

What Can Be Known about the Radiometric Response from Images?

Michael D. Grossberg and Shree K. Nayar

Columbia University, New York NY 10027, USA,
{mdog,nayar}@cs.columbia.edu,
<http://www.cs.columbia.edu/CAVE>

Abstract. Brightness values of pixels in an image are related to image irradiance by a non-linear function, called the radiometric response function. Recovery of this function is important since many algorithms in computer vision and image processing use image irradiance. Several investigators have described methods for recovery of the radiometric response, without using charts, from multiple exposures of the same scene. All these recovery methods are based solely on the correspondence of gray-levels in one exposure to gray-levels in another exposure. This correspondence can be described by a function we call the brightness transfer function. We show that brightness transfer functions, and thus images themselves, do not uniquely determine the radiometric response function, nor the ratios of exposure between the images. We completely determine the ambiguity associated with the recovery of the response function and the exposure ratios. We show that all previous methods break these ambiguities only by making assumptions on the form of the response function. While iterative schemes which may not converge were used previously to find the exposure ratio, we show when it can be recovered directly from the brightness transfer function. We present a novel method to recover the brightness transfer function between images from only their brightness histograms. This allows us to determine the brightness transfer function between images of different scenes whenever the change in the distribution of scene radiances is small enough. We show an example of recovery of the response function from an image sequence with scene motion by constraining the form of the response function to break the ambiguities.

1 Radiometric Calibration without Charts

An imaging system usually records the world via a brightness image. When we interpret the world, for example if we try to estimate shape from shading, we require the scene radiance at each point in the image, not just the brightness value. Some devices produce a brightness which is a linear function of scene radiance, or at least image irradiance. For most devices, such as consumer digital, video, and film cameras, a non-linear radiometric response function gives brightness in terms of image irradiance.¹

¹ We are ignoring spatially varying linear factors, for example, due to the finite aperture. We will assume that the response function is normalized both in domain (irradiance) and range (brightness).

Some vision applications such as tracking may not require precise linear irradiance values. Nevertheless, when one estimates the illumination space of an object as in [1], estimates the BRDF from images [2], determines the orientation of surface normals [3], or merges brightness images taken at different exposures to create high dynamic range images [4], one must find irradiance values from brightness values by determining the radiometric response of the imaging system.

We can recover the radiometric response function by taking an image of a uniformly illuminated chart with patches of known reflectance, such as the Macbeth chart. Unfortunately, using charts for calibration has drawbacks. We may not have access to the imaging system so we may not be able to place a chart in the scene. Changes in temperature change the response function making frequent recalibration necessary for accurate recovery of the response. Additionally, we must uniformly illuminate the chart which may be difficult outside of a laboratory environment.

The problems associated with using charts have lead researchers to develop methods to recover the radiometric response from registered images of a scene taken with different exposures. Mann and Picard assume the response function has the form of a gamma curve [5]. They estimate its parameters assuming they know the exposure ratios between the images. Debevec and Malik also assume the ratio of exposures is known, but they recover the log of the inverse radiometric response without a parametric form [6]. To obtain a solution they impose a smoothness constraint on the response. Mitsunaga and Nayar assume that the inverse response function can be closely approximated by a polynomial [4]. They are then able to recover the coefficients of that polynomial and the exposure ratios, from rough estimates of those ratios. Tsin, Ramesh and Kanade [7] and separately, Mann [8] recover both the response and exposure ratios by combining the iterative approach from [4], with the non-parametric recovery in [6].

The essential information all methods use for recovery is how brightness gray-levels in one image correspond to brightness gray-levels in another. Ideally a function we call the brightness transfer function describes this correspondence.² Figure 1 illustrates the role of the brightness transfer function. All chart-less recovery methods [4, 5, 6, 7, 8] use the constraint that all irradiances change by the exposure ratios to recover response function and exposure ratios from exactly the information contained in the brightness transfer function.

Previous authors have not completely addressed whether there actually is enough information in the images to recover the response function and exposure ratios from images without charts. Put another way, since the brightness transfer function contains all the information about the response function, we could ask: given the exposure ratio, is there a unique response function for each brightness transfer function? We show that solutions to this inverse problem are not unique. We demonstrate that due to a fractal-like ambiguity any method to recover the response function must constrain the function, for example with regularity,

² Most methods use all pairs of gray-levels at corresponding points in the images. Mann [8] shows that the brightness transfer function summarizes this information.

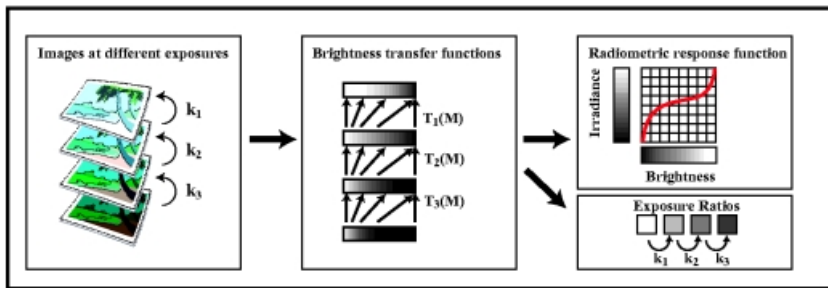


Fig. 1. A diagram showing the role of the brightness transfer function. All the information in the images, relevant for chart-less recovery is contained in the brightness transfer functions. These describe how a brightness in one image corresponds to a brightness in another image. Thus, the problem of recovering the radiometric response function falls into two parts: recovery of the brightness transfer functions from images, and recovery of the radiometric response function and the exposure ratios from the brightness transfer functions.

to break this ambiguity. Beyond this ambiguity, is it possible to recover the response function and the exposure ratios simultaneously and uniquely? We show that there are families of solutions, arising from what we call an exponential ambiguity. Again, only by making assumptions on the response function can we expect a unique solution. Can we recover the exposure without recovering the response? We will show when this is possible.

Given that it is possible to recover the response function and exposure ratios from the brightness transfer function by making assumptions on the form of the response, how do we recover the brightness transfer function from images? Previous work compared registered images of a static scene taken at different exposures to recover the brightness transfer functions. Is it necessary for the scene to be static and do we require any spatial information to recover the brightness transfer function? It is not, since we show that the brightness transfer function can be obtained from 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 brightness transfer function. To illustrate this, we recover the response function from a sequence of images with scene motion.

2 The Fundamental Constraint for Chart-Less Recovery

The brightness value at a point in the image should allow us to determine the scene radiance. The ideal brightness value I is linear in the scene radiance L . The ideal brightness is related to scene radiance by $I = LPe$, where P is a factor due to the optics of the system, and e is the exposure, following the notation of [4]. 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.³ 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 the ideal brightness I as image plane irradiance.

A function f called the radiometric response function relates the actual measured brightness value $M = f(I)$ at a photosensitive element to the image plane irradiance I . Imaging system designers often intentionally create a non-linear radiometric response function f , to compress the dynamic range, for example. Since measured brightness indicates relative irradiance, we can assume that the response function f monotonically increases.⁴ The minimum irradiance is 0, while the maximum irradiance, I_{max} is a single, unrecoverable parameter. Thus we normalize domain of f , irradiance I , to go from 0 to 1. We also normalize the range of f , brightness M , so that $f(1) = 1$ and $f(0) = 0$.⁵ 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 constraints that permit us to extract f from images without assuming the knowledge of scene reflectances.

As a special case of what we mean, suppose we take images A and B of the same scene with different exposures e_A and e_B . If image A has image irradiance I_A at a point and the corresponding point in image B has the irradiance I_B , then $I_A/e_A = I_B/e_B$. The exposure ratio $k := e_B/e_A$ expresses the relationship between the two images, $I_B = kI_A$. Relating this back to measurements in images A and B we have $f(I_A) = M_A$ and $f(I_B) = M_B$. Monotonicity of radiometric response function makes it invertible. Let $g := f^{-1}$ be the inverse of f . Then, we have the equation,

$$g(M_B) = kg(M_A). \quad (1)$$

All chart-less methods base recovery of g and k on this equation. In each pair of images, each corresponding pair of pixel brightness values gives one constraint. The exposure ratio k is constant for each pair of images. When k is known, this gives a linear set of equations g . If g is a polynomial, then equation (1) becomes linear in the coefficients of the polynomial. Mitsunaga and Nayar solve for these coefficients [4]. Debevec and Malik [6] and Mann [8] take the log of both sides of equation (1). Rather than start with a parameterized model of $\log g$ they discretely sample it at the gray-level values, treating it as a vector. By imposing a regularity condition on the discrete second derivatives of $\log g$, they are able to obtain a solution.

³ Details of P for a simple perspective camera can be found in Horn [9]. Whereas Mitsunaga and Nayar [4] 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.

⁴ A monotonically decreasing response function is also possible, as in negative films. We can, however, re-normalize f so that f increases with irradiance, for example by multiplication by -1 .

⁵ The minimum brightness value in digital imaging systems is often effectively greater than zero due to non-zero mean thermal noise called dark current. By taking an image with the lens covered, this effect may be estimated and subtracted.

When we know g but not k , we can solve equation (1) for k . Mitsunaga and Nayar [4] and Mann [8] use an iterative scheme in which they first solve for g with an initial guess for k . Updating their estimates, they iteratively solve for k and g .

3 Brightness Transfer Functions

The pairs of measurements M_A and M_B at corresponding points in different images of the same scene constitute all the information available from which to recover the response function in chart-less recovery. Mann [8] pointed out that all this information is contained in a two variable histogram he calls the *comparagram*. If (M_A, M_B) are any two pairs of brightness values, then $J(M_A, M_B)$ is the number of pixels which have brightness value M_A in image A and M_B at the corresponding point in image B .

The comparagram encodes how a gray-level in image A corresponds to gray-level in image B . For real images, a probability distribution most accurately models this correspondence, rather than a function. A function fails to model all the pixel pairs because of noise, quantization of the brightness values, spatial quantization, and saturated pixels. Ignoring these considerations for a moment, from equation (1), we see that a function should ideally relate the brightness values in the images

$$M_B = T(M_A) := g^{-1}(kg(M_A)) \quad (2)$$

which we call the *brightness transfer function*. This function describes how to transfer brightness values in one image into the second image. We can estimate the brightness transfer function T from J .⁶ For a collection of images $A_1, A_2, A_3, \dots, A_l$ the brightness transfer functions of all possible pairs $T_{m,n}$ summarize the correspondence of gray-levels between the images. Once we estimate T , we have a modified version of equation (1), given by

$$g(T(M)) = kg(M). \quad (3)$$

This equation has an advantage over equation (1): because the function T contains all information about the gray-level correspondence between images, we can study the mathematical problem of existence and uniqueness of solutions to equation (3).

To study solutions to equation (3) we must first derive some properties of brightness transfer functions. From equation 2 we know that the monotonicity of g implies the monotonicity of T and so T^{-1} exists. Define $T^0(M) := M$, $T^n(M) := T(T^{n-1}(M))$ and $T^{-n}(M) := T^{-1}(T^{1-n}(M))$.

Theorem 1 (Properties of the Brightness Transfer Function). *Let g be a smooth monotonically increasing function with smooth inverse. Suppose that $g(0) = 0$ and $g(1) = 1$, and $k > 1$, then the function $T(M) := g^{-1}(kg(M))$*

⁶ Mann [8] calls T the *comparametric function*, and the process of going from J to T , *comparametric regression*.

has the following properties: (a) T monotonically increases, (b) $T(0) = 0$, (c) $\lim_{n \rightarrow \infty} T^{-n}(M) = 0$, and (d) if $k > 1$, then $M \leq T(M)$. [see appendix A for proof].

Assuming that $k > 1$ just means that we order our images so that exposure increases. For example, for a pair of images with $0 < k < 1$, we can replace equation (3) with the equation $f(T^{-1}(M)) = (1/k)f(M)$, where $(1/k) > 1$. By replacing k with $1/k$ and T with T^{-1} , we have reordered our images so that $k > 1$. To order the images themselves note that when $k > 1$, then $g^{-1}(kI) \geq g^{-1}(I)$ since g^{-1} monotonically increases. In other words every brightness value in one image corresponds to a brighter value in the other image. Therefore, for $k > 1$, T goes from the image with darker average pixel value, to the image with lighter average pixel value.⁷

4 Fractal Ambiguity

If the exposure ratio $k > 1$, the brightness transfer function T expands the M -axis as we see from theorem 1(d). Equation (3) says by stretching the domain of g with T , and multiplying the range of g by k , the function g becomes self-similar. This kind of self-similarity is the hallmark of fractals. The left side of equation (3) relates the value of g at points in $(T^{-1}(1), 1]$, to the value of g at points in $(T^{-2}(1), T^{-1}(1)]$ on the right side. Nothing, however, relates the value of g at points in $(T^{-1}(1), 1]$ to *each-other*. This means as long as $g(1) = 1$, $g(T^{-1}(1)) = 1/k$, and g is continuous and monotonic, then g can have arbitrary values on $(T^{-1}(1), 1)$, and still be a solution to equation (3). We call this ambiguity to equation (3), the fractal ambiguity.⁸ More formally, we state this by saying that we can build a solution g starting with any function $s(M)$ on $[T^{-1}(1), 1]$:

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

We understand this fractal ambiguity by looking at a sampling of s . On the left in figure 2, take any three points $p_1, p_2, p_3 \in (T^{-1}(1), 1]$. We can choose the values of s hence g at these points essentially arbitrarily. The only restriction *a-priori* is that s be monotonic, and thus $s(T^{-1}(1)) = 1/k \leq s(p_n) \leq 1$. Each point $p_1 \in (T^{-1}(1), 1]$ gives rise to a sequence of points $p_1 \geq T^{-1}(p_1) \geq T^{-2}(p_1) \geq \dots$. Equation (3) determines this sequence from p_1 . It places no restrictions on the relationship between the values of s at points in $(T^{-1}(1), 1]$. On the right in figure 2 we see alternative solutions to equation (3). The fractal ambiguity presents a problem relating different sequences of points, each obtained by choosing a

⁷ Note that we do not have to register images to compare the average pixel values.

⁸ Mann [8], has noted the relationship of his version of equation (3) to dynamical systems, and that the equations allow “ripples” in the solutions




Fig. 2. Left: Graph with arbitrary increasing values chosen at points p_1, p_2, p_3 . Each point, p_1 , generates a sequence of points. Right: Two arbitrary curves s on $(T^{-1}(1), 1]$ extended uniquely to $[0, 1]$. This means neither equation (3), nor its log, tell us anything about the values of g , in $(T^{-1}(1), 1]$. Unique solutions of g can only come from prior assumptions on the form or smoothness of g .

point p_1 and a value $s(p_1)$ in $(T^{-1}(1), 1]$. The choice of initial value $s(p_1)$ determines each sequence, however, only up to a single multiple. Only by constraining the form of g , or imposing a regularity constraint, can we break this ambiguity. For example, suppose we take a sequence of points in the plane $(1, 1)$, $(T^{-1}(1), 1/k)$, $(T^{-2}(1), 1/k^2), \dots (T^{-N}(1), 1/k^N)$. If we assume that g is a low degree $m \ll N$ polynomial, then we can solve for the least squares polynomial fit for the points. This breaks the ambiguity by allowing the polynomial to interpolate between $T^{-1}(1)$ and 1 based on the best fit to the rest of the points. Choosing multiple exposure ratios k reduces this ambiguity only to the extent that for some exposure ratio, $1 - T^{-1}(1)$ is smaller.

Implications: Since the extension from s to g is unique, solving equation (3) when k is known is unique up to the fractal ambiguity. The fractal ambiguity suggests that it would be best to choose an exposure ratio as close to 1 as practical. Unfortunately, in the presence of noise, T may be difficult to recover accurately when k is close to 1. Nevertheless, minimizing $1 - T^{-1}(1)$, which shrinks as k approaches 1, minimizes the fractal ambiguity. The fractal ambiguity also shows that the values of any solution g , on $(T^{-1}(1), 1)$, to equation (3) must come from *a-priori* assumptions on g .

5 The Exponential Ambiguity

If solutions to equation (3) for known exposure ratios are unique up to the fractal ambiguity what happens as we try to solve for both g and k together? As an example, suppose we have two imaging systems: System A has an inverse response $g_A(M) = M^\gamma$ and system B has a linear inverse response $g_B(M) = M$. Each takes two images of the same scene. Both systems have identical initial exposures. We change the exposure by a factor of 2^γ for system A , and 2 for

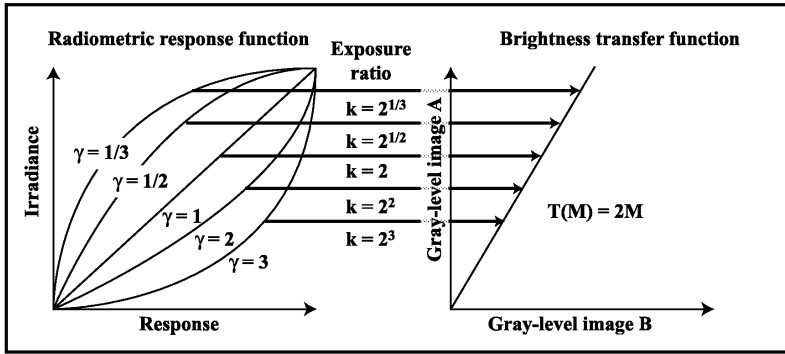


Fig. 3. Several different response functions giving rise to the same brightness transfer function T . The brightness transfer function $T(M) = 2M$ results from an inverse response function M^γ , and exposure ratio between images of 2^γ , independent of γ . This shows we can only recover the radiometric response and exposure simultaneously by making assumptions on the solutions, g and k , that break this ambiguity.

system B . System A has brightness transfer function is $T(M) = g_A^{-1}(2^\gamma g_A(M) = (2^\gamma M^\gamma)^{-\gamma} = 2M$. System B also has brightness transfer function $T(M) = 2M$. Systems A and B will produce different images and brightness values but the correspondence between gray levels in one image and gray-levels in the other, the brightness transfer function is identical. As illustrated in figure 3, it is therefore impossible to recover k and g simultaneously from T alone, without making *a-priori* assumptions on g and k .

It is easy to see that this exponential ambiguity does not only come from gamma curves. For example, if g and k are solutions to $T(M) = g^{-1}(kg(M))$ 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 device with an inverse response function 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 the identical brightness transfer functions. The following theorem shows that there are no other ambiguities in equation (3):

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

$$\begin{aligned} g_1(T^{-1}(M)) &= k_1^{-1}g_1(M) \\ g_2(T^{-1}(M)) &= k_2^{-1}g_2(M) \end{aligned} \tag{4}$$

Define $\beta(M) := g_2(g_1^{-1}(M))$, which is an ambiguity in solutions g, k to equation (3), then $\beta(M) = KM^\gamma$, and that $k_1^\gamma = k_2$, for some constants γ , and K [see appendix C for proof.]

Mitsunaga and Nayar [4] discussed this ambiguity in their method for simultaneous recovery of g and k . They resolved it by assuming that g is polynomial.

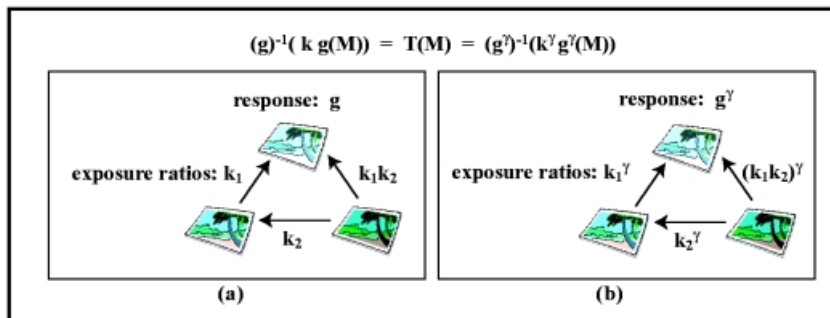


Fig. 4. Directed graph of images at different exposures. The exponential ambiguity cannot be broken by taking more images. If the inverse response function and all the exposure ratios are raised to the same power, the brightness transfer functions do not change. We cannot tell situation (a) from situation (b) from $T(M)$.

This restriction limits the possible values of γ . They assumed also they have rough estimates of the exposure ratios. They show that the multiple solutions are far enough apart so that from the rough estimate of exposure ratios the correct solution can be found with an iterative method. In [7] they 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. Without assumptions on the response functions, no method can resolve this ambiguity.

It is also worth noting that using multiple images and thus multiple T does not break this ambiguity. In figure 4, we see that the images form a directed graph with arrows going from darker images to brighter images. If we start with images at different exposures with exposure ratios k_1, k_2 , and $(k_1 k_2)$, and inverse response function g as in figure 4(a), and raise all the exposure ratios and the inverse response function by γ , we get an identical brightness transfer functions figure 4(b). 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 [4]. We should be wary of applying any algorithm for recovering 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. It is not surprising that if we know the response, we can simply recover the exposure k from equation (3). It is surprising that it is possible to recover the exposure ratio when we don't know the response. If we differentiate both sides of equation (3) with

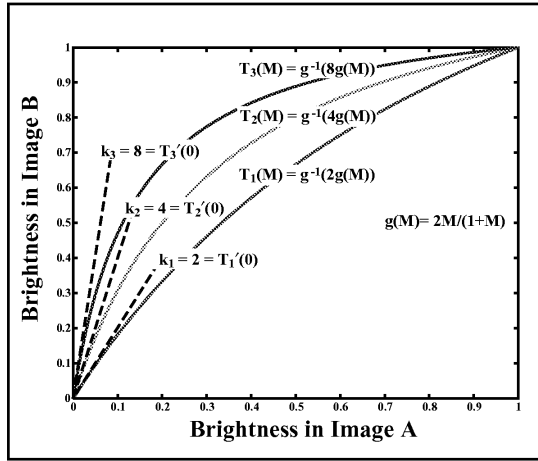


Fig. 5. Graph showing that the exposures ratios k are equal to the slope of the brightness transfer function at the origin, $T'(0)$. As an example, we choose the inverse radiometric response $g(M) = 2M/(M + 1)$ The curves are brightness transfer functions for exposure ratios 2, 4, and 8. This shows that if the brightness transfer function can be estimated near $M = 0$, the exposure ratios may be recovered without recovering the response function.

respect to M , we get $g'(T(M))T'(M) = kg'(M)$. Evaluating this at $M = 0$, and using the fact that $T(0) = 0$ we get $g'(0)T'(0) = kg'(0)$. We have $T'(0) = k$, when $g'(0) \neq 0$. This tells us that we can, in theory, directly estimate the exposure ratio from T , as illustrated in figure 5. This seems to contradict the exponential ambiguity of section 5. It does not, however, because $(g(M)^\gamma)' = \gamma g(M)^{\gamma-1}g'(M)$, so since $g(0) = 0$, if $\gamma > 1$, then g^γ has a zero derivative at $M = 0$, and if $\gamma < 1$, then g^γ has an infinite derivative at $M = 0$. In either case we can not cancel $g'(0)$.

Unfortunately, in practice, estimating $T'(0)$ may not be practical. The SNR is often very low for small gray-levels, making estimation of $T(M)$ near $M = 0$ difficult.

Implications: It is theoretically possible to recover the exposure ratio from T alone as long as $g'(0) > 0$. In practice $g(M)$ must be well behaved with a limited amount of noise near $M = 0$ for this recovery to be robust.

7 Recovering the Brightness Transfer Function from Histograms

As discussed in section 2, all the information we can recover about the exposure ratios, and the radiometric response function comes from the correspondence of gray-levels between images of the same static scene taken at different exposures. Previous work recovered this correspondence by registering the images. The brightness values at corresponding points may be collected in a compara-

gram, J . From the comparagram, J the brightness transfer function, $T(M)$ can be estimated.

The need for registration restricts the camera motion and requires static scenes. We will show we can broaden the kinds of images we may use by eliminating the need for spatial correspondence. In this context previous methods needed spatial correspondence between images to obtain gray-level correspondence between images. In fact, we may use any feature of gray-levels that allows us to establish this correspondence, and thus $T(M)$. For example, we will now show how we can recover $T(M)$ using the histogram, a simple statistic of gray-levels.

Suppose we have two images, μ_A , and μ_B , of the same scene taken with different exposures. We interpret both $\mu_A(x, y)$, and $\mu_B(x, y)$ as functions of brightness with x and y coordinates in the image. We denote the set of all points in the image taking values between 0 and M_A by $\mu_A^{-1}([0, M_A]) := \{(x, y) | 0 \leq \mu_A(x, y) \leq M_A\}$. The continuous histogram of brightness values, can be defined as the unique function h_A such that

$$\int_{\mu_A^{-1}([0, M_A])} dx dy = \int_0^{M_A} h_A(u) du. \quad (5)$$

This integral, the cumulative histogram $H_A(M_A)$, can be interpreted as the area of image points with gray-levels less than M_A . Ignoring saturation and quantization for the moment, each gray-level M_B in image B corresponds to a gray value M_A in image A , $M_B = T(M_A)$. The set of image points in image A with gray-levels less than M_A , must be the same as the set in image B with gray-levels less than M_B , since they correspond to the same set of scene points. Hence, these sets must have equal area, so $H_A(M_A) = H_B(M_B) = H_B(T(M_A))$.

Finally, replacing $M_A = u$ and solving for T we have

$$T(u) = H_B^{-1}(H_A(u)). \quad (6)$$

We can interpret this equation in terms of histogram modeling, as in [10]. 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_A , and h_B so that $H_A(1) = H_B(1) = 1$ so the image has unit area. When intensities in, for example image μ_A are remapped via the function $H_A(u)$ is 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 image μ_A to have the same histogram as image μ_B , we use the function $H_B^{-1}(H_A(u))$ to remap the intensities.⁹

⁹ Histogram specification will not give information about $T(M)$ for gray-levels M where the histogram is zero. This is not a weakness with this method but rather with all chart-less 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 brightnesses.

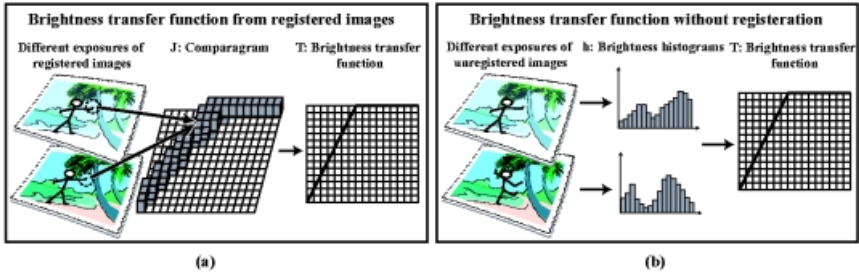


Fig. 6. (a) Diagram showing recovery of the brightness transfer function from registered images at different exposures. The comparagram counts the number of pixels which have gray-level M_A in image A and gray-level M_B at the corresponding pixel in image B . From this we can find a best fit brightness transfer function T . (b) Diagram showing recovery of brightness transfer function T from differently exposed images without registration. We first compute histograms of the images. Histogram specification gives the brightness transfer function between images.

We have shown that if we assume our images μ_A and μ_B are different exposures of the same scene can recover the brightness transfer function T from the histograms of the image, via histogram specification. By omitting the step of registering the images, we achieve some computational savings as well as remove a potential source of error. More importantly we can relax the assumption that the scene remains static. Often scene motion will not change the histogram significantly. While scene points may move around spatially in the image, as long as the distribution of scene radiances remains roughly constant, our method can be applied. The same is true for camera motion provided the distribution of scene radiances remains roughly constant. Scene radiances will not remain constant for arbitrary scene changes, for example, if the illumination changes significantly. Nevertheless, by not requiring registration our method works on a much wider class of images than previous methods.

Implication: Histogram specification gives a simple way to recover T when we expect the histograms of scene radiance to remain approximately constant between images taken under different exposure. Eliminating registration makes it possible to recover the brightness transfer 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

In this section we recover the brightness transfer function from images with scene motion using our histogram specification method. In order to recover the radiometric response function and exposure ratios we must make assumptions that break the ambiguities of chart-less recovery. We modify the method in Mitsunaga and Nayar [4] to recover the radiometric response function from the brightness transfer functions, rather than images. Our goal in recovering the




Fig. 7. Five different exposures of similar scenes taken with a Nikon 990 Coolpix camera. We vary exposures in the images by changing the integration time: (a) 1/500 sec, (b) 1/125 sec, (c) 1/60 sec, (d) 1/30 sec, and (d) 1/15 sec. Not only does the figure in the foreground change pose but the plants in the background are blowing in the wind. There is no point to point correspondences between the images. Previous algorithms for recovering the response curve or exposure ratios from images, cannot handle this case.

response is to simply verify our recovery of the brightness transfer functions. We also show that under the assumption that the response function has non-zero derivative at zero, we can obtain rough estimates of the exposure ratios from our recovered brightness transfer functions.

Figure 7 shows a sequence of 5 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 sec, (b) 1/125 sec, (c) 1/60 sec, (d) 1/30 sec, and (d) 1/15 sec. Larger apertures decrease the uniformity of image irradiance at the image plane so we used a small aperture setting of $F = 8$.

In the images, the man with the cricket bat moves in the scene. The plants in the background move with the breeze.¹⁰ Our assumption, however, is that although scene points move, the overall distribution of intensities in the scene remains approximately constant between the images. On the other hand, this motion is sufficient to make point to point correspondence between the images impossible. Thus previous chart-less methods cannot be used on such image sequences.

We computed the cumulative histograms for all the images in figure 8. We inverted the cumulative histograms using linear interpolation. For each image pair μ_A , and μ_B , we computed T using the histogram specification $T(M) = H_A^{-1}(H_B(M))$. Figure 8 shows 4 brightness transfer functions $T(M)$ corresponding to image pairs (a)-(b), (b)-(c), (c)-(d), and (d)-(e), from figure 7. The exposure ratios, computed from the integration times reported by the camera for these images, are 4, 2.08, 2, and 2. As pointed out in section 6, we can obtain rough estimates of these ratios from inspection of $T'(M)$ at $M = 0$. From the graphs we estimated $T'(0)$ at 4.3 ± 1 for image pairs (a)-(b), $2.4 \pm .4$ for (b)-(c), $1.8 \pm .4$ for (c)-(d), and $1.8 \pm .4$ for (d)-(e). These estimates are very

¹⁰ The illumination does not change.

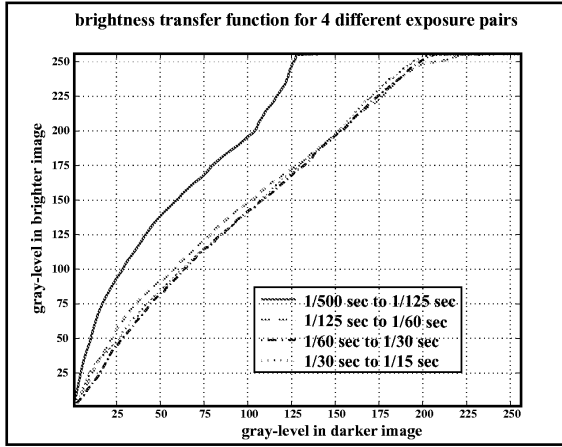


Fig. 8. Brightness transfer curves obtained by histogram specification, from image pairs (a)-(b), (b)-(c), (c)-(d), (d)-(e), in figure 7. We estimate that the curve corresponding to the exposure pair 1/500 sec to 1/125 sec roughly has slope 4.3 ± 1 near the origin. Near the origin, the curve for (b)-(c) has slope $2.4 \pm .4$ and the other curves slopes $1.8 \pm .4$, which provide very rough estimates of their exposure ratios.

rough because it is precisely this part of the curve which is most sensitive to noise.

The algorithm in Mitsunaga-Nayar [4] was designed for registered images rather than recovery of the response from $T(M)$. Also, we did not attempt to use our rough exposure estimates as inputs to their iterative scheme for recovering the exposure ratios. Rather, since we are just verifying the brightness transfer functions, we used their assumptions on g to break the ambiguities and recover the inverse response curve. For each pair of images μ_A and μ_B we generated pairs of brightness values M_A and M_B using T . We combined the pairs $(n/255, T(n/255))$ for 256 gray-levels $0 \leq n \leq 255$, with the pairs $(T^{-1}(n/255), n/255)$. Each pair (M_A, M_B) gives us a constraint from equation 1.

The cumulative histograms, H_A, H_B , from which we computed T , are not uniform. This means our certainty about our estimation of T at various gray-levels depends on how many pixels in the images have those gray-levels. We need to weight our constraints by our certainty. For the pairs $(M_A, T(M_A))$, the number $h_A(M_A) = C$ is the number of pixels with value equal to M_A . The number C gives us some indication of our confidence in T at gray-level M_A . To weight the least squares problem we multiplied the constraint by \sqrt{C} , so the constraint became $g(T(M_A))\sqrt{C} = kf(M_A)\sqrt{C}$. Similarly we weighted the constraints for the pairs $(T^{-1}(M_B), M_B)$ by $\sqrt{H_B(M_B)}$. Putting the constraints together we had 4 pairs of images \times 256 gray-levels \times 2 kinds of brightness pairs (those where $n/255$ is a gray-level in A or B) to get 2048 constraint equations, weighted by the root of the histograms. We broke the ambiguity as was done in




Fig. 9. The inverse radiometric response curves recovered from the image sequence in figure 7 compared to data from a Macbeth chart. The image sequence was used to obtain the brightness transfer curves from figure 8. To recover the inverse response curve we broke the ambiguity in the constraint equation by assuming g to be a polynomial of order 6. The excellent agreement with the Macbeth chart shows that by using our method to extract the brightness transfer functions, we can recover the inverse radiometric response even with some scene motion.

[4]; we assumed that g is a sixth order polynomial and solved the 2048 equation linear system for the coefficients of the polynomial g .

In figure 9, we show the recovery of the RGB response curves. 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 cameras 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 by using histogram specification to recover the brightness transfer function, we can recover the radiometric response curves, even in the presence of modest scene motion.

A Appendix: Proof of the Properties of T

- (a) Evaluating at zero, we find $T(0) = g^{-1}(kg(0)) = g^{-1}(0) = 0$.
- (b) Since g is smooth and monotonically increasing $g' \geq 0$. From equation 3 we have $T'(M) = kg'(M)/g'(T(M))$. Thus $T'(M) \geq 0$ so T is monotonically increasing.
- (c) Since g is monotonically increasing, if $M_1 \leq M_2$ then $g(M_1) \leq g(M_2)$. Since $k > 1$ then $g(M) \leq kg(M) = g(T(M))$. Since g^{-1} is also monotonically increasing, $M = g^{-1}(g(M)) \leq g^{-1}(g(T(M))) = T(M)$.
- (d) Consider the sequence of decreasing points $M > T^{-1}(M) > T^{-2}(M) > \dots$. We know that these points are bounded from below by 0. Thus the sequence must converge to a limit point M^* . At this point $T(M^*) = M^*$. This means $g(M^*) = g(T(M^*)) = kg(M^*)$. Since $k > 1$, it must be that $g(M^*) = 0$, thus $M^* = 0$

B Appendix: Proof of the Fractal Ambiguity Proposition

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

$$g(M) = \frac{1}{k^{r(M)}} s(T^{r(M)}(M)) \text{ for } M > 0, g(0) = 0 \tag{7}$$

Now observe that $g(T(M)) = \frac{1}{k^{r(T(M))}} s(T^{r(T(M))}(T(M))) = \frac{1}{k^{r(M)-1}} s(T^{r(M)-1}(T(M))) = \frac{k}{k^{r(M)}} s(T^{r(M)}(M)) = kg(M)$. Thus we see that g satisfies $g(T(M)) = kg(M)$. Since s continuous and monotonic, then so is $g(M)$ inside the union of the disjoint intervals $(T^{-n-1}(1), T^{-n}(1))$. Since $s(1) = 1$, $s(T^{-1}(1)) = 1/k$, and h is monotonic, so is g from equation 7. Because $s(M)$ is continuous both at $h(1)$ and $h(T^{-1}(1))$, so is g at $T^{-n}(1)$. We have $\lim_{n \rightarrow \infty} (g(T^{-n}(1))) = \lim_{n \rightarrow \infty} 1/k^n = 0$. Thus g is continuous at 0.

C Appendix: Proof of the Exponential Ambiguity

If g_1, g_2 are monotonic functions then so is β . Note that $\beta(g_2(M)) = g_1(M)$. Since $\beta(k_1^{-n} g_1(M)) = \beta(g_1(T^{-n}(M))) = g_2(T^{-n}(M)) = k_2^{-n} g_2(M) = k_2^{-n} \beta(g_1(M))$, we can simplify this equation by calling $c = g_1(M)$, $\gamma = \ln k_2 / \ln k_1$, 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 $p \geq T^{-1}(p) \geq T^{-2}(p) \geq \dots$. The response g_1 has values $g_1(p) \geq (1/k_1)g_1(p) \geq (1/k_1^2)g_1(p) \geq \dots$, 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 $K = \beta(g_1(p))$. 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 $\beta(M) = KM^\gamma$. If $\beta(M) = g_2(g_1^{-1}(M))$ for all M not just along the sequence, then since $\beta(1) = 1, K = 1$.

Acknowledgments. This work was supported by a National Science Foundation ITR Award (IIS-00-85864).

References

1. Belhumeur, P., Kriegman, D.: What is the set of images of an object under all possible illumination conditions? *IJCV* **28** (1998) 245–260
2. Eric, S.M.: Image-based brdf measurement (1998)
3. Zhang, R., Tsai, P., Cryer, J., Shah, M.: Shape from shading: A survey. *PAMI* **21** (1999) 690–706
4. Mitsunaga, T., Nayar, S.K.: Radiometric self calibration. In: Proc CVPR. Volume 2. (1999) 374–380

5. Mann, S., Picard, R.: Being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures. In: In Proceedings of IS&T, 46th annual conference. (1995) 422–428
6. Debevec, P.E., Malik, J.: Recovering high dynamic range radiance maps from photographs. In: Computer Graphics, Proc. SIGGRAPH. (1997) 369–378
7. Tsin, Y., Ramesh, V., Kanade, T.: Statistical calibration of the ccd imaging process. In: ICCV01. (2001) I: 480–487
8. Mann, S.: Comparametric imaging: Estimating both the unknown response and the unknown set of exposures in a plurality of differently exposed images. In: Proc. CVPR, IEEE Computer Society (2001)
9. Horn, B.: Robot Vision. The MIT Press (1986)
10. Jain, A.: Fundamentals Of Digital Image Processing. Prentice Hall, Engle wood Cliffs (1989)