} g_2((B)) = k_2^{-n} \beta(g_1((B)))$, we can simplify this equation by calling $c = g_1((B))$, $\gamma = \ln k_2 / \ln k_1$, and $k_1^{-n} = a$. Then, the equation becomes $\beta(ac) = a^\gamma \beta(c)$.

Note that $\gamma = \ln k_2 / \ln k_1$ implies $k_1^\gamma = k_2$. Now, for a sequence of points of $b \geq \tau^{-1}(b) \geq \tau^{-2}(b) \geq \dots$, the response g_1 has values $g_1(b) \geq (1/k_1)g_1(b) \geq (1/k_1^2)g_1(b) \geq \dots$, while the response g_2 has the sequence $K \geq (1/k_1^\gamma)K \geq (1/k_1^{2\gamma})K \geq \dots$, where $K = \beta(g_1(b))$. Since these sequences are only determined up to a factor of scale, we have shown that these sequences can have at most an ambiguity up to an exponential $\beta(B) = KB^\gamma$. If $\beta(B) = g_2(g_1^{-1}(B))$ for all B (not just along the sequence) then, since $\beta(1) = 1$, $K = 1$.

REFERENCES

[1] N. Asada, A. Amano, and M. Baba, "Photometric Calibration of Zoom Lens Systems," *Proc. Int'l Conf. Pattern Recognition,* vol. A, pp. 189-190, 1996.

[2] A.F. Barros and F.M. Candocia, "Image Registration in Range Using a Constrained Piecewise Linear Model," *Proc. Int'l Conf. Acoustics, Speech, and Signal Processing,* vol. 4, pp. 3345-3348, 2002.

[3] P.N. Belhumeur and D.J. Kriegman, "What Is the Set of Images of an Object under All Possible Illumination Conditions?" *Int'l J. Computer Vision,* vol. 28, no. 3, pp. 245-260, July 1998.

[4] F.M. Candocia, "A Least Squares Approach for the Joint Domain and Range Registration of Images," *Proc. Int'l Conf. Acoustics, Speech, and Signal Processing,* vol. 4, pp. 3237-3240, 2002.

[5] Y.C. Chang and J.F. Reid, "RGB Calibration for Color Image Analysis in Machine Vision," *Proc. Int'l Conf. 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," *Proc. ACM SIGGRAPH,* pp. 369-378, 1997.

[7]   P.E. Debevec, "Rendering Synthetic Objects into Real Scenes," *Proc. ACM SIGGRAPH,* pp. 189-198, 1998.

[8]   R. Dror, personal communication, 2002.

[9]   H. Farid, "Blind Inverse Gamma Correction," *Proc. Int'l Conf. Image Processing,* vol. 10, no. 10, pp. 1428-1433, Oct. 2001.

[10] M.D. Grossberg and S.K. Nayar, "What Can Be Known about the Radiometric Response from Images?" *Proc. European Conf. Computer Vision,* vol. 4, pp. 189-205, 2002.

[11] G.H. Hardy and E.M. Wright, *An Introduction to the Theory of Numbers.* Oxford Univ. Press, 1979.

[12] B. Horn, *Robot Vision.* The MIT Press, 1986.

[13] A.K. Ja, *Fundamentals of Digital Image Processing.* Prentice Hall, 1989.

[14] L.T. Maloney and B.A. Wandell, "Color Constancy: A Method for Recovering Surface Spectral Reflectance," *J. Optical Society of Am. A,* vol. 3, pp. 29-33, 1986.

[15] S. Mann, "Pencigraphy with AGC: Joint Parameter Estimation in Both Domain and Range of Functions in Same Orbit of the Projective-Wyckoff Group," *Proc. Int'l Conf. Image Processing,* vol. 3, pp. 193-196, 1996.

[16] S. Mann, "Comparametric Equations with Practical Applications in Quantigraphic Image Processing," *Proc. Int'l Conf. Image Processing,* vol. 9, no. 8, pp. 1389-1406, Aug. 2000.

[17] S. Mann, "Comparametric Imaging: Estimating Both the Unknown Response and the Unknown Set of Exposures in a Plurality of Differently Exposed Images," *Proc. CS Conf. Computer Vision and Pattern Recognition,* Dec. 2001.

[18] S. Mann and R. Picard, "Being 'Undigital' with Digital Cameras: Extending Dynamic Range by Combining Differently Exposed Pictures," *Proc. 46th Ann. IS&T Conf.,* pp. 422-428, 1995.

[19] S. Mann and R.W. Picard, "Video Orbits of the Projective Group: A Simple Approach to Featureless Estimation of Parameters," *Proc. Int'l Conf. Image Processing,* vol. 6, no. 9, pp. 1281-1295, Sept. 1997.

[20] S.E. Marschner, "Image-Based BRDF Measurement," PhD thesis, Stanford Univ., Calif., 1998.

[21] T. Mitsunaga and S.K. Nayar, "Radiometric Self Calibration," *Proc CS Conf. Computer Vision and Pattern Recognition,* vol. 2, pp. 374-380, June 1999.

[22] 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.

[23] F.E. Nicodemus, J.C. Richmond, J.J. Hsia, I.W. Ginsberg, and T. Limperis, "Geometrical Considerations and Nomenclature for Reflectance," NBS Monograph 160, Nat'l Bureau of Standards, Washington, D.C., Oct. 1977.

[24] Y.Y. Schechner and S.K. Nayar, "Generalized Mosaicing," *Proc. Int'l Conf. Computer Vision,* vol. 1, pp. 17-24, 2001.

[25] J. Stewart, *Calculus,* fourth ed. Brooks/Cole, 1999.

[26] Y. Tsin, V. Ramesh, and T. Kanade, "Statistical Calibration of the CCD Imaging Process," *Proc. Int'l Conf. Computer Vision,* vol. 1, pp. 480-487, 2001.

[27] 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.

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 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.

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. He 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** http://computer.org/publications/dlib.