

Image Quality Assessment Based on a Degradation Model

Niranjan Damera-Venkata, *Student Member, IEEE*, Thomas D. Kite, Wilson S. Geisler, Brian L. Evans, *Senior Member, IEEE*, and Alan C. Bovik, *Fellow, IEEE*

Abstract—We model a degraded image as an original image that has been subject to linear frequency distortion and additive noise injection. Since the psychovisual effects of frequency distortion and noise injection are independent, we decouple these two sources of degradation and measure their effect on the human visual system. We develop a distortion measure (DM) of the effect of frequency distortion, and a noise quality measure (NQM) of the effect of additive noise. The NQM, which is based on Peli's contrast pyramid, takes into account the following:

- 1) variation in contrast sensitivity with distance, image dimensions, and spatial frequency;
- 2) variation in the local luminance mean;
- 3) contrast interaction between spatial frequencies;
- 4) contrast masking effects.

For additive noise, we demonstrate that the nonlinear NQM is a better measure of visual quality than peak signal-to-noise ratio (PSNR) and linear quality measures. We compute the DM in three steps. First, we find the frequency distortion in the degraded image. Second, we compute the deviation of this frequency distortion from an allpass response of unity gain (no distortion). Finally, we weight the deviation by a model of the frequency response of the human visual system and integrate over the visible frequencies. We demonstrate how to decouple distortion and additive noise degradation in a practical image restoration system.

Index Terms—Computational vision, human visual system modeling, image quality.

I. INTRODUCTION

IMAGES may be corrupted by degradation such as linear frequency distortion, noise, and blocking artifacts. These sources of degradation may arise during image capture or processing, and have a direct bearing on visual quality. In this paper, we model degradation to develop efficient methods for minimizing the visual impact of degradation. We model a degraded image as an original image which has been subject to two independent sources of degradation—linear frequency distortion and additive noise injection. This model is commonly

used in image restoration. Based on the model, we develop methods to measure the quality of images and demonstrate how one may use the quality measures in quantifying the performance of image restoration algorithms.

We model the distortion (relative to the original image) as linear and spatially invariant. We model the noise as spatially varying additive noise. We refer to a *degraded image* as an image degraded by the two-source degradation model. When we speak of the quality of the restored image, we consider the degraded image to be the image we are processing with the restoration algorithm. We will then quantify the degradation *in the restored image* as compared with the original, uncorrupted image.

We develop two complementary quality measures that separately measure the impact of frequency distortion and noise injection on the human visual system (HVS). This decoupled approach allows a designer to explore the fundamental trade-offs between distortion and noise to improve restoration algorithms, which is not possible with a scalar-valued quality measure. Previous scalar-valued image quality measures have been based on signal-to-noise ratio (SNR) as well as linear and nonlinear models of the HVS.

SNR measures, such as peak SNR (PSNR), assume that distortion is only caused by additive signal-independent noise. As a consequence, noise measures applied directly to a restored image and its original do not measure visual quality. Quality measures based on linear HVS models [1]–[4] assess image quality in three steps. First, an error image is computed as the difference between the original image and the restored image. Second, the error image is weighted by a frequency response of the HVS given by a lowpass contrast sensitivity function (CSF). Finally, a signal-to-noise ratio is computed. These quality measures can take into account the effects of image dimensions, viewing distance, printing resolution, and ambient illumination. They do not include nonlinear effects of contrast perception, such as local luminance, contrast masking, and texture masking [5]–[7].

Daly's visible differences predictor [5] assesses still image quality using a nonlinear HVS model consisting of an amplitude nonlinearity, a lowpass CSF, and a hierarchy of detectors. Daly's predictor produces an *error image* which characterizes the regions in the test image that are visually different from the original image. The degree of visual difference at each point is quantified by the intensity at that point. The results of the Daly model need to be interpreted by visual inspection of the error image. Daly's model is well suited for compression. Lubin's sarnoff visual discrimination model [6], which is also based on a nonlinear HVS model, quantifies a wide variety of distortions, including blocking and quantization effects which are common

Manuscript received October 7, 1998; revised August 13, 1999. This work was supported by Hewlett-Packard and a U.S. National Science Foundation CAREER Award Grant MIP-9702707. The authors conducted this research at the Center for Vision and Image Sciences, University of Texas. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Glenn Healey.

N. Damera-Venkata, B. L. Evans, and A. C. Bovik are with the Department of Electrical and Computer Engineering, University of Texas, Austin, TX 78712 USA (e-mail: damera@vision.ece.utexas.edu; bevans@vision.ece.utexas.edu; bovik@vision.ece.utexas.edu).

W. S. Geisler is with the Department of Psychology, University of Texas, Austin, TX 78712 USA (e-mail: geisler@vision.ece.utexas.edu).

T. D. Kite is with Audio Precision, Beaverton, OR 97075 USA (e-mail: tom@vision.ece.utexas.edu).

Publisher Item Identifier S 1057-7149(00)02677-4.

in image compression. Teo and Heeger’s perceptual distortion metric [8] is similar in spirit to Lubin’s model. These computationally intensive approaches return either a single parameter or an error map to represent visual quality.

We develop two measures of degradation—distortion measure (DM) and noise quality measure (NQM)—based on the observation that the psychovisual effects of filtering and noise are separate. Instead of computing a residual image, we compute a *model restored image* by passing the original image through the restoration algorithm using the same parameters as were used while restoring a degraded image. We compute the DM in three steps. First, we find the frequency distortion in the restored image by comparing the restored and the model restored images. Second, we compute the deviation of this frequency distortion from an allpass response of unity gain (no distortion). Finally, we weight the deviation by a lowpass CSF and integrate over the visible frequencies.

We compute the NQM in two steps. First, we process the original image and the modeled restored image separately through a contrast pyramid. The contrast pyramid, which is based on Peli’s work [9], computes the contrast in an image at every pixel and at spatial frequencies separated by an octave, and models the following nonlinear spatially varying visual effects:

- 1) variation in contrast sensitivity with distance, image dimensions, and spatial frequency;
- 2) variation in the local luminance mean;
- 3) contrast interaction between spatial frequencies;
- 4) contrast masking effects.

Second, we form the NQM by computing the SNR of the restored degraded image with respect to the model restored image.

The NQM is similar to Lubin’s model, but exhibits several key differences. The NQM ignores the orientation sensitivity of the HVS. Based on visual tests, Mitsa and Varkur [3] conclude that ignoring orientation sensitivity, i.e., assuming a uniform retina, has very little effect on visual quality. This agrees with Peli [9]. By omitting orientation sensitivity, we greatly reduce computational cost by avoiding directional filtering, skew Hilbert transforms, and model calibration and contrast normalization. Moreover, we use a cosine-log filterbank instead of the Gaussian pyramid in implementing the contrast pyramid. This approach is justified in Section VI. Contrast masking is taken directly into account by using the contrast pyramid.

Section II reviews several quality measures. Section III decouples frequency distortion from noise injection in restored images and defines a distortion transfer function for image restoration systems. Section IV develops the DM which weights the distortion transfer function to quantify the psychovisual effect of frequency distortion. Section V reviews several definitions of contrast, and describes a consistent definition by Peli [9] that underlies the NQM. Section VI defines the nonlinear NQM. Section VII illustrates the performance of the NQM using test images. Section VIII concludes the paper.

II. QUALITY MEASURES FOR DEGRADATION BY ADDITIVE NOISE

Objective measures that correlate with the visual difference between two images are key in ranking and optimizing image

restoration algorithms. Quality measures should vary monotonically with visual quality. Section II-A reviews SNR and PSNR measures. Section II-B reviews linear quality measures which weight the noise in frequency according to a model of the frequency response of the HVS.

A. Conventional Quality Measures: SNR and PSNR

Both SNR and PSNR are mean-squared (l_2 -norm) error measures. SNR is defined as the ratio of average signal power to average noise power. For an $M \times N$ image

$$\text{SNR (dB)} = 10 \log_{10} \left(\frac{\sum_{i,j} x(i,j)^2}{\sum_{i,j} (x(i,j) - y(i,j))^2} \right) \quad (1)$$

for $0 \leq i \leq M - 1$ and $0 \leq j \leq N - 1$, where $x(i, j)$ denotes pixel (i, j) of the original (“clean”) image and $y(i, j)$ denotes pixel (i, j) of the noisy image. PSNR is defined as the ratio of peak signal power to average noise power

$$\text{PSNR (dB)} = 10 \log_{10} \left(\frac{D^2 MN}{\sum_{i,j} (x(i,j) - y(i,j))^2} \right) \quad (2)$$

for $0 \leq i \leq M - 1$ and $0 \leq j \leq N - 1$ where D is the maximum peak-to-peak swing of the signal (255 for 8-bit images). We assume that the noise $x(i, j) - y(i, j)$ is uncorrelated with the signal.

In many image processing applications, such as compression and halftoning, degradation in the processed image is not due solely to additive uncorrelated noise. As a consequence, the correlation between SNR or PSNR and visual quality is known to be poor [10]. Fig. 2(a) shows the *lena* image corrupted by additive white noise, and Fig. 2(b) shows the *lena* image corrupted with additive highpass noise (“blue noise” [11]), which is characteristic of halftoning by error diffusion. Even though both images have an SNR of 10.00 dB, Fig. 2(b) has higher subjective quality. Noise measures are commonly misused in standards and the literature to evaluate image quality when the image has been corrupted by degradation other than additive noise, e.g., blocking artifacts.

B. Linear Quality Measures: Frequency-Domain Weighting

Several quality measures [2]–[4] perceptually weight the frequency domain by using the frequency response of a linear model of the HVS. The HVS, however, is a nonlinear, spatially varying system. A measure of the nonlinear HVS response to a single frequency is called the contrast threshold function (CTF). The CTF is measured over the visible radial spatial frequencies from 0 to 60 cycles/degree [12]. The CTF is the minimum amplitude necessary to just detect a sine wave of a given angular spatial frequency [13]. Inverting a CTF gives a frequency response, called the contrast sensitivity function (CSF), which is a linear spatially invariant approximation to the HVS.

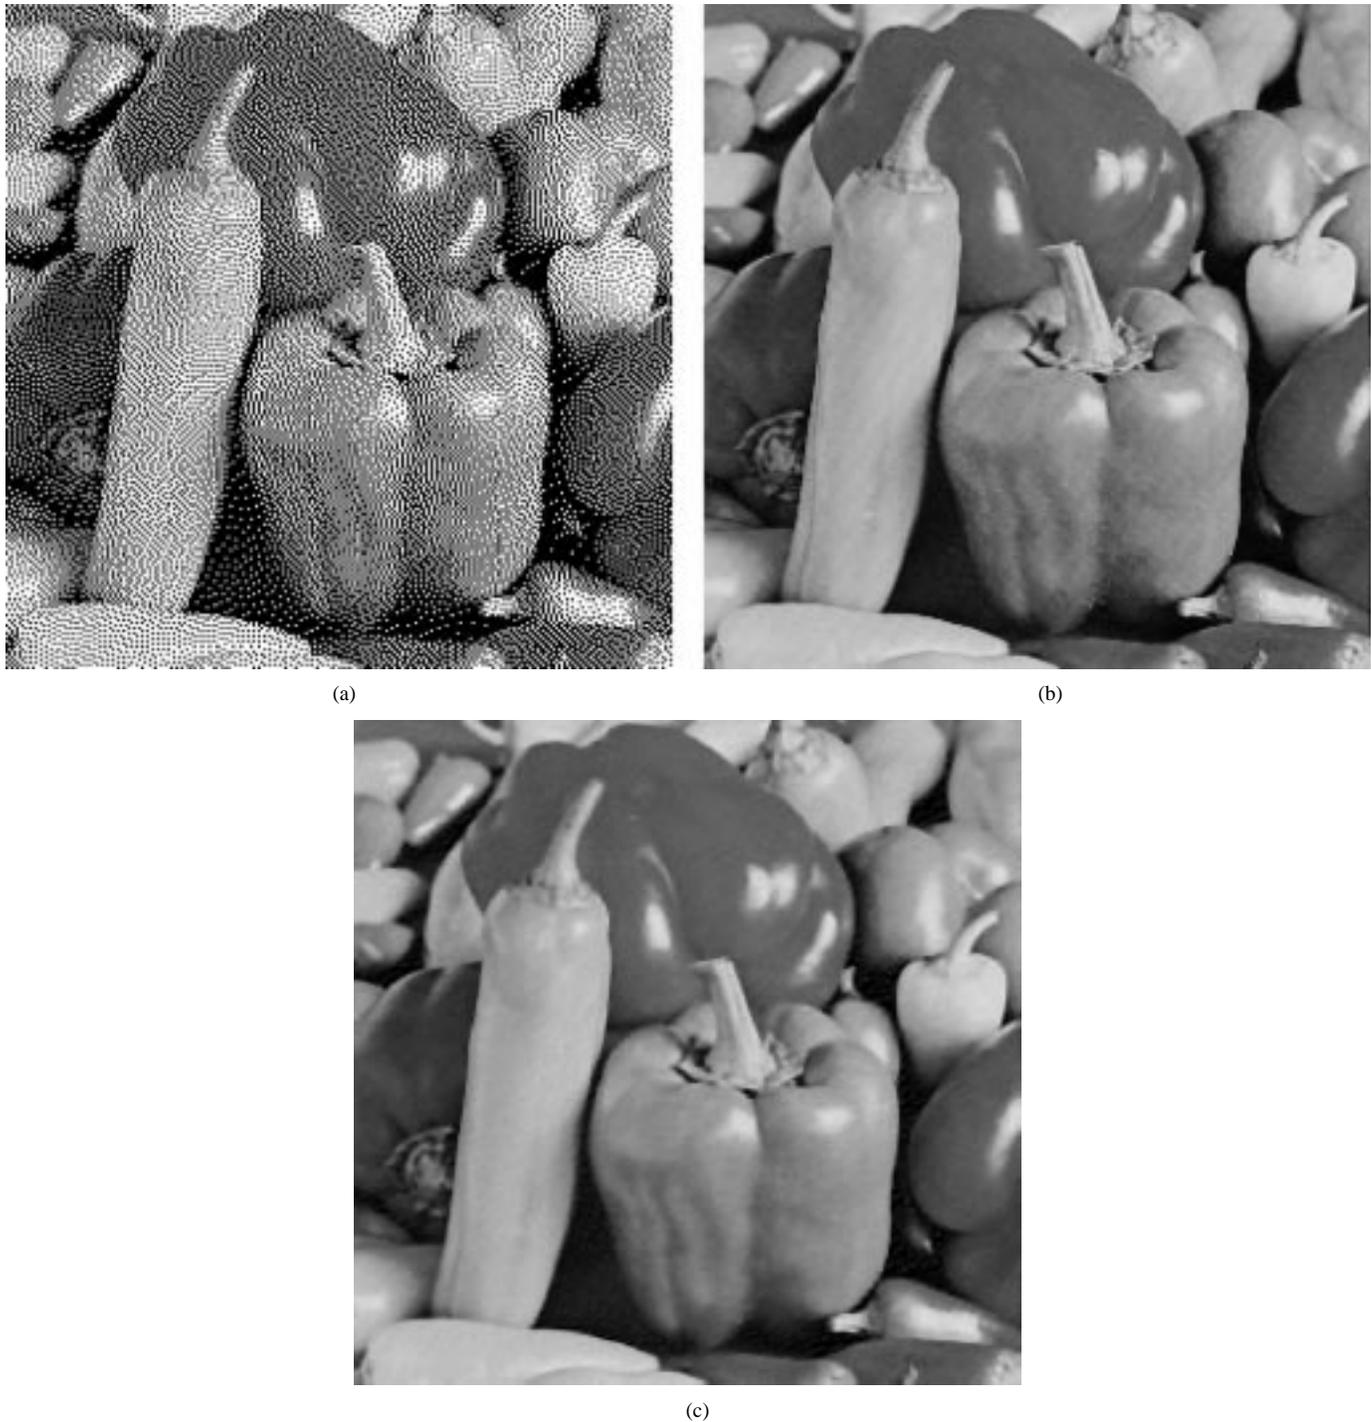


Fig. 1. Example of image restoration system using the *peppers* image as the original. (a) Degraded image (1 bit/pixel), (b) original image (8 bits/pixel), and (c) restored image (8 bits/pixel). We used the error diffusion halftoning algorithm [21] to produce the degraded image, and the inverse halftoning algorithm [20] to restore the image.

Fig. 3 shows a bandpass CSF [14], a lowpass CSF [3], [4], and a CTF. Based on psychovisual tests, the lowpass CSF model is better for complex images when viewed under suprathreshold conditions [3], [4]. The bandpass model is derived from experiments with the subject fixated; under normal conditions, eye movements restore the lost low frequency sensitivity [15]. Peli *et al.* [16] provide an excellent discussion of the measurement and choice of CSF for use in practical applications.

The CSF can incorporate information about the printing device and viewing conditions in quality measures. Lin [2] uses the lowpass CSF to weight the Fourier transforms of the original image and the degraded image, and then computes a root mean square error in the frequency domain between the two weighted images. Mitsa [3] models the processing of cortical simple cells in the eye as a bank of Gabor bandpass filters. The error image is decomposed in the filterbank and each bandpass filter output is weighted according to the lowpass CSF.



(a)

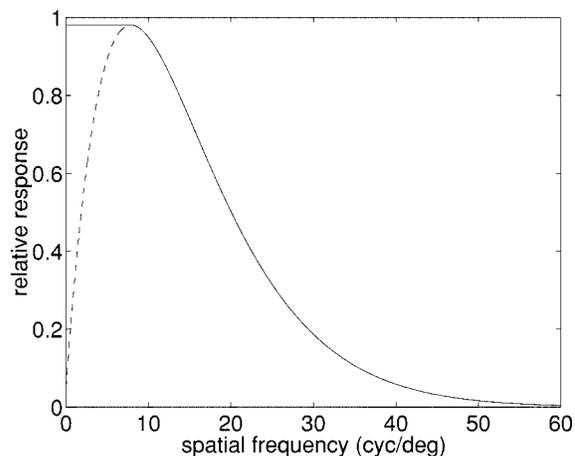


(b)

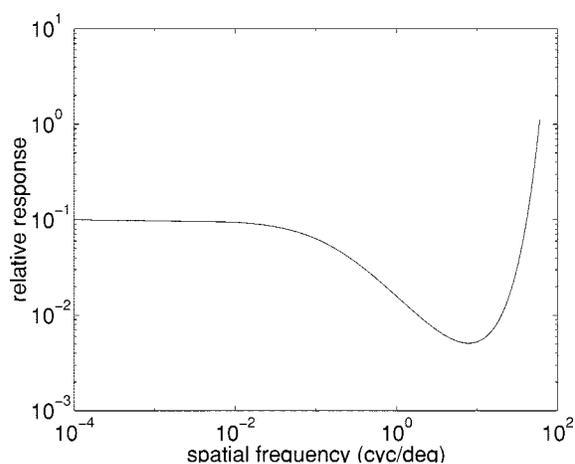
Fig. 2. Two corrupted *Lena* images with the same SNR with respect to the original but with different visual quality. (a) White noise added and (b) filtered white noise added.

By using the CSF as the weighting function, we define weighted SNR (WSNR) as the ratio of the average *weighted* signal power to the average *weighted* noise power. The images in Fig. 2(a) and (b) have WSNR values of 11.22 dB and 28.67 dB, respectively, when viewed at a 4° visual angle. This ordering corresponds to their relative visual quality.

Because a CSF is a linear spatially invariant approximation of the HVS, it cannot quantify nonlinear and spatially varying effects. It cannot model the change in perceived contrast due to amplitude components at other spatial frequencies [9], [17]. It also ignores the change in perceived noise level with local image content. The visibility of a pixel depends on the local background contrast. This effect, called contrast masking, is ignored by the CSF. Therefore, before applying any noise measure such as SNR, PSNR, or WSNR, it is crucial to simulate the nonlinear, spatially varying response of the HVS to the original image and the processed image.



(a)



(b)

Fig. 3. HVS response to a sine wave at different frequencies. In (a), the bold line denotes the lowpass modification to the CSF to account for suprathreshold viewing and the dotted line shows the original bandpass CSF. (a) Contrast sensitivity functions and (b) contrast threshold function.

III. DECOUPLING FREQUENCY DISTORTIONS AND NOISE DEGRADATION

Before applying a noise measure such as SNR, PSNR, or WSNR, it is necessary to account for the sources of degradation other than additive noise [18]. Otherwise, the other sources of degradation will be erroneously incorporated into the noise measure, as demonstrated by Fig. 4. Fig. 4(a) is the original *lena* image. Fig. 4(b) sharpens the original image with a 3×3 filter. We add highpass noise to Fig. 4(b) to produce Fig. 4(c). The SNR of Fig. 4(c) relative to Fig. 4(b) is 10 dB. Fig. 4(d) shows the difference (residual) between Fig. 4(a) and (c). Because the residual is correlated with the original image, it is inappropriate to compute an SNR measure of Fig. 4(c) relative to Fig. 4(a). It is appropriate to compute an SNR measure for Fig. 4(c) relative to Fig. 4(b), since their difference is noise that is independent of the original image.

Table I lists WSNR figures for the image in Fig. 4(c) for five viewing distances. The third column shows the WSNR relative to Fig. 4(a), while the fourth column shows the WSNR relative to Fig. 4(b). As expected, the values in the



Fig. 4. Effect of sharpening on SNR measures. (a) Original image, (b) sharpened original, (c) sharpened original + highpass noise, and (d) residual of (c) and (a). Since the residual of (c) and (a) shown in (d) contains information from (a), applying an SNR measure of (c) relative to (a) would be inappropriate. Since the residual (c) and (b) consists of signal-independent noise, applying an SNR measure of (c) relative to (b) would be appropriate.

third column are lower than those in the fourth column, because the residual includes power from the original image. The WSNR figures relative to the sharpened original are correct because the residual is uncorrelated with the original image. The results of Table I show the importance of removing as much image power as possible from the residual before computing the WSNR of an image.

In this section, we separate sources of degradation in restored images into *noise injection* and *frequency distortion*. This de-

coupling enables both effects to be quantified and restoration algorithms to be assessed. Section III-A defines a correlation measure between images which we use to quantify the amount of signal components present in noise. Section III-B derives an effective transfer function for a restoration system called a distortion transfer function (DTF). Section III-C gives an example of computing the DTF for a practical image restoration system. The DTF is the basis for the distortion quality measure described in Section IV.

TABLE I
WSNR FIGURES USING INCORRECT AND
CORRECT RESIDUALS

Viewing Distance (mm)	Max. Angular Frequency (cyc/deg)	WSNR w/r to	
		original (dB)	sharpened (dB)
200	6.9	5.8	10.2
400	13.7	8.8	13.1
600	20.6	12.5	17.3
800	27.5	15.7	21.5
1000	34.4	18.1	25.6

TABLE II
VARIATION OF SNR AND WSNR WITH CORRELATION OF RESIDUAL

Gain α	C_{RI}	SNR (dB)	WSNR (dB)
1.000	0.000	28.0	31.7
1.005	0.019	27.9	31.6
1.010	0.038	27.7	31.1
1.015	0.058	27.4	30.5
1.020	0.077	27.0	29.7
1.025	0.096	26.5	28.9
1.030	0.115	26.0	28.0

A. Correlation of the Residual with the Original Image

To quantify the degree to which a residual image R is correlated with an original image I , we use the magnitude of the correlation coefficient between them [19]

$$C_{RI} = \frac{|\text{Cov}[R, I]|}{\sigma_R \sigma_I} \quad (3)$$

where Cov refers to covariance, and σ_R and σ_I are the standard deviations of images R and I , respectively. By using an absolute value in the numerator, we ensure that $0 \leq C_{RI} \leq 1$, with 0 indicating no correlation and 1 indicating linear correlation. Thus, C_{RI} can be considered to be a measure of linear correlation between two images. The covariance is defined as

$$\text{Cov}[R, I] = E[(R - \mu_R)(I - \mu_I)] \quad (4)$$

where μ_R and μ_I denote the means of R and I , respectively.

A residual image should consist only of independent additive noise, and should therefore have zero correlation with the original. In practice, the correlation will not be exactly zero; this may cause signal-to-noise ratio measures to be in error. We analyze the effect of correlation on WSNR. We generate an “original image” I , composed of lowpass filtered noise, and a white noise image N of the same size. We create a noisy, corrupted image

$$J = \alpha I + N \quad (5)$$

where α is a gain factor. The residual image is $R = J - I = (\alpha - 1)I + N$. We force a prescribed linear correlation between R and I by choosing α , measure the correlation, and compute SNR and WSNR for J relative to I .

Table II shows the results for values of α ranging from 1.000 to 1.030. As α increases, the correlation C_{RI} increases, and the SNR and WSNR decrease, as expected. The WSNR falls by approximately 3 dB as the correlation increases from zero to 0.100. This large variation underscores the importance of the correlation of the residual and the original image being approximately zero for the WSNR figure to be accurate. We consider $C_{RI} < 0.020$ to be approximately zero.

B. The Distortion Transfer Function

We model the blurring in restoration algorithms to create a noise-free model restored image that exhibits the same blurring as the restored image. We can then obtain a residual between

the restored image and the model restored image that is additive noise. We model the blur by computing an effective transfer function for the image restoration system as follows:

- compute the two-dimensional (2-D) fast Fourier transform (FFT) of the original image and the model restored image;
- divide the model FFT by the original image FFT point-for-point, for spatial frequencies where the original image FFT is nonzero. Where the FFT of the original image is zero, the corresponding frequencies in the computed transfer function are set to unity;
- compute the absolute value (magnitude) of the complex quotient to find the 2-D transfer function; and
- radially average the transfer function over annuli of radius f_r [11] to obtain a one-dimensional (1-D) distortion transfer function (DTF).

The resulting 1-D DTF reflects the blurring in the restored image.

C. Computing a Model Restored Image in a Practical System

We address the issue of computing the model restored image for a practical image restoration method. The model restored image has similar linear distortion characteristics to the restored image, but it is noise-free. We first process the degraded image using the restoration scheme. This results in an image with both linear distortion and additive noise. The parameters used are saved, and the *original image* is processed with the saved parameters to produce the model restored image.

We illustrate this approach with an example. We consider an algorithm which attempts to restore a degraded image using spatially adaptive linear filters [20]. This algorithm actually performs an operation known as inverse halftoning, in which a 1 bit/pixel quantized image is to be restored to an 8 bits/pixel grayscale image. We use the degradation model to model the restoration algorithm. We consider the restored image, and attempt to quantify its frequency distortion with respect to the original.

We compute the DTF for the restoration algorithm in [20]. The algorithm adaptively smooths quantization noise and preserves edge information by using a 7×7 spatially varying FIR lowpass filter. We assess the frequency distortion of the algorithm in two steps. First, we save the filter used at each pixel, while restoring a 1 bit/pixel image. We must be confident that

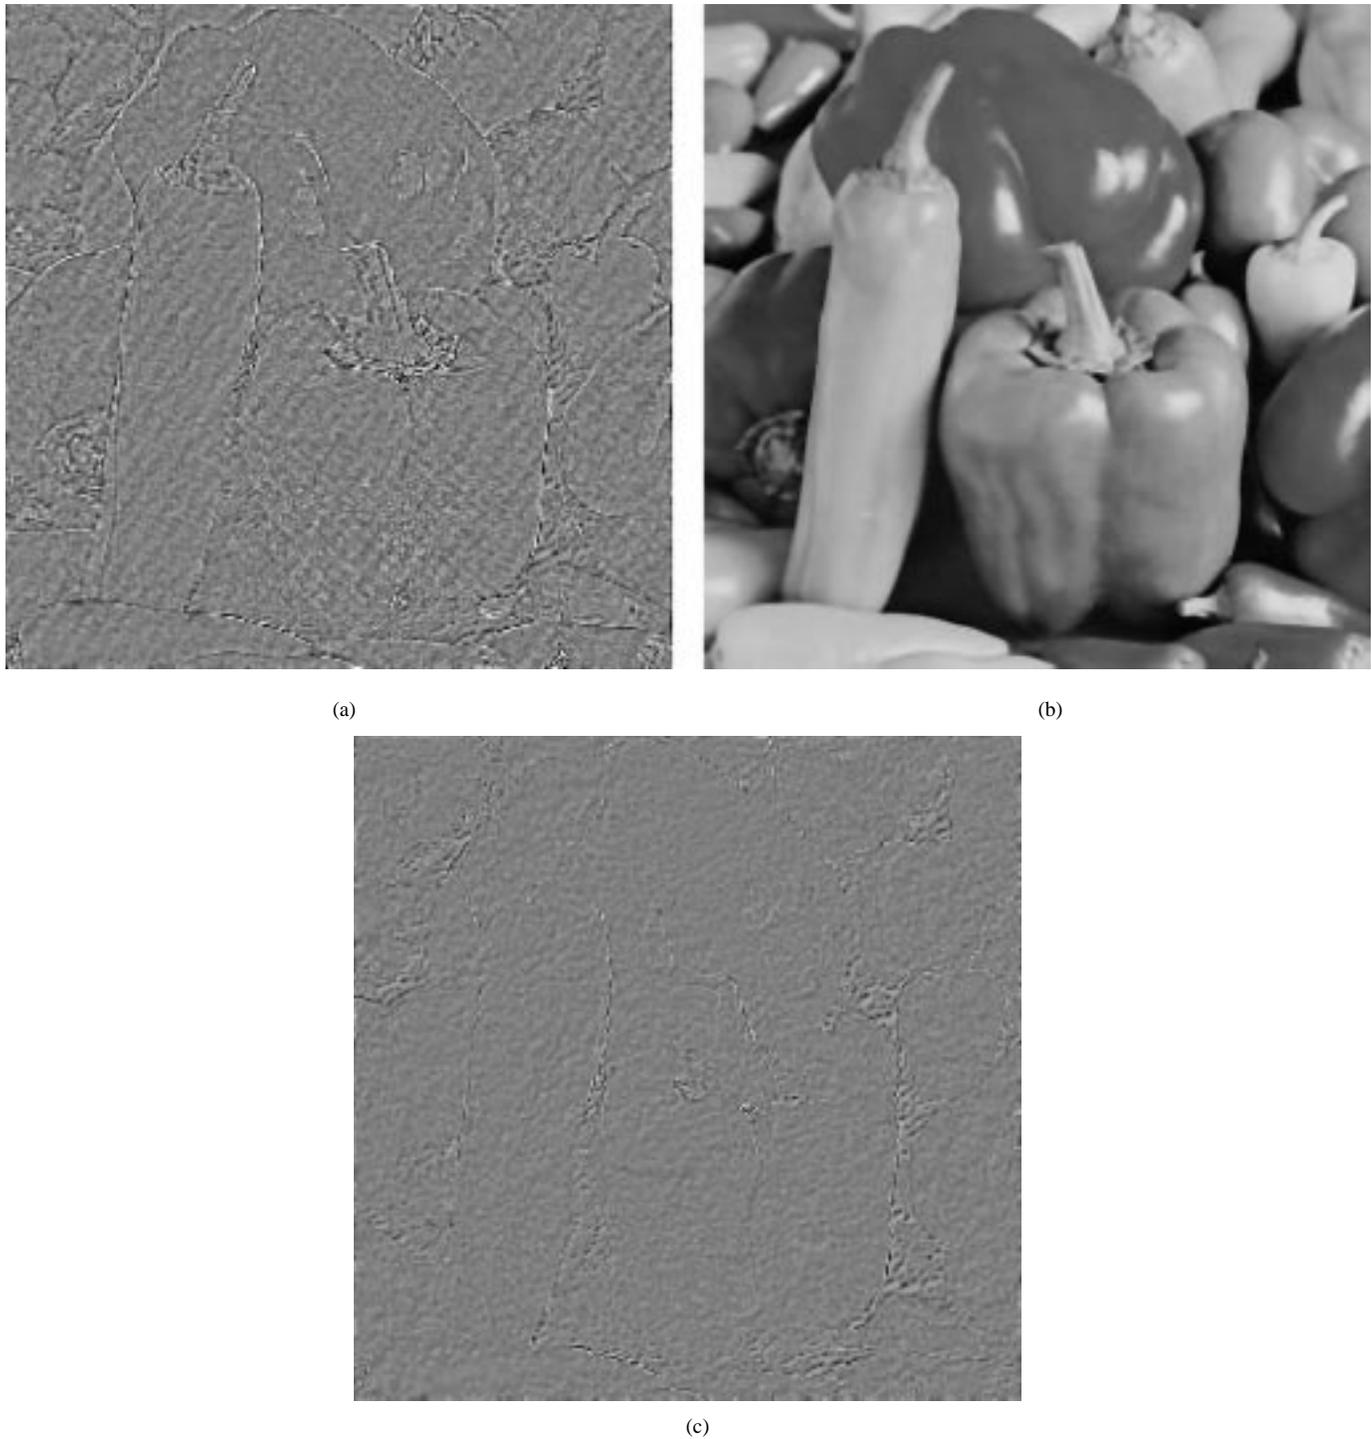


Fig. 5. Result of modeling a restored peppers image. (a) Residual of the restored image [Fig. 1(c)] minus the original image [Fig. 1(b)]. (b) Model restored image having the same linear frequency distortion as the degraded image but without noise. (c) Residual of the restored image [Fig. 1(c)] minus the model restored image [Fig. 5(b)]. The residual in (c) represents noise that is uncorrelated with respect to the model restored image. In all three images, a gain of four was applied for display purposes.

the degraded image has similar sharpness as the original image [18]. It may be necessary to preprocess the original with a linear filter to achieve this. Second, we apply the saved filters to the preprocessed original to produce the model restored image that has the same spatial blur as the restored image, but does not include the injected noise (quantization noise, in this case).

Fig. 1(a) shows the Floyd–Steinberg halftone [21] of the original *peppers* image in Fig. 1(b) which we are trying to

restore. We compute the restored image, as shown in Fig. 1(c), and save the FIR filter coefficients used at each pixel. Fig. 5(a) shows the residual between the restored image and the original image. Strong image edges exist because the restored image is blurred. Fig. 5(b) shows the model restored image, computed from Fig. 1(b) using the same filters used to create Fig. 1(c). Fig. 1(c) shows the residual between the restored image and the model restored image. The image components are greatly

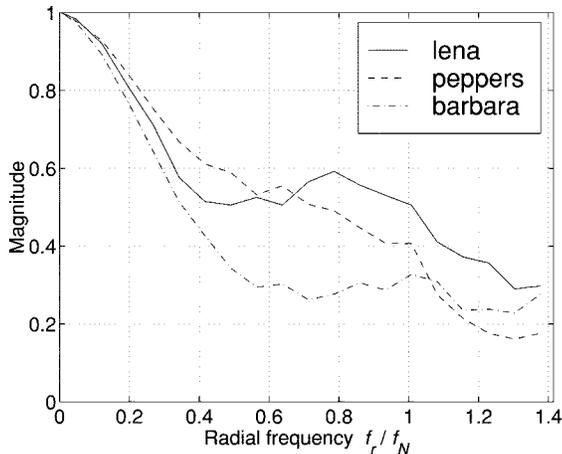


Fig. 6. Distortion transfer function of image restoration systems is a function of radial frequency $f_r = \sqrt{f_x^2 + f_y^2}$ [20]. The magnitude of f_r is the average transfer function magnitude over an annulus in the frequency domain with the average radius f_r .

reduced relative to Fig. 5(a). Fig. 6 shows the DTF's for the *lena*, *peppers*, and *barbara* images. All show the marked high frequency suppression that is characteristic of blurring.

We validate the model for frequency distortion by using the correlation measure given by (3). Table III shows the correlation between the original image and two residual images: the difference between the restored image and the original image, and the difference between the restored and the model restored image. For the test images, the average correlation is 0.317 for the actual residual and 0.010 for the modeled residual. On average, image components have been suppressed by factor of 33 in the modeled residual.

The low correlation of the original image and the modeled residual permits the use of modeled degraded images as a basis for WSNR measurements. Table IV shows WSNR measurements for five test images, assuming a maximum spatial frequency in the x and y directions of 20 cycles/degree, which corresponds to a typical combination of image resolution, size, and viewing distance. The first row shows the WSNR of the restored image relative to the original image, while the second row shows the WSNR of the restored image relative to the model restored image. The second of these figures is a true measure of the weighted noise content of the restored images, since the first figure includes image distortions. As expected, WSNR is higher when the restored image is compared to the model restored image. It is also more stable across images, varying by 1.25 dB over the test set, compared to a variation of over 8.50 dB when image distortion is not taken into account.

IV. DISTORTION MEASURE

The previous section demonstrates the importance of decoupling the frequency distortion and noise introduced by image restoration algorithms. In this section, we quantify the psycho-visual effect of frequency distortion by using the lowpass CSF discussed in Section II and the DTF introduced in Section III-B. First, we weight the deviation from unity of the DTF by the CSF

TABLE III
CORRELATION COEFFICIENTS FOR INVERSE HALFTONE RESIDUALS

Residual	Correlation Coefficient $C_{\text{original,difference}}$					
	<i>w/r to</i>	<i>barbara</i>	<i>boats</i>	<i>lena</i>	<i>mandrill</i>	<i>peppers</i>
Original		0.364	0.259	0.189	0.546	0.229
Model		0.007	0.006	0.011	0.018	0.007

TABLE IV
WSNR MEASURES FOR INVERSE HALFTONES

Ref.	WSNR (dB)				
	<i>barbara</i>	<i>boats</i>	<i>lena</i>	<i>mandrill</i>	<i>peppers</i>
Image					
Original	20.47	25.36	26.93	19.02	27.69
Model	32.29	33.02	32.74	31.93	31.77

to form a visual distortion function. Then, we compute the area under the visual distortion function to form a distortion measure

$$DM = \int_0^{f_{\max}} \left[1 - \text{DTF} \left(\frac{f_r}{f_N} \right) \right] \text{CSF}(f_r) df_r. \quad (6)$$

Here, f_r is the radial frequency $f_r = \sqrt{f_x^2 + f_y^2}$, where f_x and f_y are the horizontal and vertical frequencies, respectively; f_N is the Nyquist frequency; and f_{\max} is the maximum radial frequency included in the DM. We choose f_{\max} to be 60 cycles/degree. The DM penalizes low frequency distortion more heavily than high frequency distortion, to model perception by the HVS. Using $DM = 1$ as an arbitrary reference, we express the DM as

$$DM \text{ (dB)} = 20 \log_{10}(DM). \quad (7)$$

On a single-processor 167 MHz Sun Ultra-2 workstation, Matlab 5 requires 10 s to compute the DM for 512×512 original and restored images.

V. CONTRAST IN COMPLEX IMAGES

Contrast is a key perceptual image attribute. Measurement and evaluation of contrast and contrast changes in arbitrary images are not uniquely defined in the literature. The processing of images in the human visual system is believed to be neither periodic nor purely local; therefore, the representation of contrast in images should be quasi-local as well. For completeness, this section reviews the development of a robust measure of contrast in complex images [9]. Section V-A summarizes measures of contrast in simple images. Section V-B reviews three measures of contrast for complex images. Section V-C presents Peli's definition of contrast and shows that this definition is consistent with subjective experiments on complex images.

A. Contrast Definitions for Simple Patterns

We review the Michelson and Weber contrast definitions for simple targets in still images. The Michelson contrast [22] of a periodic pattern, such as a sinusoidal grating, is

$$C_M = \frac{L_{\max} - L_{\min}}{L_{\max} + L_{\min}} \quad (8)$$

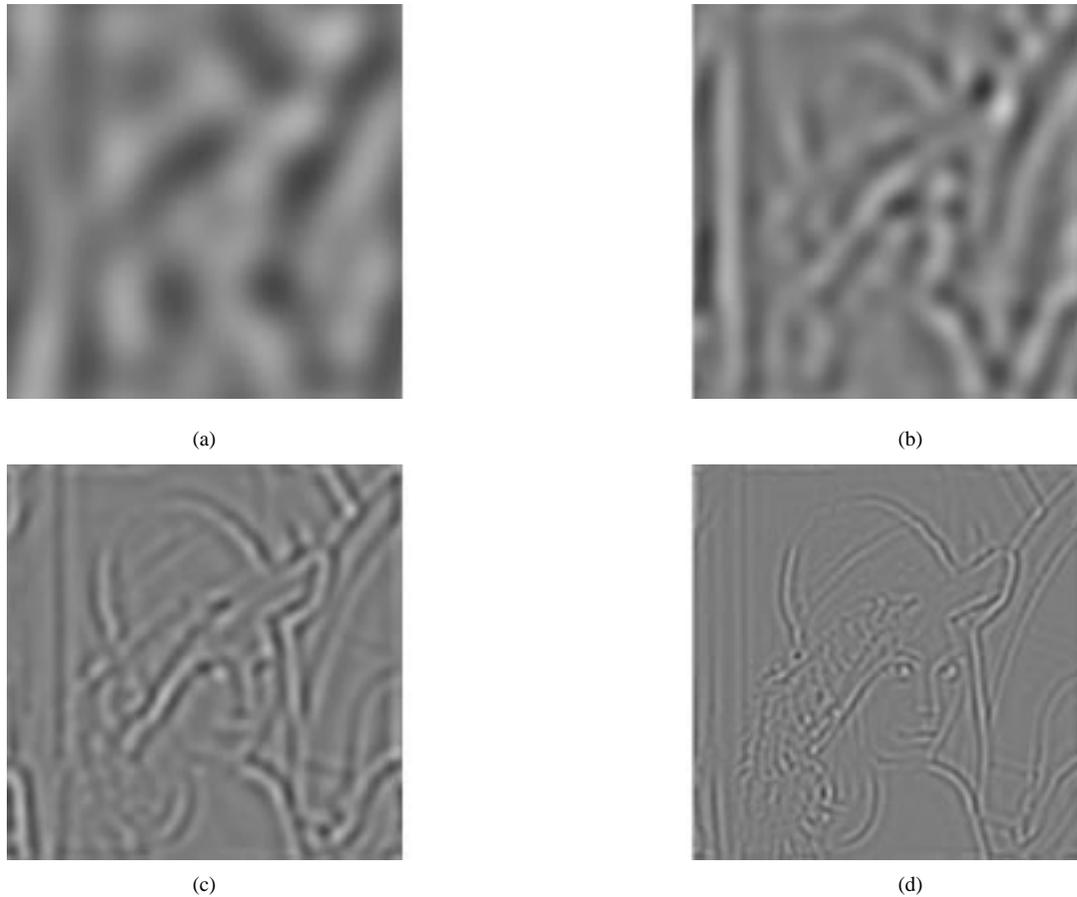


Fig. 7. Bandpass images at (a) 4, (b) 8, (c) 16, and (d) 32 cycles/image.

where L_{\max} and L_{\min} are the maximum and minimum luminance values, respectively. The Weber contrast, which measures local contrast of a single target of uniform luminance observed against a uniform background, is

$$C_W = \frac{\Delta L}{L} \quad (9)$$

where ΔL is the change in the target luminance from the uniform background luminance L . One usually assumes a large background with a small test target, so that the average luminance will be close to the background luminance. This assumption does not hold for complex images.

The Michelson contrast definition is inconsistent with the Weber contrast definition. In the Weber contrast definition, $\Delta L = 1/2(L_{\max} - L_{\min})$ and $L = L_{\min}$. Using these relations, we express the Michelson contrast as

$$C_M = \frac{\Delta L}{L + \Delta L} \quad (10)$$

to reveal that the Michelson and Weber contrast definitions disagree [9]. The numerator terms in (8) and (10) are the same but the denominator terms are only equal when $\Delta L = 0$, which is a trivial case. It is difficult to find a consistent definition of contrast for complex images.

B. Contrast Definitions for Complex Images

Many definitions of contrast in a complex scene are restricted to the assessment of contrast changes in an image displayed

in two different ways. Ginsburg [23] defines image contrast spanning all 256 gray levels as 100%; therefore, linearly compressing the image to span gray levels 0–127 reduces the contrast to 50%. With this definition, the mean luminance of the image decreases with contrast. If the minimum intensity remains zero, then Michelson's definition in (8) leaves contrast unchanged relative to compression of the graylevel range.

Hess and Pointer [24] define contrast in terms of horizontal and vertical spatial frequencies ω_1 and ω_2 as

$$C(\omega_1, \omega_2) = \frac{2A(\omega_1, \omega_2)}{A(0, 0)} \quad (11)$$

where $A(\omega_1, \omega_2)$ is the amplitude of Fourier component (ω_1, ω_2) , and $A(0, 0)$ is the DC value of the image. This definition has been applied globally to images and to nonoverlapping subimages. This approach does not capture the local nature of contrast changes.

Badcock [25] measures local contrast for complex grating patterns composed of first and third harmonics. Hess and Pointer [26] only calculate the contrast around peaks of the first harmonic and not around valleys. This implicitly ignores the effect of the local luminance mean on the contrast of the higher harmonic [9], [17], which we describe next.

C. Local Bandlimited Contrast in Complex Images

The definition of local bandlimited contrast proposed by Peli [9] provides a consistent definition of contrast. In order to de-

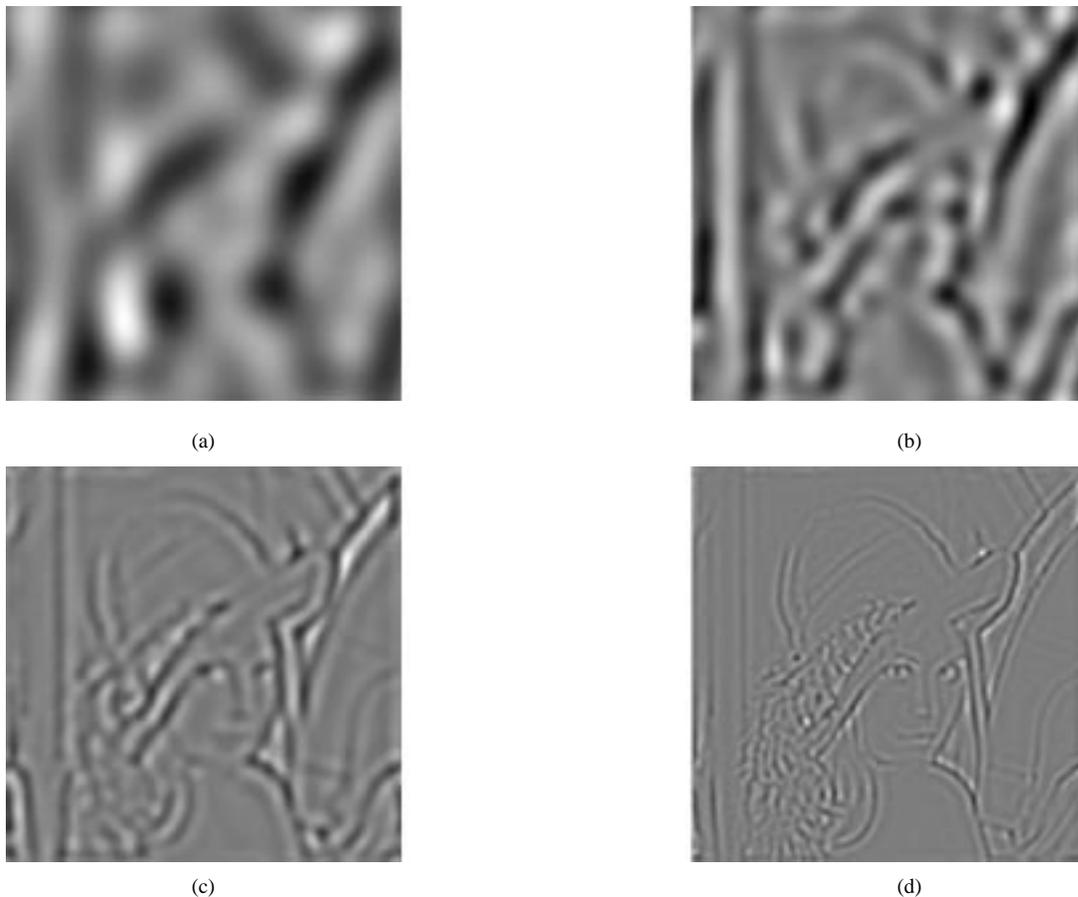


Fig. 8. Simulated contrast images at (a) 4, (b) 8, (c) 16, and (d) 32 cycles/image.

TABLE V
VARIATION IN SPATIAL FREQUENCY

Metric	Fig. 2(a)	Fig. 2(b)
SNR (dB)	10.00	10.00
WSNR (dB)	11.22	28.67
NQM (dB)	20.47	32.65

fine bandlimited contrast for a complex image, a bandlimited version of the image is obtained by filtering the image with a bank of bandpass filters. In the filter bank, we select a one-octave bandwidth to model the the cortical bandpass frequency channels [27].

In the spatial domain, a filtered image can be represented by

$$a(x, y) = f(x, y) * h(x, y) \quad (12)$$

where

* linear convolution;

$f(x, y)$ input image;

$h(x, y)$ impulse response of the filter.

For every bandpass filtered image $a(x, y)$, we define the corresponding local luminance mean image $l(x, y)$, which is a low-pass filtered version of the original image that contains the fre-

quency components below the band of $h(x, y)$. At the given band of spatial frequencies,

$$c(x, y) = \frac{a(x, y)}{l(x, y)} \quad (13)$$

defines the *local bandlimited contrast*: it assigns a local contrast at every point in the image and at every frequency channel. Since human contrast sensitivity is dependent on spatial frequency, the contrast for each spatial frequency band is calculated separately. The contrast at each point in the image is calculated separately to account for contrast variation across the image. As (13) indicates, brightness changes in the image affect the apparent contrast. Decreasing intensity over a local area increases the contrast more over dark areas than light areas for the same spatial frequency.

Local bandlimited contrast relates contrast at a particular band of frequencies with amplitudes at lower spatial frequencies. Thomas [17] validates this approach by using a 1-D grating consisting of a fundamental frequency and the eighth harmonic. He finds that the apparent contrast of the high-frequency component changes across the image even though its amplitude is fixed. The contrast is greater at the valleys than at the peaks, as predicted by local bandlimited contrast: high-frequency bands will have a greater contrast over dark areas where the corresponding points in $l(x, y)$ are low. Details below threshold in the bandpass filtered image are assumed to have no relevance in perception [28], but may be



(a)



(b)

Fig. 9. Effect of noise position on visibility. (a) Lena with spatially localized white noise added at lower center and (b) with white noise added to feathers at lower left.

TABLE VI
VARIATION IN SPATIAL POSITION

Metric	Fig. 9(a)	Fig. 9(b)
SNR (dB)	30.17	30.17
LQM (WSNR in dB)	40.28	40.28
NQM (dB)	30.26	32.00

above threshold in the contrast image, add to image sharpness, and aid in recognition. This shows the importance of including contrast effects in a quality measure.

VI. THE NOISE QUALITY MEASURE

In this section, we present a nonlinear noise quality measure (NQM) that not only accounts for many of the phenomena not

measured by LQM's, but also can potentially be extended to include other nonlinear factors. We simulate the appearance of the original and restored images to an observer. The SNR is then computed for the difference of the two simulated images as a measure of image quality. To produce the simulated images, nonlinear space-frequency processing is performed based on Peli's contrast pyramid [9]. While retaining the essential components of this scheme, we modify the pyramid in the following ways:

- 1) we define a threshold that varies for each spatial frequency band and each pixel in the bandpass images, to account for contrast masking;
- 2) we derive global thresholds for each channel based on the inverse of the CSF in [14] to incorporate information about the viewing medium, and ambient luminance [13];
- 3) we account for suprathreshold contrast discrimination effects explicitly using a contrast masking threshold.

If $O_s(x, y)$ and $I_s(x, y)$ denote the simulated versions of the model restored image and the restored images, respectively, then the NQM is

$$\text{NQM (dB)} = 10 \log_{10} \left(\frac{\sum_x \sum_y O_s^2(x, y)}{\sum_x \sum_y (O_s(x, y) - I_s(x, y))^2} \right). \quad (14)$$

On a single-processor 167 MHz Ultra-2 workstation, Matlab 5 requires 90 s to compute the NQM for 512×512 model restored and restored images.

To implement the contrast pyramid, we use a bank of cosine-log bandpass filters defined by

$$G_i(r) = \frac{1}{2} [1 + \cos(\pi \log_2(r) - \pi i)] \quad (15)$$

where each filter is centered at a frequency of 2^i cycles/image, where i is an integer. These filters, which have a one-octave bandwidth, are symmetric on a log-frequency axis, and their outputs in the frequency domain sum to one. Gabor filters, which are traditionally used, are not symmetric on a log-frequency axis and a summation of the filter frequency responses is not flat [9]. Since the filter outputs are summed to obtain simulated images, the unweighted summation of all channel responses should be constant over all spatial frequencies.

We use a bank of six filters centered at $\{0, 2, 4, 8, 16, 32\}$ cycles/image. The zero spatial frequency (DC) filter is a low-pass, shifted version of the filter centered at 2 cycles/image and is given by

$$G_{\text{dc}}(r) = \frac{1}{2} [1 + \cos(\pi \log_2(r + 2) - \pi)]. \quad (16)$$

The Fourier transform of the image is multiplied by the frequency-domain transfer functions of the bandpass filters to decompose the image into spatial frequency bands. In the frequency domain, the image may be represented as

$$F(\omega_1, \omega_2) = L_0(\omega_1, \omega_2) + \sum_{i=1}^n A_i(\omega_1, \omega_2) + H_n(\omega_1, \omega_2) \quad (17)$$



(a)



(b)

Fig. 10. Effect of adding a random function of the noise. (a) Lena with Gaussian noise added and (b) with a function [see (33)] of the noise in (a) added.

TABLE VII
EFFECT OF ADDING A SPATIALLY RANDOM FUNCTION OF ADDITIVE NOISE

Metric	Fig. 10(a)	Fig. 10(b)
SNR (dB)	16.40	16.40
LQM (WSNR in dB)	27.16	27.17
NQM (dB)	18.50	18.56

where L_0 and H_n are the low frequency and high frequency residuals, respectively. In the spatial domain, this becomes

$$f(x, y) = l_0(x, y) + \sum_{i=1}^n a_i(x, y) + h_n(x, y) \quad (18)$$

where $l_0(x, y)$ is the image obtained by filtering with the low-pass filter. The $a_i(x, y)$ terms are computed by using the filter transfer functions $G_i(\omega_1, \omega_2)$, defined as

$$a_i(x, y) = \text{IFFT}[F(\omega_1, \omega_2)G_i(\omega_1, \omega_2)]. \quad (19)$$

The high frequency residual may be omitted from consideration because the eye is virtually insensitive to these frequencies, as shown in Fig. 3(a).

For every $a_i(x, y)$, we define a local luminance mean image $l_i(x, y)$ given by

$$l_i(x, y) = l_0(x, y) + \sum_{j=1}^{i-1} a_j(x, y) \quad (20)$$

while the contrast image for each spatial frequency band is

$$c_i(x, y) = \frac{a_i(x, y)}{l_i(x, y)}. \quad (21)$$

Figs. 7 and 8 show the bandpass and contrast *lena* images, respectively. The contrast is greater in the dark regions of low luminance mean than suggested by the corresponding bandpass images. We retain all points in the corresponding bandpass image that have values above the threshold in the contrast image, and zero points below the threshold.

We derive the thresholds from the CTF. The thresholds are a function of angular frequency, and therefore depend on the angle subtended at the eye by the image. The threshold of the i th channel is

$$t_i = \text{CTF} \left(\frac{2^i}{\theta} \right) \quad (22)$$

where θ is the viewing angle in degrees subtended by the image. This is related to the viewing distance d and the image size l in each dimension by

$$\theta = 2 \tan^{-1} \left(\frac{l}{2d} \right) \approx \frac{l}{d} \quad (23)$$

where d and l have the same units. We make use of the CTF derived from threshold measurements directly, instead of using the CSF, which is a linear weighting function derived from non-linear threshold measurements. In deriving the thresholds for each channel, we use the unmodified CTF because we will account for suprathreshold contrast discrimination explicitly by deriving a set of contrast masking thresholds. By making the dependence of the channel thresholds on the CTF explicit, we can use the CTF to take into account the effect of physical parameters of the display and the surrounding environment on visual perception at various spatial frequencies [13].

Suprathreshold viewing effects are related to contrast discrimination tasks [29]. We are interested in distinguishing between contrast components that are above threshold (and that can therefore be discerned). The contrast value that may be discriminated from a background contrast depends on the background contrast [29]. Bradley and Ozawa [30] further establish that if we normalize the contrast to be discriminated and the background contrast at a particular spatial frequency by the detection thresholds for that spatial frequency, then the just-discriminable contrast is a fixed function of the background

contrast. In fact, the function may be approximated with a straight line with a slope of approximately 0.86 [30]. The same function is obtained for all spatial frequencies. So, normalized suprathreshold contrast discrimination may be regarded as invariant to spatial frequencies.

In our contrast pyramid, suprathreshold effects may be taken into account directly if we consider the contrast of the simulated model restored image as background, and the corresponding contrast component in the simulated restored image as a value to be discriminated. We can therefore ascertain whether the two contrast components will be distinguishable. If they are not distinguishable, then the two values in the corresponding bandpass images are set to be equal.

Using a linear fit to the suprathreshold contrast discrimination function of Bradley and Ozhawa [30] gives the just-discriminable contrast $T_i(x, y)$ as a function of background contrast $c_i(x, y)$ and the detection threshold of the i th channel

$$T_i(x, y) = \text{CTF}(i) \left(0.86 \left(\frac{c_i(x, y)}{\text{CTF}(i)} - 1 \right) + 0.3 \right). \quad (24)$$

The indices of the contrast components in the i th channel that have undergone imperceptible change from the contrast of the model restored image may be represented as belonging to the set

$$S_i = \{(x, y): |c_i^I(x, y) - c_i^O(x, y)| - T_i(x, y) < 0\} \quad (25)$$

where the superscript O refers to the model restored image and the superscript I refers to the restored image. The thresholding is completed by setting the corresponding elements in the bandpass image of the restored image equal to those in the bandpass image of the model restored image

$$a_i^I(x, y) = a_i^O(x, y) \quad \forall x, y \in S_i. \quad (26)$$

We apply the global thresholds of (22) to the channel images as follows:

$$V_i^O = \{(x, y): |c_i^O(x, y)| - t_i < 0\} \quad (27)$$

$$V_i^I = \{(x, y): |c_i^I(x, y)| - t_i < 0\} \quad (28)$$

$$a_i^O(x, y) = 0 \quad \forall x, y \in V_i^O \quad (29)$$

$$a_i^I(x, y) = 0 \quad \forall x, y \in V_i^I. \quad (30)$$

We then compute the net simulated bandpass images of the model restored image and restored image as

$$O_s(x, y) = l_0^O(x, y) + \sum_{j=0}^n a_j^O(x, y) \quad (31)$$

$$I_s(x, y) = l_0^I(x, y) + \sum_{j=1}^n a_j^I(x, y). \quad (32)$$

Using $O_s(x, y)$ and $I_s(x, y)$ in (14), we compute the NQM.

VII. VALIDATING THE NOISE QUALITY MEASURES

To validate the noise quality measure, we conducted a two-alternative forced choice experiment. A total of ten images were altered by adding noise. All images were 256×256 pixels in size, and were viewed at a 4 degree visual angle, after printing.

Three observers viewed the images, and we compared their responses with the predictions of the noise quality measure.

Three tests were performed. In the first test, we added Gaussian noise with different spatial frequency distributions (white and highpass), but of the same power, to the same original image chosen from a set. The noise power was different for each image in the set. The two degraded images and the original were viewed by the observers. Each observer was allotted 10 s in which to choose which degraded image he or she preferred. A typical image pair is shown in Fig. 2. SNR, LQM, and NQM values for this image are tabulated in Table V. For this test, both the LQM and the NQM produced 100% correlation with visual results, while SNR predictions were uncorrelated with visual results, since the noise power was the same in both degraded images.

In the second test, we added white noise filtered by the contrast sensitivity function to two different spatial locations in the original image. This ensures that the SNR and LQM values are identical for both of the resulting images, while the visual quality may be different. The two-alternative forced choice method outlined above was used to compute the correlation coefficients. SNR and LQM were completely uncorrelated with the visual results while the NQM produced a correlation of 80%. Fig. 9 shows a typical image pair. Table VI tabulates the SNR, LQM, and NQM values for one image used in this test.

In the third test, we print two degraded images for each of the original images. The degradation consists of spatially varying distortion and spatially varying noise. The degradation in one image is additive noise $n_a(x, y)$ generated from a uniform, Gaussian, or Laplacian random process. The degradation in the other image is also additive noise $n_b(x, y)$ that is a spatially varying function of $n_a(x, y)$ plus a spatially varying random process

$$n_b(x, y) = n_a(x, y) + r(x, y) \sqrt{\frac{x}{20}} \quad (33)$$

where $r(x, y)$ is a sample of a zero-mean truncated Gaussian random variable in the range $[-1, 1]$. A typical example of a resulting image pair is shown in Fig. 10. Table VII tabulates the SNR, LQM, and NQM values for this image pair. SNR produced a correlation of 0%, the LQM 60%, and the NQM 80%.

VIII. CONCLUSION

This paper develops a distortion measure (DM) and a noise quality measure (NQM) to quantify the impact on the human visual system (HVS) of frequency distortion and noise injection in image restoration. We derive a 2-D distortion transfer function for modeling the linear distortion effects present in restored images. After we radially average the distortion transfer function, we apply perceptual weighting to generate the DM. For the NQM, we formulate a nonlinear quasi-local processing model of the HVS by modifying Peli's contrast pyramid to measure

- 1) variation in contrast sensitivity with distance, image dimensions, and spatial frequency;
- 2) variations in the local luminance mean;
- 3) contrast interaction between spatial frequencies;
- 4) contrast masking effects.

We also show how the DM may be calculated in a practical image restoration system.

The DM and NQM quantify the two key sources of degradation in restored images—frequency distortion and noise injection. Measures based on SNR and linear HVS models do not account for frequency distortion and ignore the essential nonlinear processing of the HVS in the spatial and frequency domains. We have demonstrated the importance of taking nonlinear effects into account in the computed quality measures. Previous measures based on nonlinear HVS models are tailored to compressed images and are computationally intensive to compute. We reduce the amount of computation by not including sensitivity to orientation in our HVS model [6]. Since our quality assessment is based on independent measures for frequency distortion and noise, one can optimize the parameters of an image restoration algorithm to minimize the visual impact of both these effects. Measures that return one parameter cannot indicate the relative visual impact of the degradations that may occur. This is of key importance. An important open problem is to define a quality metric based on the two quality measures for frequency distortion and noise injection.

REFERENCES

- [1] T. D. Kite, B. L. Evans, A. C. Bovik, and T. L. Sculley, "Digital halftoning as 2-D delta-sigma modulation," in *Proc. IEEE Int. Conf. Image Proc.*, vol. 1, Oct. 1997, pp. 799–802.
- [2] Q. Lin, "Halftone image quality analysis based on a human vision model," in *Proc. SPIE*, vol. 1913, Feb. 1993, pp. 378–389.
- [3] T. Mitsa, K. L. Varkur, and J. R. Alford, "Frequency channel based visual models as quantitative quality measures in halftoning," in *Proc. SPIE*, vol. 1913, Feb. 1993, pp. 390–401.
- [4] T. Mitsa and K. L. Varkur, "Evaluation of contrast sensitivity functions for the formulation of quality measures incorporated in halftoning algorithms," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing*, vol. 3, Mar. 1992, pp. 313–316.
- [5] S. Daly, "The visible differences predictor: An algorithm for the assessment for image fidelity," in *Proc. SPIE Conf. on Human Vision, Visual Processing, Digital Display*, vol. 1666, San Jose, CA, Feb. 1992, pp. 2–15.
- [6] J. Lubin, "A visual discrimination model for imaging system design and evaluation," in *Vision Models for Target Detection and Recognition*, Singapore: World Scientific, 1995, pp. 245–283.
- [7] P. Teo and D. Heeger, "A model of perceptual image fidelity," in *Proc. IEEE Conf. Image Processing*, vol. 2, Oct. 1995, pp. 343–345.
- [8] ———, "Perceptual image distortion," in *Proc. IEEE Conf. Image Processing*, vol. 2, Nov. 1994, pp. 982–986.
- [9] E. Peli, "Contrast in complex images," *J. Opt. Soc. Amer. A*, vol. 7, pp. 2032–2039, Oct. 1990.
- [10] T. Mitsa, "Image quality metrics for halftone images," in *Proc. SPIE*, vol. 1778, Mar. 1992, pp. 196–207.
- [11] R. Ulichney, *Digital Halftoning*. Cambridge, MA: MIT Press, 1987.
- [12] B. A. Wandell, *Foundations of Vision*. Sunderland, MA: Sinauer, 1995.
- [13] P. Barten, "Evaluation of subjective image quality with the square-root integral method," *J. Opt. Soc. Amer. A*, vol. 7, pp. 2024–2031, Oct. 1990.
- [14] J. Mannos and D. Sakrison, "The effects of visual fidelity criterion on the encoding of images," *IEEE Trans. Inform. Theory*, vol. IT-20, pp. 525–535, July 1974.
- [15] T. Cornsweet, *Visual Perception*, New York: Academic, 1970.
- [16] E. Peli, L. Arnold, and R. Goldstein, "Contrast sensitivity to patch stimuli: effects of spatial bandwidth and temporal presentation," *Spatial Vis.*, vol. 7, pp. 1–14, June 1993.
- [17] J. P. Thomas, "Independent processing of suprathreshold spatial gratings as a function of their separation in spatial frequency," *J. Opt. Soc. Amer. A*, vol. 6, pp. 1102–1111, 1989.
- [18] T. Kite, "Design and quality assessment of forward and inverse error-diffusion halftoning algorithms," Ph.D. dissertation, Dept. Elect. Comput. Eng., Univ. Texas, Austin, TX, 1998.
- [19] R. Williams, *Electrical Engineering Probability*. St. Paul, MN: West, 1991.
- [20] T. D. Kite, N. Damera-Venkata, B. L. Evans, and A. C. Bovik, "A high quality, fast inverse halftoning algorithm for error diffused halftones," in *Proc. IEEE Int. Conf. Image Processing*, vol. 2, Oct. 1998, pp. 59–63.
- [21] R. Floyd and L. Steinberg, "An adaptive algorithm for spatial grayscale," in *Proc. Soc. Image Display*, vol. 17, 1976, pp. 75–77.
- [22] A. A. Michelson, *Studies in Optics*. Chicago, IL: Univ. of Chicago Press, 1927.
- [23] A. P. Ginsburg, "Visual information processing based on spatial filters constrained by biological data," Ph.D. dissertation, Univ. Cambridge, Cambridge, U.K., 1978.
- [24] R. F. Hess, A. Bradley, and L. Pirotrowski, "Contrast-coding in amblyopia. I. Differences in the neural basis of human amblyopia," in *Proc. R. Soc. Lond.*, vol. 217, 1983, pp. 309–330.
- [25] D. R. Badcock, "Spatial phase or luminance profile discrimination?," *Vision Res.*, vol. 24, pp. 613–623, 1984.
- [26] R. F. Hess and J. S. Pointer, "Evidence for spatially local computations underlying discrimination of periodic patterns in fovea and periphery," *Vision Res.*, vol. 27, pp. 1343–1360, 1987.
- [27] R. D. Valois, D. G. Albrecht, and L. G. Thorell, "Spatial frequency selectivity of cells in macaque visual cortex," *Vision Res.*, vol. 22, pp. 545–559, 1982.
- [28] A. Fiorentini, L. Maffei, and G. Sandini, "The role of high spatial frequencies in face perception," *Perception*, vol. 12, pp. 195–201, 1983.
- [29] J. Nachmias and R. Sansbury, "Grating contrast discrimination may be better than detection," *Vision Res.*, vol. 14, pp. 1039–1042, Oct. 1974.
- [30] A. Bradley and I. Ozhawa, "A comparison of contrast detection and discrimination," *Vision Res.*, vol. 26, pp. 991–997, June 1986.



Niranjan Damera-Venkata (S'00) received the B.S.E.E. degree from the University of Madras, Madras, India, in July 1997 and the M.S.E.E. degree from the University of Texas, Austin, in May 1999. He is currently pursuing the Ph.D. degree in electrical engineering at the University of Texas.

His research interests include document image processing, symbolic design and analysis tools, image and video quality assessment, and fast algorithms for image processing.

Mr. Damera-Venkata is a member of Sigma Xi. He won a 1998–1999 Texas Telecommunications Engineering Consortium Graduate Fellowship from the University of Texas.



Thomas D. Kite received the B.S. degree in physics from Oxford University, Oxford, U.K., and the M.S. and Ph.D. degrees in electrical engineering from the University of Texas, Austin, in 1991, 1993, and 1998, respectively. His M.S. thesis was in digital audio and his Ph.D. dissertation was in image halftoning.

He is now a DSP Engineer at Audio Precision, Beaverton, OR.



Wilson S. Geisler is a Professor with the Department of Psychology and the Director of the Center for Vision and Image Sciences at the University of Texas, Austin. He is currently serving as section editor for *Vision Research*. He has broad research interests within the general areas of human vision, visual neurophysiology, machine vision, and image processing.

Dr. Geisler is a Fellow of the Optical Society of America, served as a member of the Visual Science B Study Section for the National Institutes of Health, and received a Career Research Excellence Award from the University of Texas. He has chaired the program planning committees for the national meetings of Optical Society of America and for the Association for Research in Vision and Ophthalmology.



Brian L. Evans (S'88–M'93–SM'97) received the B.S.E.E.C.S. degree from the Rose-Hulman Institute of Technology, Terre Haute, IN, in May 1987, and the M.S.E.E. and Ph.D. degrees from the Georgia Institute of Technology, Atlanta, in December 1998 and September 1993, respectively.

From 1993 to 1996, he was a Postdoctoral Researcher with the University of California, Berkeley, where he worked with the Ptolemy Project. (Ptolemy is a research project and software environment focused on design methodology for signal processing, communications, and controls systems.)

In addition to Ptolemy, he has played a key role in the development and release of six other computer-aided design frameworks. He is the primary architect of the *Signals and Systems Pack* for Mathematica, which has been on the market since October 1995. He is currently an Assistant Professor with the Department of Electrical and Computer Engineering, University of Texas, Austin (UT Austin). He is also the Director of the Embedded Signal Processing Laboratory within the Center for Vision and Image Sciences. His research interests include real-time embedded systems; signal, image and video processing systems; system-level design; symbolic computation; and filter design. At UT Austin, he developed and currently teaches multidimensional digital signal processing, embedded software systems, and real-time digital signal processing laboratory.

Dr. Evans is an Associate Editor of the IEEE TRANSACTIONS ON IMAGE PROCESSING, and the recipient of a 1997 National Science Foundation CAREER Award.



Alan C. Bovik (S'80–M'80–SM'89–F'96) received the B.S., M.S., and Ph.D. degrees in electrical engineering in 1980, 1982, and 1984, respectively, from the University of Illinois, Urbana-Champaign.

He is currently the General Dynamics Endowed Fellow and Professor in the Department of Electrical and Computer Engineering and the Department of Computer Sciences, University of Texas, Austin, where he is also the Associate Director of the Center for Vision and Image Sciences, which is an independent research unit that brings together

electrical engineering, computer science, and psychology professors, staff, and students. This paper is a product of the interdisciplinary work at the Center. His current research interests include digital video, image processing, computer vision, wavelets, three-dimensional microscopy, and computational aspects of biological visual perception. He has published more than 250 technical articles in these areas and holds U.S. patents for the image and video compression algorithms VPIC and VPISC.

Dr. Bovik is the Editor-in-Chief of the IEEE TRANSACTIONS ON IMAGE PROCESSING and is on the Editorial Board for the PROCEEDINGS OF THE IEEE. He is the Founding General Chairman, First IEEE International Conference on Image Processing, which was held in Austin in November 1994.