

Comparametric Equations with Practical Applications in Quantigraphic Image Processing

Steve Mann

Abstract—It is argued that, hidden within the flow of signals from typical cameras, through image processing, to display media, is a homomorphic filter. While homomorphic filtering is often desirable, there are some occasions where it is not. Thus, cancellation of this implicit homomorphic filter is proposed, through the introduction of an antihomomorphic filter. This concept gives rise to the principle of quantigraphic image processing, wherein it is argued that most cameras can be modeled as an array of idealized light meters each linearly responsive to a semi-monotonic function of the quantity of light received, integrated over a fixed spectral response profile. This quantity is neither radiometric nor photometric, but, rather, depends only on the spectral response of the sensor elements in the camera. A particular class of functional equations, called comparametric equations, is introduced as a basis for quantigraphic image processing. Comparametric equations are fundamental to the analysis and processing of multiple images differing only in exposure. The well-known “gamma correction” of an image is presented as a simple example of a comparametric equation, for which it is shown that the underlying quantigraphic function does not pass through the origin. For this reason it is argued that exposure adjustment by gamma correction is inherently flawed, and alternatives are provided. These alternatives, when applied to a plurality of images that differ only in exposure, give rise to a new kind of processing in the “amplitude domain” (as opposed to the time domain or the frequency domain). While the theoretical framework presented in this paper originated within the field of wearable cybernetics (wearable photographic apparatus) in the 1970s and early 1980s, it is applicable to the processing of images from nearly all types of modern cameras, wearable or otherwise. This paper is a much revised draft of a 1992 peer-reviewed but unpublished report by the author, entitled “Lightspace and the Wyckoff principle.”

Index Terms—Comparametric equation, comparametric plot, image processing, lightspace, personal imaging, photography, quantigraphic imaging, wearable cybernetics, Wyckoff principle.

I. INTRODUCTION

THE theory of quantigraphic image processing, with comparametric equations, arose out of the field of wearable cybernetics, within the context of so-called mediated reality (MR) [1] and personal imaging [2]. However, it has potentially much more widespread applications in image processing than just the wearable photographic personal assistant for which it was developed. Accordingly, a general formulation that does not necessarily involve a wearable photographic system will be given.

Manuscript received March 31, 1999; revised March 1, 2000. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Michael Frater.

The author is with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, Ont., M5S 3G4, Canada.

Publisher Item Identifier S 1057-7149(00)06143-1.

II. WYCKOFF PRINCIPLE AND THE RANGE OF LIGHT

The quantity of light falling on an image sensor array, or the like, is a real valued function $q(x, y)$ of two real variables x and y . An image is typically a degraded measurement of this function, where degradations may be divided into two categories, those that act on the domain (x, y) and those that act on the range q . Sampling, aliasing, and blurring act on the domain, while noise (including quantization noise) and the nonlinear response function of the camera act on the range q .

Registering and combining multiple pictures of the same subject matter will often result in an improved image of greater definition. There are four classes of such improvement:

- 1) increased spatial resolution (domain resolution);
- 2) increased spatial extent (domain extent);
- 3) increased tonal fidelity (range resolution);
- 4) increased dynamic range (range extent).

A. What is Good for the Domain is Good for the Range

The notion of producing a better picture by combining multiple input pictures has been well-studied with regards to the domain (x, y) of these pictures. Horn and Schunk, for example, provide means of determining optical flow [3], and many researchers have then used this result to spatially register multiple images in order to provide a single image of increased spatial resolution and increased spatial extent. Subpixel registration methods such as those proposed by [4] and [5] attempt to increase domain resolution. These methods depend on slight (subpixel) shift from one image to the next. Image compositing (mosaicking) methods such as those proposed by [6]–[8] attempt to increase domain extent. These methods depend on large shifts from one image to the next.

Methods that are aimed at increasing domain resolution and domain extent tend to also improve tonal fidelity, to a limited extent, by virtue of a signal averaging and noise reducing effect. However, we shall see in what follows, a generalization of the concept of signal averaging called quantigraphic signal averaging. This generalized signal averaging allows images of different exposure to be combined to further improve upon tonal fidelity (range resolution), beyond improvements possible by traditional signal averaging. Moreover, the proposed methodology drastically increases dynamic range (range extent). Just as spatial shifts in the domain (x, y) improve the image, we will also see how exposure shifts (shifts in the range, q) can, with the proposed methodology, result in even greater improvements to the image.

B. Extending Dynamic Range and Improvement of Range Resolution by Combining Differently Exposed Pictures of the Same Subject Matter

The principles of quantigraphic image processing and the notion of using differently exposed pictures of the same subject matter to make a picture composite of extended dynamic range was inspired by the pioneering work of Charles Wyckoff who invented so-called “extended response film” [9], [10].

Most everyday scenes have a far greater dynamic range than can be recorded on a photographic film or electronic imaging apparatus. However, a set of pictures, that are identical except for their exposure, collectively show us much more dynamic range than any single picture from that set, and also allow the camera’s response function to be estimated, to within a single constant scalar unknown [6], [11], [12].

A set of functions

$$f_i(\mathbf{x}) = f(k_i q(\mathbf{x})) \quad (1)$$

where k_i are scalar constants, is known as a Wyckoff set [6], [12]. A Wyckoff set of functions, $f_i(\mathbf{x})$ describes a set of images differing only in exposure, when $\mathbf{x} = (x, y)$ is the continuous spatial coordinate of the focal plane of an electronic imaging array (or piece of film), q is the quantity of light falling on the array (or film), and f is the unknown nonlinearity of the camera’s (or combined film’s and scanner’s) response function. Generally, f is assumed to be a pointwise function, e.g., invariant to \mathbf{x} .

C. Photoquantity q

The quantity, q , in (1), is called the *photoquantigraphic* quantity [13], or just the photoquantity (or photoq) for short. This quantity is neither radiometric (e.g. neither *radiance* nor *irradiance*) nor photometric (e.g. neither *luminance* nor *illuminance*). Most notably, since the camera will not necessarily have the same spectral response as the human eye, or, in particular, that of the photopic spectral luminous efficiency function as determined by the CIE and standardized in 1924, q is neither brightness, lightness, luminance, nor illuminance. Instead, quantigraphic imaging measures the quantity of light integrated over the spectral response of the particular camera system

$$q = \int_0^\infty q_s(\lambda) s(\lambda) d\lambda \quad (2)$$

where $q_s(\lambda)$ is the actual light falling on the image sensor and s is the spectral sensitivity of an element of the sensor array. It is assumed that the spectral sensitivity does not vary across the sensor array.

D. Camera as an Array of Lightmeters

The quantity q reads in units that are quantifiable (e.g. linearized or logarithmic), in much the same way that a photographic light meter measures in quantifiable (linear or logarithmic) units. However, just as the photographic light meter imparts to the measurement its own spectral response (e.g., a light meter using a selenium cell will impart the spectral response of selenium cells to the measurement) quantigraphic imaging

accepts that there will be a particular spectral response of the camera, which will define the quantigraphic unit q . Each camera will typically have its own quantigraphic unit. In this way, the camera may be regarded as an array of lightmeters, each being responsive to the quantigral

$$q(x, y) = \int_0^\infty q_{ss}(x, y, \lambda) s(\lambda) d\lambda \quad (3)$$

where q_{ss} is the spatially varying spectral distribution of light falling on the image sensor.

Thus, varying numbers of photons of lesser or greater energy (frequency times Planck’s constant) are absorbed by a given element of the sensor array, and, over the temporal quantization time of a single frame in the video sequence (or the exposure time of a still image) result in the photoquantity given by (3).

In the case of a color camera, or other color processes, $q(x, y)$ is simply a vector quantity. Color images may arise from as little as two channels, as in the old bichromatic (orange and blue) motion pictures, but more typically arise from three channels, or sometimes more as in the four color offset printing, or even the high quality Hexachrome printing process. A typical color camera might, for example, include three channels, e.g., $[q_r(x, y), q_g(x, y), q_b(x, y)]$, where each component is derived from a separate spectral sensitivity function. Alternatively, another space such as YIQ, YUV, or the like, may be used, in which, for example, the Y (luminance) channel has full resolution and the U and V channels have reduced (e.g., half in each linear dimension giving rise to one quarter the number of pixels) spatial resolution and reduced quantizational definition. In this paper, the theory will be developed and explained for greyscale images, where it is understood that most images are color images, for which the procedures are applied either to the separate color channels, or by way of a multichannel quantigraphic analysis. Thus in both cases (greyscale or color) the continuous spectral information $q_s(\lambda)$ is lost through conversion to a single number q or to typically three numbers, q_r, q_g, q_b . Although it is easiest to apply the theory of this paper to color systems having distinct spectral bands, there is no reason why it cannot also be applied to more complicated polychromatic, possibly tensor, quantigrals.

Ordinarily cameras give rise to noise, e.g., there is noise from the sensor elements and further noise within the camera (or equivalently noise due to film grain and subsequent scanning of a film, etc.). Thus a goal of quantigraphic imaging is to attempt to estimate the photoquantity q , in the presence of noise. Since $q_s(\lambda)$ is destroyed, the best we can do is to estimate q . Thus q is the fundamental or “atomic” unit of quantigraphic image processing.

E. Accidentally Discovered Componder

Most cameras do not provide an output that varies linearly with light input. Instead, most cameras contain a dynamic range compressor, as illustrated in Fig. 1. Historically, the dynamic range compressor in video cameras arose because it was found that televisions did not produce a linear response to the video signal. In particular, it was found that early cathode ray screens provided a light output approximately equal to voltage raised to the exponent of 2.5. Rather than build a circuit into every

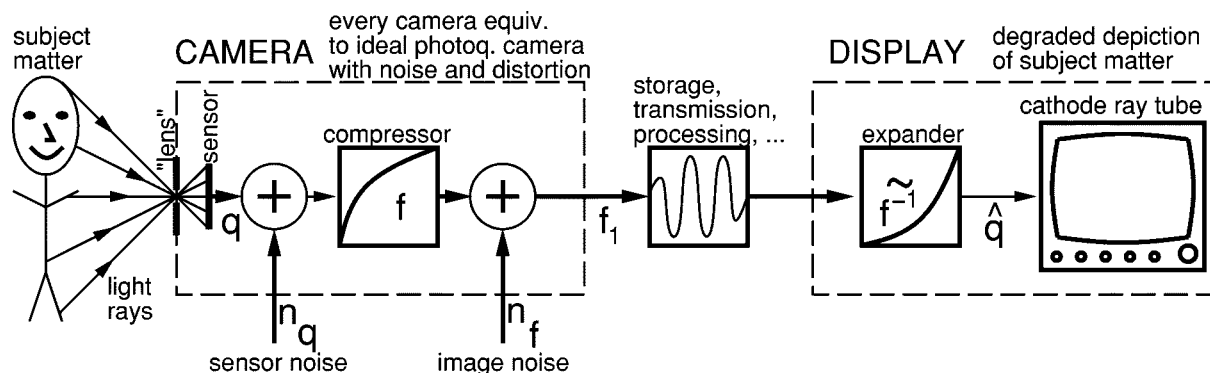


Fig. 1. Typical camera and display: light from subject matter passes through lens (typically approximated with simple algebraic projective geometry, e.g. an idealized “pinhole”) and is quantified in units “q” by a sensor array where noise n_q is also added, to produce an output which is compressed in dynamic range by a typically unknown function f . Further noise n_f is introduced by the camera electronics, including quantization noise if the camera is a digital camera and compression noise if the camera produces a compressed output such as a JPEG image, giving rise to an output image $f_1(x, y)$. The apparatus that converts light rays into $f_1(x, y)$ is labeled CAMERA. The image f_1 is transmitted or recorded and played back into a DISPLAY system where the dynamic range is expanded again. Most cathode ray tubes exhibit a nonlinear response to voltage, and this nonlinear response is the expander. The block labeled “expander” is generally a side effect of the display, and is not usually a separate device. It is depicted as a separate device simply for clarity. Typical print media also exhibit a nonlinear response that embodies an implicit “expander.”

television to compensate for this nonlinearity, a partial compensation (exponent of 1/2.22) was introduced into the television camera at much lesser total cost since there were far more televisions than television cameras in those days before widespread deployment of video surveillance cameras and the like. Indeed, the original model of television is suggested by the names of some of the early players: ABC (American Broadcasting Corporation); NBC (National Broadcasting Corporation); etc.. Names like this suggest that they envisioned a national infrastructure in which there would be one or two television cameras and millions of television receivers.

Through a very fortunate and amazing coincidence, the logarithmic response of human visual perception is approximately the same as the inverse of the response of a television tube (e.g. human visual response turns out to be approximately the same as the response of the television camera) [14], [15]. For this reason, processing done on typical video signals will be on a perceptually relevant tone scale. Moreover, any quantization on such a video signal (e.g. quantization into 8 bits) will be close to ideal in the sense that each step of the quantizer will have associated with it a roughly equal perceptual change in perceptual units.

Fig. 2 shows plots of the compressor (and expander) used in video systems together with the corresponding logarithm $\log(q + 1)$, and antilogarithm $\exp(q) - 1$, plots of the human visual system and its inverse. (The plots have been normalized so that the scales match.)

With images in print media, there is a similarly expansive effect in which the ink from the dots bleeds and spreads out on the printed paper, such that the mid tones darken in the print. For this reason printed matter has a nonlinear response curve similar in shape to that of a cathode ray tube (e.g., the nonlinearity expands the dynamic range of the printed image). Thus cameras designed to capture images for display on video screens have approximately the same kind of built-in dynamic range compression suitable for print media as well.

It is interesting to compare this naturally occurring (and somewhat accidental) development in video and print media with the deliberate introduction of companders (compressors




Fig. 2. The power law dynamic range compression implemented inside most cameras has approximately the same shape of curve as the logarithmic function, over the range of signals typically used in video and still photography. Similarly, the power law response of typical cathode ray tubes, as well as that of typical print media, is quite similar to the antilog function. Therefore, the act of doing conventional linear filtering operations on images obtained from typical video cameras, or from still cameras taking pictures intended for typical print media, is, in effect, homomorphic filtering with an approximately logarithmic nonlinearity.

and expanders) in audio. Both the accidentally occurring compression and expansion of picture signals and the deliberate use of logarithmic (or mu-law) compression and expansion of audio signals serve to allow 8 bits to be used to often encode these signals in a satisfactory manner. (Without dynamic range compression, 12 to 16 bits would be needed to obtain satisfactory reproduction.)

Most still cameras also provide dynamic range compression built into the camera. For example, the Kodak DCS-420 and DCS-460 cameras capture internally in 12 bits (per pixel per color) and then apply dynamic range compression, and finally output the range-compressed images in 8 bits (per pixel per color).

F. Why Stockham was Wrong

When video signals are processed, using linear filters, there is an implicit homomorphic filtering operation on the photoquantity. As should be evident from Fig. 1, operations of storage,




Fig. 3. *The anti-homomorphic filter*: Two new elements \hat{f}^{-1} and \hat{f} have been inserted, as compared to Fig. 1. These are *estimates* of the the inverse and forward nonlinear response function of the camera. Estimates are required because the exact nonlinear response of a camera is generally not part of the camera specifications. (Many camera vendors do not even disclose this information if asked.) Because of noise in the signal f_1 , and also because of noise in the estimate of the camera nonlinearity f , what we have at the output of \hat{f}^{-1} is not q , but, rather, an estimate, \hat{q} . This signal is processed using linear filtering, and then the processed result is passed through the estimated camera response function, \hat{f} , which returns it to a compressed tone scale suitable for viewing on a typical television, computer, or the like, or for further processing.

transmission, and image processing take place between approximately reciprocal nonlinear functions of dynamic range compression and dynamic range expansion.

Many users of image processing methodology are unaware of this fact, because there is a common misconception that cameras produce a linear output, and that displays respond linearly. In fact there is a common misconception that nonlinearities in cameras and displays arise from defects and poor quality circuits, when in actual fact these nonlinearities are fortuitously present in display media and deliberately present in most cameras.

Thus, the effect of processing signals such as f_1 in Fig. 1 with linear filtering is, whether one is aware of it or not, homomorphic filtering.

Stockham advocated a kind of homomorphic filtering operation in which the logarithm of the input image was taken, followed by linear filtering (e.g. linear space invariant filters), followed by taking the antilogarithm [16].

In essence, what Stockham didn't appear to realize, is that such homomorphic filtering is already manifest in simply doing ordinary linear filtering on ordinary picture signals (whether from video, film, or otherwise). In particular, the compressor gives an image $f_1 = f(q) = q^{1/2.22} = q^{0.45}$ (ignoring noise n_q and n_f) which has the approximate effect of $f_1 = f(q) = \log(q + 1)$ (e.g., roughly the same shape of curve, and roughly the same effect, e.g., to brighten the mid-tones of the image prior to processing), as shown in Fig. 2. Similarly a typical video display has the effect of undoing (approximately) this compression, e.g. darkening the mid-tones of the image after processing with $\hat{q} = \hat{f}^{-1}(f_1) = f_1^{2.5}$.

Thus in some sense what Stockham did, without really realizing it, was to apply dynamic range compression to already range compressed images, then do linear filtering, then apply dynamic range expansion to images being fed to already expansive display media.

G. On the Value of Doing the Exact Opposite of What Stockham Advocated

There exist certain kinds of image processing for which it is preferable to operate linearly on the photoquantity q . Such operations include sharpening of an image to undo the effect of the point spread function (PSF) blur of a lens. It is interesting to note

that many textbooks and papers that describe image restoration (e.g. deblurring an image) fail to take into account the inherent nonlinearity deliberately built into most cameras.

What is needed to do this deblurring and other kinds of quantigraphic image processing is an *anti-homomorphic filter*. The manner in which an anti-homomorphic filter is inserted into the image processing path is shown in Fig. 3.

Consider an image acquired through an imperfect lens that imparts a blurring to the image. The lens blurs the actual spatio-spectral (spatially varying and spectrally varying) quantity of light $q_{ss}(x, y, \lambda)$, which is the quantity of light falling on the sensor array just prior to being *measured* by the sensor array

$$\tilde{q}_{ss}(x, y, \lambda) = \iint B(x - u, y - v) q_{ss}(u, v, \lambda) du dv. \quad (4)$$

This blurred spatio-spectral quantity of light $\tilde{q}_{ss}(x, y, \lambda)$ is then photoquantified by the sensor array

$$\begin{aligned} q(x, y) &= \int_0^\infty \tilde{q}_{ss}(x, y, \lambda) s(\lambda) d\lambda \\ &= \int_0^\infty \int_{-\infty}^\infty \int_{-\infty}^\infty B(x - u, y - v) q_{ss}(u, v, \lambda) \\ &\quad \cdot s(\lambda) du dv d\lambda \\ &= \int_{-\infty}^\infty \int_{-\infty}^\infty B(x - u, y - v) \\ &\quad \cdot \left(\int_0^\infty q_{ss}(u, v, \lambda) s(\lambda) d\lambda \right) du dv \\ &= \int_{-\infty}^\infty \int_{-\infty}^\infty B(x - u, y - v) q(u, v) du dv \end{aligned} \quad (5)$$

which is just the blurred photoquantity q .

Thus the antihomomorphic filter of Fig. 3 can be used to better undo the effect of lens blur than traditional linear filtering which simply applies linear operations to the signal f_1 and therefore operates homomorphically rather than linearly on the photoquantity q .

Thus we see that in many practical situations, there is an articulable basis for doing exactly the opposite of what Stockham advocated (e.g., expanding the dynamic range of the image before processing and compressing it afterward as opposed to what Stockham advocated which was to compress the dynamic range before processing and expand it afterward).

H. Using Differently Exposed Pictures of the Same Subject Matter to Get a Better Estimate of q

Because of the effects of noise (quantization noise, sensor noise, etc.), in practical imaging situations, the Wyckoff set that describes a plurality of pictures that differ only in exposure (1) should be rewritten

$$f_i(\mathbf{x}) = f(k_i q(\mathbf{x}) + n_{q_i}) + n_{f_i} \quad (6)$$

where each image has associated with it a separate realization of a quantigraphic noise process n_q and an image noise process n_f which includes noise introduced by the electronics of the dynamic range compressor f , and other electronics in the camera that affect the signal *after* its dynamic range has been compressed. In the case of a digital camera, n_f also includes quantization noise (applied after the image has undergone dynamic range compression). Furthermore, in the case of a camera that produces a data-compressed output, such as the Kodak DC260 which produces JPEG images, n_f also includes data-compression noise (JPEG artifacts, etc., which are also applied to the signal after it has undergone dynamic range compression). Refer again to Fig. 1.

If it were not for noise, we could obtain the photoquantity q from any one of a plurality of differently exposed pictures of the same subject matter, e.g. as

$$q = \frac{1}{k_i} f^{-1}(f_i) \quad (7)$$

where the existence of an inverse for f follows from the semimonotonicity assumption. Semimonotonicity follows from the fact that we expect pixel values to either increase or stay the same with increasing quantity of light falling on the image sensor.¹ However, because of noise, we obtain an advantage by capturing multiple pictures that differ only in exposure. The dark (“underexposed”) pictures show us highlight details of the scene that would have been overcome by noise (e.g., washed out) had the picture been “properly exposed.” Similarly, the light pictures show us some shadow detail that would not have appeared above the noise threshold had the picture been “properly exposed.”

Each image thus provides us with an estimate of the actual photoquantity q

$$q = \frac{1}{k_i} (f^{-1}(f_i - n_{f_i}) - n_{q_i}) \quad (8)$$

where n_{q_i} is the quantigraphic noise associated with image i , and n_{f_i} is the image noise for image i . This estimate of q , \hat{q}_i may be written

$$\hat{q}_i = \frac{1}{\hat{k}_i} \hat{f}^{-1}(f_i) \quad (9)$$

where \hat{q}_i is the estimate of q based on considering image i , and \hat{k}_i is the estimate of the exposure of image i based on considering a plurality of differently exposed images. The estimated

\hat{q}_i is also typically based on an estimate of the camera response function f , which is also based on considering a plurality of differently exposed images. Although we could just assume a generic function $f(q) = q^{0.45}$, in practice, f varies from camera to camera. We can, however, make certain assumptions about f that are reasonable for most cameras, such as the fact that f does not decrease when q is increased (that f is semimonotonic), and that it is usually smooth, and that $f(0) = 0$. In what follows, it will be shown how k and f are estimated from multiple differently exposed pictures. For the time being, let us suppose that they have been successfully estimated, so that we can calculate \hat{q}_i from each of the input images i . Such calculations, for each input image i , give rise to a plurality of estimates of q , which in theory would be identical, were it not for noise. However, in practice, because of noise, the estimates \hat{q}_i are each corrupted *in different ways*. Therefore it has been suggested, that multiple differently exposed images may be combined together to provide a single estimate of q which can then be turned into an image of greater dynamic range, greater tonal resolution, and lesser noise [6], [12]. In particular, the criteria under which collective processing of multiple differently exposed images of the same subject matter will give rise to an output image which is acceptable at every point (x, y) in the output image, are summarized below.

Wyckoff signal/noise criteria: $\forall (x_0, y_0) \in (x, y)$, $\exists k_i q(x_0, y_0)$ such that

- 1) $k_i q(x_0, y_0) \gg n_{q_i}$;
- 2) $c_i(q(x_0, y_0)) \gg c_i((1/k_i)f^{-1}(n_{f_i}))$.

The first criterion indicates that for every pixel in the output image, at least one of the input images provides sufficient exposure at that pixel location to overcome sensor noise, n_{q_i} . The second criterion states that of those at least one input images, at least one of them provides an exposure that falls favorably (e.g. is neither overexposed nor underexposed) on the response curve of the camera, so as not to be overcome by camera noise n_{f_i} .

The manner in which differently exposed images of the same subject matter are combined is illustrated, by way of an example involving three input images, in Fig. 4.

Moreover, it has been shown [11] that the constants k_i as well as the unknown nonlinear response function of the camera can be determined, up to a single unknown scalar constant, given nothing more than two or more pictures of the same subject matter, in which the pictures differ only in exposure. Thus the reciprocal exposures used to tonally register (tonally align) the multiple input images are estimates, $1/\hat{k}_i$, in Fig. 4. These exposure estimates are generally made by applying an estimation algorithm to the input images, either while simultaneously estimating f , or as a separate estimation process (since f only has to be estimated once for each camera but the exposure k_i is estimated for every picture i that is taken).

Owing to the large dynamic range that some Wyckoff sets can cover, small errors in f tend to have adverse effects on the overall estimate \hat{q} . Thus it may be preferable to estimate f as a separate process (e.g. by taking hundreds of exposures with the camera under computer program control). Once f is known (previously measured), then k_i can be estimated for a particular set of images.

¹Except in rare instances where the illumination is so intense as to damage the imaging apparatus, as, for example, when the sun burns through photographic negative film and appears black in the final print or scan.

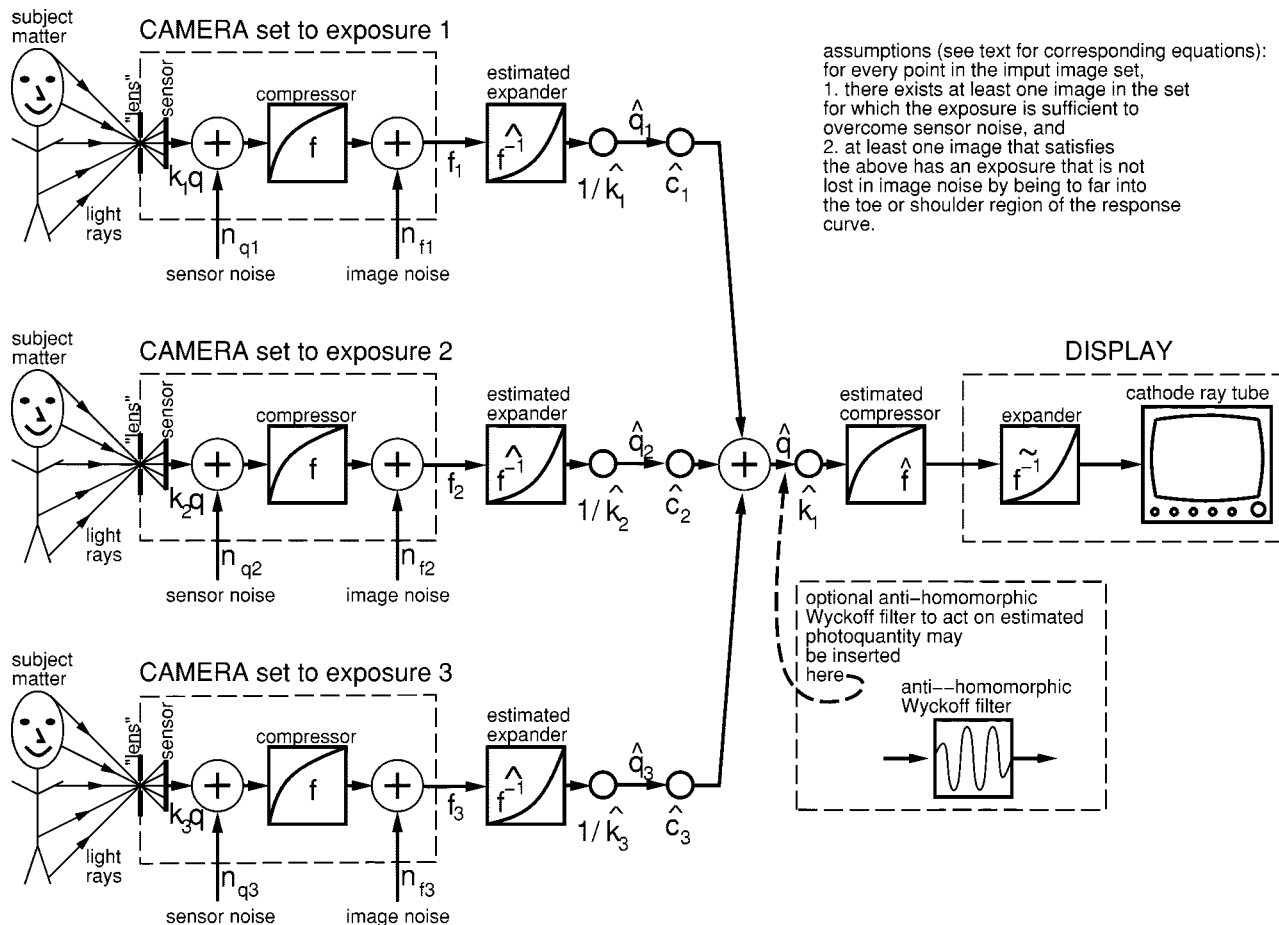


Fig. 4. *Wyckoff principle*: Multiple differently exposed images of the same subject matter are captured by a single camera. In this example there are three different exposures. The first exposure (CAMERA set to exposure 1), gives rise to an exposure $k_1 q$, the second to $k_2 q$ and the third to $k_3 q$. Each exposure has a different realization of the same noise process associated with it, and the three noisy pictures that the camera provides are denoted f_1 , f_2 , and f_3 . These three differently exposed pictures comprise a noisy Wyckoff set. In order to combine them into a single estimate, the effect of f is undone with an estimate, \hat{f} that represents our best guess of what the function f is. While many video cameras use something close to the standard $f = kq^{0.45}$ function, it is preferable to attempt to estimate f for the specific camera in use. Generally this estimate is made together with an estimate of the exposures k_i . After re-expanding the dynamic ranges with \hat{f}^{-1} , the inverse of the estimated exposures $1/\hat{k}_i$ are applied. In this way, the darker images are made lighter and the lighter images are made darker so that they all (theoretically) match. At this point the images will all appear as if they were taken with identical exposure, except for the fact that the pictures that were brighter to start with will be noisy in lighter areas of the image and those that had been darker to start with will be noisy in dark areas of the image. Thus rather than simply applying ordinary *signal averaging*, a weighted average is taken. The weights are the spatially varying *certainty functions*, $c_i(x, y)$. These certainty functions turn out to be the derivative of the camera response function shifted up or down by an amount k_i . In practice, since f is an estimate, so is c_i , so it is denoted \hat{c}_i in the figure. The weighted sum is $\hat{q}(x, y)$, the estimate of the photoquantity $q(x, y)$. To view this quantity on a video display, it is first adjusted in exposure, and may be adjusted to a different exposure level than any of the exposure levels used in taking the input images. In this figure, for illustrative purposes, it is set to the estimated exposure of the first image, \hat{k}_1 . The result is then range-compressed with \hat{f} for display on an expansive medium (DISPLAY).

The final estimate for q , depicted in Fig. 4, is given by

$$\hat{q}(x, y) = \frac{\sum_i \hat{c}_i \hat{q}_i}{\sum_i \hat{c}_i} = \frac{\sum_i \frac{\hat{c}_i(\hat{q}(x, y))}{\hat{k}_i} \hat{f}^{-1}(f_i(x, y))}{\sum_i \hat{c}_i(\hat{q}(x, y))} \quad (10)$$

where \hat{c}_i is given by

$$\hat{c}_i(\log(q(x, y))) = \frac{df_i(x, y)}{d \log \hat{q}(x, y)} = \frac{d\hat{f}(\hat{k}_i \hat{q}(x, y))}{d \log \hat{q}(x, y)} \quad (11)$$

from which we can see that $c_i(\log(q))$ are just shifted versions of $c(\log(q))$, e.g. dilated versions of $c(q)$.

The intuitive significance of the certainty function is that it captures the slope of the response function which indicates

how quickly the output (pixel value or the like) of the camera varies for given input. In the case of a noisy camera, especially a digital camera where quantization noise is involved, generally the output of a camera will be most reliable where it is most sensitive to a fixed change in input light level. This point where the camera is most responsive to changes in input is at the peak of the certainty function c . The peak in c tends to be near the middle of the camera's exposure range. On the other hand, where the camera exposure input is extremely large or small (e.g. the sensor is very much overexposed or very much underexposed), the change in output for a given input is much less. Thus the output is not very responsive to the input and the change in output can be easily overcome by noise. Thus c tends to fall off toward zero on either side of its peak value.

The certainty functions are functions of q . We may also write the uncertainty functions, which are functions of pixel value in

the image (e.g., functions of greyvalue in f_i), as

$$U(x, y) = \frac{dF(f_i(x, y))}{df_i(x, y)} \quad (12)$$

and its reciprocal is the certainty function C in the domain of the image (e.g., the certainty function in *pixel coordinates*)

$$C(x, y) = \frac{df_i(x, y)}{dF(f_i(x, y))} \quad (13)$$

where $F^{-1} = \log f$. Note that C is the same for all images (e.g. for all values of image index i), whereas c_i was defined separately for each image. For any i , the function c_i is a shifted (dilated) version of any other certainty function, c_j , where the shift (dilation) depends on the log exposure, K_i (the exposure k_i).

The final estimate of q (10) is simply a weighted sum of the estimates from q obtained from each of the input images, where each input image is weighted by the certainties in that image.

1) *Exposure interpolation and extrapolation*: The architecture of this process is shown in Fig. 5, which depicts an image acquisition section (in this illustration, of three images), followed by an analysis section (to estimate q), followed by a resynthesis section to generate an image again at the output (in this case four different possible output images are shown).

The output image can look like any of the input images, but with improved signal to noise ratio, better tonal range, better color fidelity, etc. Moreover, an output image can be an interpolated or extrapolated version in which it is lighter or darker than any of the input images. It should be noted that this process of interpolation or extrapolation provides a new way of adjusting the tonal range of an image. The process is illustrated in Fig. 5. The image synthesis portion may also include various kinds of deblurring operations, as well as other kinds of image sharpening and lateral inhibition filters to reduce the dynamic range of the output image without loss of fine details, so that it can be printed on paper or presented to an electronic display in such a way as to have optimal tonal definition.

III. COMPARAMETRIC IMAGE PROCESSING: COMPARING DIFFERENTLY EXPOSED IMAGES OF THE SAME SUBJECT MATTER

As previously mentioned, comparison of two or more differently exposed images may be done to determine q , or simply to tonally register the images without determining q . Also, as previously mentioned, tonal registration is more numerically stable than estimation of q , so there are some advantages to comparometric analysis and comparometric image processing in which one of the images is selected as a reference image, and others are expressed in terms of this reference image, rather than in terms of q . Typically the dark images are lightened, and/or the light images are darkened so that all the images match the selected reference image. Note that in such lightening and darkening operations, full precision is retained for further comparometric processing. Thus all but the reference image will be stored as an array of floating point numbers.

A. Misconceptions About Gamma Correction

So-called *gamma correction* (raising the pixel values in an image to an exponent) is often used to lighten or darken images. While gamma correction does have important uses, such as lightening or darkening images to compensate for incorrect display settings, it will now be shown that when one uses gamma correction to lighten or darken an image to compensate for incorrect exposure, that whether one is aware of it or not, one is making an unrealistic assumption about the camera response function.

Proposition III.1: Tonally registering differently exposed images of the same subject matter by gamma correcting them with exponent $\gamma = k^\Gamma$ is equivalent to assuming that the nonlinear response function of the camera is $f(q) = \exp(q^\Gamma)$.

Proof: The process of gamma correcting an image may be written

$$g(q) = f(kq) = (f(q))^\gamma \quad (14)$$

where f is the original image, and g is the lightened or darkened image. Solving for f , the camera response function, we obtain

$$f(q) = \exp(q^\Gamma). \quad (15)$$

□

We see that the response function (15) does not pass through the origin, e.g. $f(0) = 1$, not zero. Since most cameras are designed so that they produce a signal level output of zero when the light input is zero, the function $f(q)$ does not correspond to a realistic or reasonable camera response function. Even media which does not itself fall to zero at zero exposure (like film, for example) is ordinarily scanned in such a way that the scanned output is zero for zero exposure, assuming that the d_{\min} (minimum density for the particular emulsion being scanned) is properly set in the scanner. Therefore it is inappropriate and incorrect to use gamma correction to lighten or darken differently exposed images of the same subject matter, when the goal of this lightening or darkening is tonal registration (making them look the “same,” apart from the effects of noise which will be accentuated in the shadow detail of the images that are lightened and the highlight detail of images that are darkened).

B. Comparometric Plots and Comparometric Equations

To understand the shortcomings of gamma correction, and to understand some alternatives, the concept of comparometric equations and comparometric plots will now be introduced.

Equation (14) is an example of what is called a *comparometric equation* [17].

Comparometric equations are a special case of the more general class of equations called *functional equations* [18] and *comparometric plots* are a special case of the more general class of plots called *parametric plots*.

The notion of a parametric plot is well-understood. For example, the parametric plot $(r \cos(q), r \sin(q))$ is a plot of a circle of radius r . Note that it does not depend explicitly on q , so long as the domain of q includes at least all points on the interval from 0 to 2π , modulo 2π .

A *comparometric plot* is a special kind of parametric plot in which a function f is plotted against itself, and in which the




Fig. 5. Quantigraphic image exposure adjustment on a Wyckoff set: Multiple (in this example, 3) differently exposed (in this example by two aperture stops) images are acquired. Estimates of q from each image are obtained. These are combined by weighted sum. The weights are the estimates of the certainty function shifted along the exposure axis by the estimated exposure for each image. From the estimated photoquantity \hat{q} , one or more output images may be generated by multiplying by the desired synthetic exposure and passing the result through the estimated camera nonlinearity. In this example, four synthetic pictures are generated. These are extrapolated and interpolated versions of the input exposures. The result is a “virtual camera” [17] in which a picture can be generated as if the user were free to select the original exposure settings that had been used on the camera originally taking the input images.

parameterization of the ordinate is a linearly scaled parameterization of the abscissa.

More precisely, the comparametric plot is defined as follows:

Definition III.1: A plot along coordinates $(f(q), f(kq))$ is called a *comparametric plot* [17] of the function $f(q)$. □

A function $f(q)$ has a family of comparametric plots, one for each value of the constant k , which is called the *comparametric ratio*.

Proposition III.2: When a function $f(q)$ is monotonic, the comparametric plot $(f(q), f(kq))$ can be expressed as a monotonic function $g(f)$ not involving q . □

Thus the plot in Definition III.1 may be rewritten as a plot $(f, g(f))$, not involving q . In this form, the function g is called the *comparametric function*, and expresses the range of the function $f(kq)$ as a function of the range of the function $f(q)$, independently of the domain, q , of the function f .

The plot g defines what is called a *comparametric equation*:

Definition III.2: Equations of the form $g(f(q)) = f(kq)$ are called comparametric equations [17]. \square

A better understanding of comparametric equations may be had by referring to the following diagram:

$$\begin{array}{ccccc}
 & & k & & \\
 & q & \longrightarrow & kq & \\
 f & \downarrow & & \downarrow & f \\
 & f(q) & \longrightarrow & f(kq) & \\
 & & g & &
 \end{array} \quad (16)$$

wherein it is evident that there are two equivalent paths to follow from q to $f(kq)$

$$g \circ f = f \circ k. \quad (17)$$

Equation (17) may be rewritten

$$g = f \circ k \circ f^{-1} \quad (18)$$

which provides an alternative definition of *comparametric equation* to that given in Definition III.2.

Equation (14) is an example of a comparametric equation, and (15) is a solution of (14).

It is often preferable that comparametric equations be on the interval from zero to one in the range of f . Equivalently stated, we desire comparametric equations to be on the interval from zero to one in the domain of g and the range of g . In this case, the corresponding plots and equations are said to be *unicomparametric*. (Actual images typically range from 0 to 255 and must thus be rescaled so that they range from 0 to 1, for unicomparametric image processing.)

Often we also impose a further constraint that $f(0) = 0$, and the constraint that $g(0) = 0$ and $g(1) = 1$.

Solving a comparametric equation is equivalent to determining the unknown camera response function from a pair of images that differ only in exposure, when the comparametric equation represents the relationship between greyvalues in the two pictures, and the comparametric ratio, k , represents the ratio of exposures (e.g., if one picture was given taken with twice the exposure of the other, then $k = 2$).

C. Zeta Correction of Images

An alternative to *gamma correction* is proposed. This alternative, called *zeta correction*, will also serve as another example of a comparametric equation.

For zeta correction, we simply adjust the exponential solution (15) of the comparametric equation given by traditional gamma correction so that the solution passes through the origin: $f(q) = \exp(q - 1)$. (For simplicity, and without loss of generality, γ has been set to k , the comparametric exposure ratio.) Thus, $g = (bf + 1)^\gamma = (bf + 1)^k$. Preferably (to be unicomparametric) we would like to have $g(1) = 1$, so we use the response function

$$f(q) = \frac{e^q - 1}{\sqrt[k]{2} - 1} \quad (19)$$

which is a solution to the corresponding comparametric equation

$$g = ((\sqrt[k]{2} - 1)f + 1)^k - 1, \quad \forall k \neq 0. \quad (20)$$

The comparametric equation (20) forms the basis for zeta correction of images

$$g = \begin{cases} ((2^\zeta - 1)f + 1)^{1/\zeta} - 1, & \forall \zeta \neq 0 \\ 2^f - 1, & \text{for } \zeta = 0. \end{cases} \quad (21)$$

Implicit in zeta correction of images is the assumption of an exponential camera response function, which, although not realistic (given that the exponential function expands dynamic range, and most cameras have compressive response functions rather than expansive response functions), is preferable to gamma correction because of the implicit notion of a response function for which $f(0) = 0$.

With standard IEEE arithmetic, values of ζ can range from approximately -50 to $+1000$.

D. Affine Comparametric Equation and Affine Correction of Images

In this section, one of the two most useful (in the author's opinion) comparametric equations is introduced, giving rise to *affine correction* of images. Affine correction is an improvement over that of zeta correction (which itself was an improvement over gamma correction).

First consider the classic model

$$f(q) = \alpha + \beta q^\gamma \quad (22)$$

used by photographers to characterize the response of a variety of photographic emulsions, including so-called *extended response* film [9]. It is well known that (22) becomes the equation of a straight line when expressed in logarithmic coordinates, if we subtract α (as many scanners such as PhotoCD attempt to do by prompting the user to scan a piece of blank film from the film trailer before scanning the rest of the roll of film)

$$\log(f(q) - \alpha) = \gamma \log(q) + \beta. \quad (23)$$

It is an interesting coincidence that the comparametric plot of this function (22) is also a straight line.

Proposition III.3: The comparametric plot corresponding to the standard photographic response function (22) is a straight line. The slope is k^γ , and the intercept is $\alpha(1 - k^\gamma)$.

Proof: $g(f(kq)) = f(kq) = \alpha + \beta(kq)^\gamma$ Re-arranging to eliminate q gives $g = k^\gamma(\alpha + \beta q^\gamma) + \alpha(1 - k^\gamma)$ so that

$$g = k^\gamma f + \alpha(1 - k^\gamma). \quad (24)$$

\square

Note that the constant β does not appear in the comparametric equation. Thus we cannot determine β from the comparametric equation. The physical (intuitive) interpretation is that we can only determine the nonlinear response function of a camera up to a single unknown scalar constant.

Note that (14) looks quite similar in form to (22), and in fact is identical if we set $\alpha = 0$ and $\beta = 1$. However, one must recall that (14) is a comparametric equation and (22) is a solution to

a (different) comparametric equation. Thus we must be careful not to confuse the two. The first corresponds to gamma correction of an image, while the second corresponds to the camera response function that is implicit in applying (24) to lighten or darken the image. To make this distinction clear, applying (24) to lighten or darken an image will be called *affine correction* (e.g. correcting by modeling the comparametric function with a straight line). The special case of *affine correction* when the intercept is equal to zero will be called *linear correction*.

Preferably *affine correction* of an image also includes a step of clipping values greater than 1 to 1, and values less than zero to zero, in the output image

$$g = \min(\max(k^\gamma f + \alpha(1 - k^\gamma), 0), 1). \quad (25)$$

If the intercept is zero and the slope is greater than one, the effect, neglecting noise, of (25), is to lighten the image in a natural manner that properly simulates the effect of having exposed the image with greater exposure. In this case, the effect is theoretically identical to that which would have been obtained by using a greater exposure on the camera, assuming the response function of the camera follows the power law $f = q^\gamma$, as many cameras do in practice. Thus it has been shown that the correct way to lighten an image is to apply *linear correction*, not *gamma correction* (apart from correction of an image to match an incorrectly adjusted display device or the like, where gamma correction is still the correct operation to apply).

Here we have worked forward, starting with the solution (22) and deriving the comparametric equation (24) of which (22) is a solution. It is much easier to generate comparametric equations from their solutions than it is to solve comparametric equations.

The above comparametric equation is both useful and simple. The simplicity is in the ease with which it is solved, and by the fact that the solution happens to be the most commonly used camera response model in photography. As we will later see, when processing images, the comparametric function can be estimated by fitting a straight line through data points describing the comparametric relation between images. However, there are two shortcomings to *affine correction*.

- 1) It is not inherently unicomparametric so it must be clipped to 1 when it exceeds 1 and clipped to zero when it falls below zero, as shown in (25).
- 2) Its solution, $f(q)$ only describes the response of cameras within their normal operating regime. Since the art of quantigraphic image processing involves a great deal of image processing done on images that have been deliberately and grossly overexposed or underexposed, there is a need for a comparametric model that captures the essence of cameras at both extremes (e.g., both overexposure and underexposure) of exposure.

E. Preferred Correction of Images

Although *affine correction* was an improvement over *zeta correction*, which itself was an improvement over *gamma correction*, *affine correction* still has the two shortcomings listed above. Therefore another form of image exposure correction is proposed, and it will be called the preferred correction. This new

exposure correction is unicomparametric (bounded in normalized units between 0 and 1) and also has a parameter to control the softness of the transition into the *toe* and *shoulder* regions of the response function, rather than the hard clipping introduced by (25).

As with affine correction, *the preferred correction* will be introduced first by its solution, from which the comparametric equation will be derived. The solution is

$$f(q) = (e^b q^a / (e^b q^a + 1))^c \quad (26)$$

which has only three parameters. Thus no extra unnecessary degrees of freedom (which might otherwise capture or model noise) have been added over and above the number of degrees of freedom in the previous model (22).

An intuitive understanding of (26) can be better had by rewriting it

$$f = \begin{cases} \exp(1/(1 + e^{-(a \log(q)+b)})^c), & \forall q \neq 0 \\ 0, & \text{for } q = 0 \end{cases} \quad (27)$$

where the soft transition into the toe (region of underexposure) and shoulder (region of overexposure) is evident by the shape of this curve on a logarithmic exposure scale.

This model may, at first, only seem like a slight improvement over (22), given our common intuition that most exposure information is ordinarily captured in the central portion that is linear on the logarithmic exposure plot. However, it is important that we unlearn what we have been taught in traditional photography, where incorrectly exposed images are ordinarily thrown away rather than used to enhance the other images! It must be emphasized that comparametric image processing differs from traditional image processing in the sense that in comparametric image processing (using the Wyckoff principle, as illustrated in Fig. 4 the images typically include some that are *deliberately* underexposed and overexposed. In fact this overexposure of some images and underexposure of other images is often deliberately taken to extremes. Therefore, the additional sophistication of the model (26) is of great value in capturing the essence of a set of images where some extend to great extremes in the *toe* or *shoulder* regions of the response function.

Proposition III.4: The comparametric equation of which the proposed photographic response function (26) is a solution, is given by

$$g(f) = \frac{f k^{ac}}{(\sqrt[ac]{f(k^a - 1) + 1})^c} \quad (28)$$

where $K = \log(k)$. This function (28) gives rise to *the preferred correction* of images (e.g., the preferred recipe for lightening or darkening an image). Again, $g(f)$ does not depend on b , which is consistent with our knowledge that the comparametric equation captures the information of $f(q)$ up to a single unknown scalar proportionality constant. \square

F. Some Solutions to Some Comparametric Equations That Are Particularly Illustrative or Useful

Some examples of comparametric equations and their solutions are summarized in Table I.

TABLE I

ILLUSTRATIVE OR USEFUL EXAMPLES OF COMPARAMETRIC EQUATIONS AND THEIR SOLUTIONS. THE THIRD FROM THE TOP AND SECOND FROM THE BOTTOM WERE FOUND TO DESCRIBE A LARGE VARIETY OF CAMERAS AND HAVE BEEN USED IN A WIDE VARIETY OF QUANTIGRAPHIC IMAGE PROCESSING APPLICATIONS. THE SECOND ONE FROM THE BOTTOM IS THE ONE THAT IS MOST COMMONLY USED BY THE AUTHOR

comparametric equation	solution (camera response function)
$g(f(q)) = f(kq)$	$f(q)$
$g = f^\gamma$	$f = \exp(q^\gamma), \gamma = k^\gamma$
$g = k^\gamma f$	$f = q^\gamma$
$g = af + b, \forall a \neq 1, \text{ or } b = 0$	$f = \alpha + \beta q^\gamma, a = k^\gamma, b = \alpha(1 - k^\gamma)$
$g = f + b$	$f = B \log(\beta q), b = B \log k$
$g = (\sqrt[\zeta]{f} + \log k)^\gamma$	$f = \log^\gamma q$
$g = \begin{cases} ((2^\zeta - 1)f + 1)^{\frac{1}{\zeta}} - 1, & \forall \zeta \neq 0 \\ 2^f - 1, & \text{for } \zeta = 0 \end{cases}$	$\frac{e^q - 1}{2^\zeta - 1}$
$g = e^b f^a = e^{\alpha(1-k^\gamma)} f^{(k^\gamma)}$	$\log f = \alpha + \beta q^\gamma$
$g = \exp((\log f)^{k^b})$	$f = \exp(a^{(q^b)})$
$g = \exp(\log^k f)$	$f = \exp(a^{q^k})$
$g = \frac{2}{\pi} \arctan(k \tan(\frac{\pi}{2} f))$	$f = \frac{2}{\pi} \arctan(q)$
$g = \frac{1}{\pi} \arctan(b\pi \log k + \tan((f - 1/2)\pi)) + \frac{1}{2}$	$f = \begin{cases} \frac{1}{\pi} \arctan(b\pi \log q) + \frac{1}{2}, & \forall q \neq 0 \\ 0, & \text{for } q = 0 \end{cases}$
$g = \left(\frac{\sqrt{f} k^a}{\sqrt{f(k^a-1)+1}} \right)^c$	$f = \left(\frac{e^b q^a}{e^b q^a + 1} \right)^c = \begin{cases} \left(\frac{1}{1+e^{-\frac{1}{c}(\log q + b)}} \right)^c, & \forall q \neq 0 \\ 0, & \text{for } q = 0 \end{cases}$
$g = \exp\left(\left(\frac{\sqrt{\log f} k^a}{\sqrt{\log f(k^a-1)+1}} \right)^c \right)$	$f = \exp\left(\left(\frac{e^b q^a}{e^b q^a + 1} \right)^c \right) = \begin{cases} \exp\left(\left(\frac{1}{1+e^{-\frac{1}{c}(\log q + b)}} \right)^c \right), & \forall q \neq 0 \\ 0, & \text{for } q = 0 \end{cases}$

G. Properties of Comparametric Equations

As stated previously, the comparametric equation only provides information about the actual photoquantity up to a single unknown scalar quantity, e.g., if $f(q)$ is a solution of comparametric equation g then so is $f(\beta q)$. In general we can think of this as a coordinate transformation from q to βq , in the domain of f . More generally, the comparametric plot $(f(q), g(f(q)))$ has the same shape as the comparametric plot $(f(h(q)), g(f(h(q))))$, for all bijective h . From this fact, we can construct a property of comparametric equations in general:

Proposition III.5: A comparametric equation $\check{g}(f(q)) = g(f(h(q))) = g(\check{f}(q))$ has solution $\check{f}(q) = f(h(q))$, for any bijective function h . □

Likewise, we can also consider coordinate transformations in the range of comparametric equations, and their effects on the solutions.

Proposition III.6: A comparametric equation $g(f) = h(\check{g})$, has solution $f = h(\check{f})$, where $\check{g}(q) = \check{f}(kq)$ is a comparametric equation with solution $\check{f}(q)$. □

Properties of comparametric equations related to their coordinate transformations are presented in Table II.

Some simple but illustrative examples of using coordinate transformation properties to solve comparametric equations are now provided.

Example 1: Let $\check{g} = a\check{f} + b$, which we know has solution $\check{f} = \alpha + \beta q^\gamma$, with $a = k^\gamma$ and $b = \alpha(1 - a)$. Let $h(\cdot) = (\cdot) + n_f$, e.g. h is a transformation which consists of simply adding noise. Thus $g(f) = h(\check{g}) = \check{g} + n_f = a\check{f} + b + n_f$, so that $g = a(f - n_f) + b + n_f$ has solution $f = h(\check{f}) = \alpha + \beta q^\gamma + n_f$. □

Example 2: From Table I, observe that the comparametric equation $\check{g} = a\check{f} + b$ has solution $\check{f} = \alpha + \beta q^\gamma$. Let $h(\cdot) = \exp(\cdot)$. We can thus solve $g(f) = h(\check{g}) = \exp(a\check{f} + b) = \exp(a \log(f) + b) = e^b f^a$ by noting that $f = h(\check{f}) = \exp(\check{f}) = \exp(\alpha + \beta q^\gamma)$. □

TABLE II

SOME PROPERTIES OF COMPARAMETRIC EQUATIONS. THIS TABLE COULD BE EXTENDED OVER SEVERAL PAGES, MUCH LIKE AN EXTENSIVE TABLE LISTING PROPERTIES OF LAPLACE TRANSFORMS, OR A TABLE OF PROPERTIES OF FOURIER TRANSFORMS, OR THE LIKE

comparametric equations	solutions (camera response functions)
$g(f(q)) = f(kq)$	$f(q)$
$\check{g}(f) = g(\check{f})$, where $\check{g}(f(q)) = g(f(h(q)))$	$\check{f}(q) = f(h(q)), \forall$ bijective h
$g(f) = \check{g}(\check{f})$, where $\check{g}(f(q)) = \check{f}(\beta q)$	$f(q) = \check{f}(\beta q)$
$g(f) = g(h(\check{f})) = h(\check{g})$, where $\check{g}(q) = f(kq)$	$f = h(\check{f})$
$h^{-1}(g) = \check{g}(\check{f})$	$f = h(\check{f})$

This solution also appears in Table I. We may also use this solution to seed the solution of the comparametric equation second from the bottom of Table I, by using $h(x) = x/(x + 1)$. The equation second from the bottom of Table I may then be further coordinate transformed into the equation at the bottom of Table I by using $h(x) = \exp(x)$. Thus properties of comparametric equations, such as those summarized in Table II, can be used to help solve comparametric equations, such as those listed in Table I.

IV. PRACTICAL IMPLEMENTATIONS

This section pertains to the practical implementation of the theory presented in previous sections.

A. Comparing Two Images That Differ Only in Exposure

Without loss of generality, consider two differently exposed pictures of the same subject matter, f_1 and f_2 , and recognize that in the absence of noise, the relationship between the two images would be

$$\frac{1}{k_1} f_1^{-1}(f_1) = q = \frac{1}{k_2} f_2^{-1}(f_2) \tag{29}$$

so that

$$f_2 = f(kf^{-1}(f_1)) = F^{-1}(F(f_1) + K) \quad (30)$$

where $k = k_2/k_1$, $K = \log(k_2) - \log(k_1)$, and $F = \log(f^{-1})$. It is evident that (30) is a comparametric equation.

This process (30) of “registering” the second image with the first differs from the image registration procedure commonly used in much of machine vision [19]–[22] and image resolution enhancement [4], [8], [7] because it operates on the *range*, $f(q(\mathbf{x}))$, (tonal range) of the image $f_i(\mathbf{x})$ as opposed to its *domain* (spatial coordinates) $\mathbf{x} = (x, y)$.

B. Comparagram

The comparagram is defined as the joint histogram between two images that differ only in exposure. The comparagram is a square matrix of size N by N , where N is the number of greylevels [11], [12]. It can be seen from (30) that the general problem of solving (30) can be done directly on the comparagram instead of the original pair of images. This comparametric approach has the added advantage of breaking the problem down into two separate simpler steps:

- 1) *comparametric regression*: finding a smooth semimonotonic function, g , that passes through most of the highest bins in the comparagram.
- 2) *solving the comparametric equation*: unrolling this function, $g(f(q)) = f(kq)$ into $f(q/q_0)$, by regarding it an *iterative map* onto itself (see Fig. 6.) The iterative map (*logistic map*) is most familiar in chaos theory [23], [24], but here, since the map is monotonic, the result is a deterministic function.

Separating this estimation process into two stages also allows us a more direct route to “registering” the image domains, if for example, we do not need to know f , but only require g , which is the recipe for expressing the range of $f(kq)$ in the units of $f(q)$. In particular, we can lighten or darken images to match one another without ever having to solve for q . The first part of the above two step process allows us to determine the relationship between two pictures that differ only in exposure, so that we can directly perform operations like image exposure interpolation and extrapolation as in Fig. 5, but skipping the intermediate step of computing q . Not all image processing applications require determining q , so there is great value in simply understanding the relationship between differently exposed pictures of the same subject matter.

C. Comparametric Regression and the Comparagram

In situations where the image data is extremely noisy, and/or where a closed-form solution for $f(q)$ is desired, a parameterized form of the comparametric function is used, in which a function $g(f)$ corresponding to a suitably parameterized response function $f(q)$ is selected. The method amounts to a *curve fitting* problem in which the parameters of g are selected so that g best fits one or more comparagrams constructed from two or more differently exposed images under analysis.




Fig. 6. Comparametric procedure for finding the pointwise nonlinearity of an image sensor from two pictures differing only in their exposures. (COMPARAMETRIC PLOT) Plot of pixel values in one image against corresponding pixel values in the other. (RESPONSE CURVE) Points on the response curve, found from only the two pictures, without any knowledge about the characteristics of the image sensor. These discrete points are only for illustrative purposes. If a logarithmic exposure scale is used, (as most photographers do) then the points fall uniformly on the $Q = \log(q/q_0)$ axis.

D. Comparametric Regression to a Straight Line

The result (24) suggests that $f(q)$ can be determined from two differently exposed images by applying linear regression to the comparagram of the images, J , treating each entry as a datapoint, and weighting by the number of bin counts $J(m, n)$ at each point. Often this is done by weighting with $(J(m, n))^\lambda$. For example, $\lambda = 0$ (assuming empty bins are not counted) provides the classic linear regression problem in which all nonempty bins are weighted equally and the slope and intercept of the best-fit line through nonempty bins is found. Generally λ is chosen somewhere between 1/4 and 2.

A simple example is presented, that of reverse-engineering the standard Kodak PhotoCD scanner issued to most major photographic processing and scanning houses. In most situations, a human operator runs the machine, and decides, by visual inspection, what “brightness” level to scan the image at (there is also an automatic exposure feature which allows the operator to preview the scanned image and decide whether or not the chosen “brightness” level needs to be overridden). By scanning the same image at different “brightness” settings, a Wyckoff set results. This allows the scanner to capture nearly the entire dynamic range of the film, which is of great utility since typical photographic negative film captures far greater dynamic range than possible with the scanner as it is ordinarily used. A photographic negative taken from a scene of extremely high contrast (a sculpture on exhibit at the List Visual Arts Center, in a completely darkened room, illuminated with a bare flash lamp from one side only) was selected because of its great dynamic range that could not be captured in any single scan. A Wyckoff set was constructed by scanning the same negative at five different “brightness” settings (Fig. 7). The settings were controlled by a slider that was calibrated in arbitrary units from -99 to $+99$, while running Kodak’s proprietary scanning software. Kodak provides no information about what these units mean. Accordingly, the goal of the experiment was to find a closed-form mathematical equation describing the effects of the “brightness” slider on the scans, and to recover the unknown nonlinearity of the scanner. In order to make the problem a little more challenging and, more importantly, to better-illustrate the principles of comparametric image




Fig. 7. These scans from a photographic negative differ only in the choice of “brightness” setting selected using the slider provided on the X-windows screen by the proprietary Kodak PhotoCD scanning software. The slider is calibrated in arbitrary units from -99 to $+99$. Five scans were done and the setting of the slider is noted above each scan.




Fig. 8. Pairwise comparagrams of the images in Fig. 8. It is evident that the data are well fitted by a straight line, which suggests that Kodak must have used the standard nonlinear response function $f(q) = \alpha + \beta q^\gamma$ in the design of their PhotoCD scanner.

processing, the *dmin* procedure of scanning a blank film at the beginning of the roll was overridden.

Comparagrams J_{01} , J_{12} , J_{23} , and J_{34} were computed from the five images (v_0 through v_4) of Fig. 7, and are displayed as density plots (e.g. treating them as images of dimension 256 by 256 pixels, where the darkness of the image is proportional to the number of counts—darkness rather than lightness to make it easier to see the pattern) in Fig. 8. Linear regression was applied to the data, and the best-fit straight line is shown passing through the data points. Because the *dmin* procedure was overridden, notice that the plots do not pass through the origin. The two leftmost plots had nearly identical slopes and intercepts, and likewise for the two rightmost, which indicates that the arbitrary Kodak units of “brightness” are self-consistent (e.g. J_{01} which describes the relationship between a scan at a “brightness” of 40 units and one of 20 units is essentially the same as J_{12} which describes the relationship between a scan at a “brightness” of 20 units and one of 0 units). Since there are three parameters in (22), k , α , and γ , which describe only two degrees of freedom (slope and intercept), γ may be chosen so that $k = \sqrt[\gamma]{a}$ works out to be linearly proportional to arbitrary Kodak units. Thus, setting $(\sqrt[\gamma]{a_{\text{left}}})/(\sqrt[\gamma]{a_{\text{right}}}) = 20/30$ (where a_{left} is the average slope of the two leftmost plots and a_{right} the average slope of the two rightmost plots) results in the value $\gamma = 0.2254$. From this we obtain $\alpha = b/(1 - a) = 23.88$. Thus, we have that

$$f(k_i q) = 23.88 + (k_i q)^{0.2254} \quad (31)$$

where k_i is in arbitrary Kodak units (e.g. $k_0 = 40$ for the leftmost image, $k_1 = 20$ for the next image, $k_2 = 0$ for the

middle image, $k_3 = -30$, and $k_4 = -60$). Thus (31) gives us a closed-form solution that describes the response curve associated with each of the five exposures $f(k_i q)$, $i \in \mathbb{Z}$, $0 < i <= 4$. The curves $f(k_i q)$ may be differentiated, and, if these derivatives are evaluated at $q_i = (1/k_i) \sqrt[\gamma]{f_i(x, y) - \alpha}$, the so-called *certainty images*, shown in Fig. 9 are obtained.

In the next example, the use of the certainty functions to construct an optimal estimate, $\hat{q}(x, y)$ will be demonstrated.

E. Comparametric Regression to the Preferred Model

For this second example, the comparametric model proposed in (28) will be used.

In many practical situations, real-world images are very noisy.

Accordingly, an example of noisy images that comprise a Wyckoff set (Fig. 10), in which an extremely poor scan was deliberately used to scan images from a publication [11], is now considered.

That the images in Fig. 10 are of very poor quality is evidenced by their comparagram [Fig. 11(a)]. Using regression of (28) to the joint comparagram combined with the knowledge (from the publication from which the images were obtained [11]) that $K = 2$, it was found that $a = 0.0017$ and $c = -3.01$. This data provides a closed-form solution for the response function. The two effective response functions, which are shifted versions of this one response function, where the relative shift is K , are plotted in Fig. 12, together with their derivatives. (Recall that the derivatives of the response functions are the certainty functions.) Since a closed-form solution has been obtained, it may be easily differentiated without the further increase in noise




Fig. 9. *Certainty functions* express the rate of change of $f(q(x, y))$ with $Q(x, y)$. The certainty functions may be used to compute the *certainty images*, $f(c_i)$. White areas in one of the *certainty images* indicate that pixel values $f(q)$ change fastest with a corresponding change in the photoquantity, Q . When using the camera as a lightmeter (e.g., a quantigraphic instrument to estimate q), it will be most sensitive where the certainty images are white. White areas of these certainty images correspond to mid-grey values (*midtone*s) of the corresponding original images in Fig. 7, while dark areas correspond to extreme pixel values (either highlights or shadows) of the original images in Fig. 7. Black areas of the *certainty image* indicate that Q changes drastically with small changes in pixel value, and thus an estimate of Q in these areas will be overcome by image noise n_{f_i} .



(a)



(b)

Fig. 10. Noisy images badly scanned from a publication. These images are identical except for exposure and a good deal of quantization noise, additive noise, scanning noise, etc. (a) Darker image shows clearly the eight people standing outside the doorway, but shows little of the architectural details of the dimly lit interior. (b) Lighter image shows the architecture of the interior, but it is not even possible to determine how many people are standing outside, let alone recognize any of them.

that usually accompanies differentiation. Otherwise, when determining the certainty functions from poor estimates of f , the certainty functions would be even more noisy than the poor estimate of f itself. The resulting certainty images, denoted by $c(f_i)$, are shown in Fig. 13. Each of the images, $f_i(x, y)$ gives

rise to an actual estimate of the quantity of light arriving at the image sensor (9) These estimates were combined by way of (10) resulting in the composite image appears shown in Fig. 14.

Note that the resulting image \hat{f}_1 looks very similar to f_1 , except that it is a floating point image array, of much greater tonal range and image quality.

Furthermore, given a Wyckoff set, a composite image may be rendered at any in-between exposure from the set (exposure interpolation), as well as somewhat beyond the exposures given (exposure extrapolation). This result suggests the “VirtualCamera” [17] which allows images to be rendered at any desired exposure, once q is computed.

This capability is somewhat similar to QuickTime VR and other image-based rendering systems, except that it operates in the range of the images \hat{f}_i rather than their domain.

V. SPATIOTONAL QUANTIGRAPHIC FILTERS

Ordinarily, most print and display media have limited dynamic range. Thus one might be tempted to argue against the utility of the Wyckoff principle based on this fact, e.g. one might ask “since televisions and print media cannot display more than a very limited dynamic range, why bother building a Wyckoff camera that can capture such dynamic ranges?” Why bother capturing the photoquantity q with more accuracy than is needed for display?

Some possible answers to this question are as follows.

- 1) Estimates of q are still useful for machine vision, and other applications that do not involve direct viewing of a final picture. An example is the wearable face recognizer [25] which determines the identity of an individual from a plurality of differently exposed pictures of that person, and then presents the identity in the form of a text label (virtual name tag) on the retina of an eye of the wearer of the eyeglass-based apparatus. Since \hat{q} need not be displayed, the problem of output dynamic range, etc., of the display (e.g. number of distinct intensity levels of the laser beam shining into a lens of the eye of the wearer) is of no consequence.
- 2) Even though the ordinary dynamic range and the range resolution (typically 8 bits) is sufficient for print media

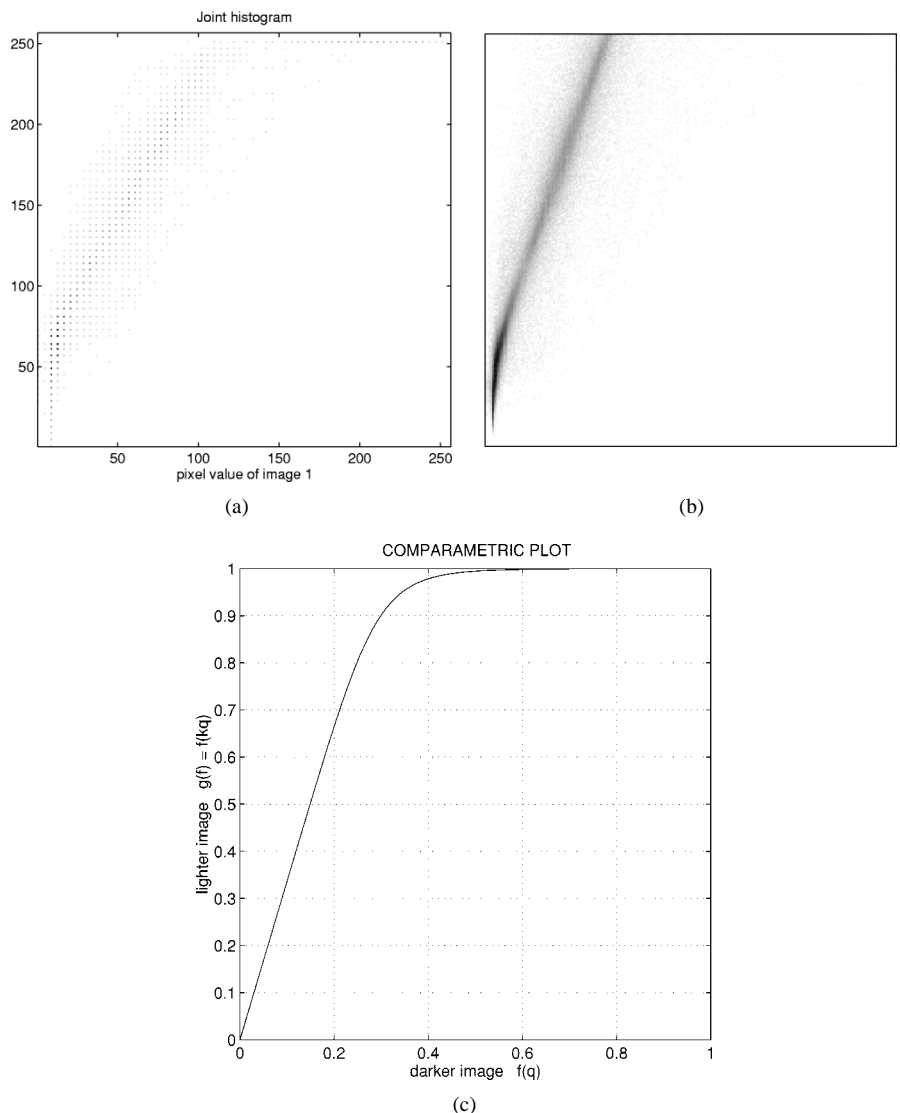


Fig. 11. *Comparametric regression*: (a) Comparagram. Note that because the images were extremely noisy, the comparagram is spread out over a fat ridge. Note also the gaps in the comparagram owing to the poor quality of the scanning process. (b) Even the comparagram of the images prior to the deliberately poor scan of them is itself quite spread out, indicating the images were quite noisy to begin with. (c) Comparametric regression is used to solve for the parameters of the comparametric function. The resulting comparametric plot is a noise-removed version of the comparagram, e.g., provides a smoothly constrained comparametric relationship between the two differently exposed images.

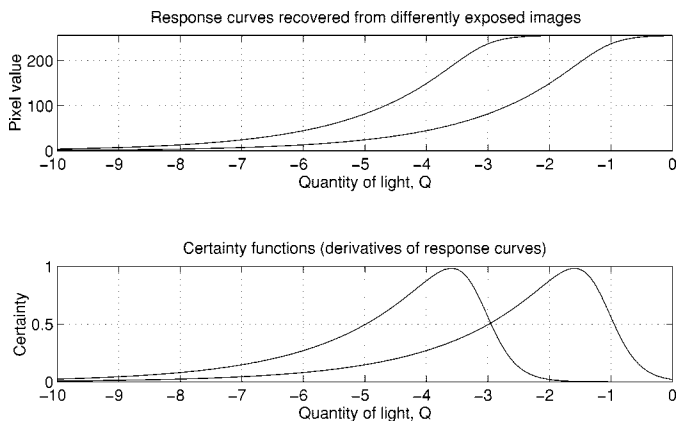


Fig. 12. Relative response functions $F^{-1}(K_i Q)$ recovered from the images in Fig. 10, plotted together with their derivatives. The derivatives of these response functions suggests a degree of confidence in the estimate $\hat{Q}_i = F(f_i) - K_i$ derived from each input image.

(given the deliberately introduced nonlinearities that best use the limited range resolution), when performing operations such as deblurring, noise artifacts become more evident. In general, sharpening involves high pass filtering, and thus sharpening will often tend to uncover noise artifacts that would normally exist below the perceptual threshold when viewed through ordinary display media. In particular, sharpening often uncovers noise in the shadow areas, making dark areas of the image appear noisy in the final print or display. Thus in addition to the benefits of performing sharpening quantigraphically by applying an anti-homomorphic filter as in Fig. 3 to undo the blur of (5), there is also further benefit from doing the generalized anti-homomorphic filtering operation at the point \hat{q} in Fig. 4, rather than just that depicted in in Fig. 3.

3) A third benefit of capturing more information than can be displayed, is that it defers the choice of which information




Fig. 13. Certainty images which will be used as weights when the weighted sum of estimates of the actual quantity of light is computed. Bright areas correspond to large degrees of certainty.



Fig. 14. Composite image made by simultaneously estimating the unknown nonlinearity of the camera as well as the true quantity of light incident on the camera's sensor array, given two input images from Fig. 10. The combined optimal estimate of \hat{q} is expressed here, in the coordinates of the lighter (rightmost) image. Although nothing has been done to appreciably enhance this image (e.g. the procedure of estimating q and then just converting it back into a picture again may seem pointless) we can note that while the image appears much like the rightmost input image, that the clipping of the highlight details has been softened somewhat. Later we will see methods of actual image enhancement done by processing \hat{q} prior to converting it back to an image again.

to get rid of. For example, a camera could be constructed in such a way that it had no exposure adjustments: neither automatic nor manual settings. Instead, the camera would be more like an array of lightmeters that would capture an array of light measurements. Decisions as to what subject matter is of importance could then be made at the time of viewing or the time of printing. Such a camera has been incorporated into eyeglasses [13], allowing the wearer to completely forget about the camera, with no need to worry about settings or adjustments. In this way the wearer can capture once-in-a-lifetime moments like a baby's first steps, and then worry about adjusting the exposure settings later. Exposure can then be adjusted in the peaceful quiet of the living room, long after the picture




Fig. 15. *Extreme example to illustrate nonmonotonic processing:* (a) An underexposed picture shows details such as the horizon and the sail of a boat, as seen through an open doorway, even though the sail is backlit with extremely bright light. (b) The picture is taken from inside an abandoned fortress with no interior lights. Light coming in from the open door is largely lost in the vastness of the dark interior, so that a much longer exposure is needed to show any detail of the inside of the fortress. (c) Sharpened (filtered) quantigraphic estimate $f(k_a S \hat{q}(A_2 \mathbf{x} + \mathbf{b}_2 / \hat{c}_2 \mathbf{x} + d_2))$ expressed in the projective coordinates of the second image in the image sequence (right hand image). A dynamic range in excess of a million to one was captured in \hat{q} , and the estimate was then quantigraphically sharpened, with filter S , resulting in a lateral inhibition effect so that the output is no longer monotonically related to the input. Notice, for example, that the sail is as dark as some shadow areas inside the fortress. Because of this filtering, a tremendous dynamic range has been captured and reduced to that of printed media.

is captured and the confusing excitement of the moment has passed. In this way exposure can be adjusted carefully in a quiet setting away from the busy and distracting action of everyday life. Since these decisions are made later, they can also be changed, as there is no need to commit to one particular exposure setting. Moreover, deferring exposure decisions may have forensic value. For ex-

ample, ordinary everyday subject matter and scenes may later become crime scenes, such that previously taken pictures in those spaces may help solve a crime. A family photographing their child's first trip to a grocery store may inadvertently capture an image of a fire exit illegally chained shut in the background. A fatal fire at some time in the future might call for evidence against the owner of the shop, where deferred choices of exposure may assist in the production of a picture exposed optimally for the fire exit in the background rather than the child in the foreground. Since the wearable apparatus transmits images to the World Wide Web, various viewers can each adjust the image interactively to suit their own display and perceptual capabilities, as well as their own preferences.

- 4) A fourth benefit from capturing a true and accurate measurement of the photoquantity, even if all that is desired is a nice looking picture (e.g. even if what is desired is not necessarily a true or accurate depiction of reality), is that additional processing may be done to produce a picture in which the limited dynamic range of the display or print medium shows a much greater dynamic range of input signal, through the use of further image processing on the photoquantity prior to display or printing.

It is this fourth benefit that will be further described, as well as illustrated through a very compelling example, in this section.

Ordinarily, humans can not directly perceive the "signal" we process numerically, but, rather, we perceive the effects of the "signal" on perceptible media such as television screens or the like. In particular, in order to display $\hat{q}(x, y)$, it is typically converted into an image $f(\hat{q}(x, y))$ and displayed, for example, on a television screen.

Fig. 14 is an attempt to display, on the printed page, a signal which contains much greater dynamic range than can be directly represented on the page. To do this, the estimate \hat{q} was converted into an image by evaluating $\hat{f}(\hat{k}_2\hat{q})$. Even though we see some slight benefit, over f_2 (one of the input images) the benefit has not been made fully visible in this print.

A. An Extreme Example with Spatiotonal Processing of Photoquantities

To fully appreciate the benefits of quantigraphic image processing, let us consider a seemingly impossible scene to photograph reasonably (in a natural way without bringing in lighting equipment of any kind).

Fig. 15 depicts a scene in which there is a dynamic range in excess of a million to one. In this case, two pictures were captured with several orders of magnitude difference between the two exposures. Thus the quantigraphic estimate \hat{q} has far greater dynamic range than can be directly viewed on a television or on the printed page. Display of $\hat{f}(\hat{k}_1\hat{q})$ would fail to show the shadow details, while display of $\hat{f}(\hat{k}_2\hat{q})$ would fail to show the highlight details.

In this case, even if we use the virtual camera architecture depicted in Fig. 5, there is no single value of display exposure k_d for which a display image $f_d = \hat{f}(k_d\hat{q})$ will capture both the inside of the abandoned fortress and the details looking outside through the open doorway.

Therefore, a strong highpass (sharpening) filter, S is applied to \hat{q} , to sharpen the photoquantity \hat{q} , as well as provide lateral inhibition similar to the way in which the human eye functions. Then the filtered result, $\hat{f}(k_d S\hat{q}(\hat{A}_2\mathbf{x} + \hat{\mathbf{b}}_2/\hat{c}_2\mathbf{x} + \hat{d}_2))$, is displayed upon the printed page [Fig. 15(c)], in the projective coordinates of the second (rightmost) image, $i = 2$. Note the introduction of spatial coordinates \mathbf{A} , \mathbf{b} , \mathbf{c} , and d . These compensate for projection (e.g. if the camera moves slightly between pictures), as described in [12], [13]. In particular, the parameters of a projective coordinate transformation are typically estimated together with the nonlinear camera response function and the exposure ratio between pictures [12], [13].

As a result of the filtering operation, notice that there is no longer a monotonic relationship between input photoquantity q and output level on the printed page. Notice, for example, that the sail is as dark as some shadow areas inside the fortress. Because of this filtering, the dynamic range of the image may be reduced to that of printed media, while still revealing details of the scene. This example answers the question "why capture more dynamic range than you can display."

ACKNOWLEDGMENT

The author would like to acknowledge assistance or contributions to this project from Kodak, Digital Equipment Corporation, Compaq, Xybernaut, WaveRider, CITO, and NSERC.

REFERENCES

- [1] S. Mann, "An historical account of the 'WearComp' and 'WearCam' projects developed for 'personal imaging'," in *IEEE Int. Symp. Wearable Computing*, Cambridge, MA, Oct. 13–14, 1997.
- [2] —, "Wearable computing: A first step toward personal imaging," *IEEE Computer*, vol. 30, no. 2, pp. 25–32, Feb. 1997.
- [3] B. Horn and B. Schunk, "Determining optical flow," *Artif. Intell.*, vol. 17, pp. 185–203, 1981.
- [4] M. Irani and S. Peleg, "Improving resolution by image registration," *Comput., Vis., Graph., Image Process.*, vol. 53, pp. 231–239, May 1991.
- [5] S. Mann and R. W. Picard, "Virtual bellows: Constructing high-quality images from video," in *Proc. IEEE 1st Int. Conf. Image Processing*, Austin, TX, Nov. 13–16, 1994, pp. 363–367.
- [6] S. Mann, "Compositing multiple pictures of the same scene," in *Proc. 46th Annu. IS&T Conf.*, Cambridge, MA, May 9–14, 1993, pp. 50–52.
- [7] S. Mann and R. W. Picard, Video orbits of the projective group; a simple approach to featureless estimation of parameters, in Tech. Rep. 338, Mass. Inst. Technol., Cambridge, MA, 1995.
- [8] R. Szeliski, "Video mosaics for virtual environments," *Comput. Graph. Applicat.*, pp. 22–30, Mar. 1996.
- [9] C. W. Wyckoff, "An experimental extended response film," *SPIE Newslett.*, pp. 16–20, June–July 1962.
- [10] —, "An experimental extended response film," Germeshausen & Grier, Inc., Boston, MA, Tech. Rep. no. B-321, Mar. 1961.
- [11] S. Mann and R. W. Picard, "Being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures," MIT Media Lab Perceptual Computing Section, Cambridge, MA, Tech. Rep. 323, 1994.
- [12] S. Mann, "'pencigraphy' with AGC: Joint parameter estimation in both domain and range of functions in same orbit of the projective-Wyckoff group," MIT Media Lab., Cambridge, MA, <http://hi.eecg.toronto.edu/icip96/index.html>, Dec. 1994.
- [13] —, "Humanistic intelligence/humanistic computing: 'wearcomp' as a new framework for intelligent signal processing," *Proc. IEEE*, vol. 86, no. 11, pp. 2123–2151, Nov. 1998.
- [14] C. Poynton, *A Technical Introduction to Digital Video*. New York: Wiley, 1996.
- [15] [Online], <http://wearcam.org/gamma.htm>.
- [16] T. G. Stockham, Jr., "Image processing in the context of a visual model," *Proc. IEEE*, vol. 60, pp. 828–842, July 1972.

- [17] S. Mann, "Personal imaging and lookpainting as tools for personal documentary and investigative photojournalism," *ACM Mobile Network*, vol. 4, no. 1, pp. 23–36, 1999.
- [18] J. Aczél, *Lectures on Functional Equations and Their Applications*. New York: Academic, 1966, vol. 19.
- [19] B. K. P. Horn, *Robot Vision*. New York: McGraw-Hill, 1986, ch. 1, 2, 6, 9, 10, 12.
- [20] B. K. P. Horn and E. J. Weldon, Jr., "Direct methods for recovering motion," *Int. J. Comput. Vis.*, 1988.
- [21] O. D. Faugeras and F. Lustman, "Motion and structure from motion in a piecewise planar environment," *Int. J. Pattern Recognit. Artif. Intell.*, vol. 2, no. 3, pp. 485–508, 1988.
- [22] S. Laveau and O. Faugeras, "3-D scene representation as a collection of images," in *Proc. IEEE Conf. Computer Vision Pattern Recognition*, Seattle, WA, June 1994.
- [23] S. C. Woon, "Period-harmonic-tupling jumps to chaos and fractal-scaling in a class of series," *Chaos, Solitons, Fractals*, vol. 5, pp. 125–130, Jan. 1995.
- [24] M. A. Berger, "Random affine iterated function systems: Curve generation and wavelets," *SIAM Rev.*, vol. 31, no. 4, pp. 614–627, 1989.
- [25] S. Mann, Wearable, tetherless computer-mediated reality: WearCam as a wearable face-recognizer, and other applications for the disabled, in TR 361, MIT Media Lab Perceptual Computing Section.
- [26] S. Mann and R. W. Picard, *IEEE Trans. Image Processing*, Sept. 1997.
- [27] S. Mann and R. W. Picard, *IS&T's 48th Annu. Conf.*, Washington, D.C., May 7–11, 1995, <http://wearcam.org/ist95.htm>, pp. 422–428.
- [28] S. Mann, *Proc. IEEE Int. Conf. Image Processing*, Lausanne, Switzerland, Sept. 16–19, 1996, pp. 193–196.
- [29] S. Mann, *AAAI Fall Symp. Developing Assistive Technology People Disabilities*, Nov. 9–11, 1996.

- [30] S. Mann, [Online], Cambridge, MA, Feb. 2, 1996, <http://wearcam.org/vmp.htm>.



Steve Mann pictured here wearing an early embodiment of his so-called "wearable computer" (WearComp) and EyeTap video camera and reality mediator (WearCam) invention, received the Ph.D. degree from the Massachusetts Institute of Technology (MIT), Cambridge, in 1997.

He is currently a faculty member with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, Ont., Canada. He has been working on his WearComp invention for more than 20 years. In 1991, he initiated what was to later become the MIT Wearable Computing Project. He also built the world's first covert fully functional WearComp with display and camera concealed in ordinary eyeglasses in 1995, for the creation of his documentary video "ShootingBack." He is also the inventor of the chirplet transform, a new mathematical framework for signal processing. He chaired the first Special Issue on Wearable Computing in *Personal Technologies Journal* and has given numerous keynote addresses on the subject, including the keynote at the first International Conference on Wearable Computing, the Virtual Reality Conference, and the McLuhan Conference on Culture and Technology, on the subject of privacy issues and wearable computers.

Dr. Mann both proposed and was the Publications Chair of the first IEEE International Symposium on Wearable Computing (ISWC'97).