Contrast sensitivity function and image discrimination

Eli Peli

Schepens Eye Research Institute, Harvard Medical School, Boston, Massachusetts 02114

Received November 29, 1999; accepted July 14, 2000; revised manuscript received July 21, 2000

A previous study tested the validity of simulations of the appearance of a natural image (from different observation distances) generated by using a visual model and contrast sensitivity functions of the individual observers [J. Opt. Soc. Am. A **13**, 1131 (1996)]. Deleting image spatial-frequency components that should be undetectable made the simulations indistinguishable from the original images at distances larger than the simulated distance. The simulated observation distance accurately predicted the distance at which the simulated image could be discriminated from the original image. Owing to the 1/*f* characteristic of natural images' spatial spectra, the individual contrast sensitivity functions (CSF's) used in the simulations of the previous study were actually tested only over a narrow range of *retinal* spatial frequencies. To test the CSF's over a wide range of frequencies, the same simulations and testing procedure were applied to five contrast versions of the images (10–300%). This provides a stronger test of the model, of the simulations, and specifically of the CSF's used. The relevant CSF for a discrimination task was found to be obtained by using 1-octave Gabor stimuli measured in a contrast detection task. The relevant CSF data had to be measured over a range of observation distances, owing to limitations of the displays. © 2001 Optical Society of America

OCIS codes: 330.4060, 330.1800, 330.1880, 330.5000, 330.6100, 330.6110.

1. INTRODUCTION

Simulating the appearance of a scene or an image to an observer is a useful design and analysis tool. Such pictorial representations have been attempted by many investigators over the years in a variety of applications in vision science[1–4] and engineering.[5–7] Such simulations are frequently generated within the context of a computational vision model. One such multiscale model of spatial vision was used to calculate local band-limited contrast in complex images.[8] This contrast measure, together with observers' contrast sensitivity functions (CSF's), expressed as thresholds, has been used to simulate the appearance of images to observers, taking into account many of the nonlinearities inherent in the visual system. The same concept of local band-limited contrast with small variations applied by Daly,[9] Duval-Destin,[10] and Lubin[6] was found to be useful in comparing image quality[9] and in other applications.[6]

In a previous study, Peli[11] tested and demonstrated the validity of the visual model, using simulations of the appearance of complex images. The simulated images were generated with the model to represent the appearance of the original images from various observation distances. Observers viewed the images (simulated by using their individual CSF's) from a wide range of distances side by side with the original image and attempted to discriminate the original from the simulated image. The distance at which discrimination performance was at threshold was compared with the simulated observation distance. Since the distances matched, the simulations were validated.

That study also sought to determine what CSF data should be used in this or any other vision models of this type. As has been shown previously, methodological changes can account for the large variability of CSF data in the literature.[12] However, we do not yet know which, if any, of the CSF's obtained with various psychophysical methods and stimuli is appropriate to the representation of complex image perception in the context of pyramidal multiscale vision models. The previous study[11] demonstrated that the CSF obtained by using grating patches with a constant size of 2 deg $\times$ 2 deg was inadequate for use in the simulation. Further, it compared the use of the CSF obtained with 1-octave Gabor patches (constant-bandwidth) stimuli in an orientation discrimination task with the CSF obtained with the same stimuli in a contrast detection task. The CSF obtained with the orientation discrimination task was not adequate either, but the CSF obtained in the detection task could not be rejected. The variable distance simulation and the testing method were shown to be sensitive, permitting clear discrimination of image appearance that resulted from a mere doubling of viewing distance and that was affected by small differences (as induced by a high-frequency residual). The main limitation of the previous study[11] was the fact that the validity of the CSF was tested at one retinal spatial frequency only, as will be explained next.

In using this and other vision models in simulations and other applications, one needs to consider both the object's contrast spectrum (given in terms of cycles per object or cycles per image) and an observer's CSF [expressed in terms of cycles per degree]. For the purpose of illustration I shall use a one-dimensional diagram. To express the object's spectrum [Fig. 1(a), line with a slope of approximately -1.0] as retinal image spectrum, one needs to know the angular size of the object at the observer's retina. The multiple scales for the horizontal axes in Fig. 1 express these relations for different observation

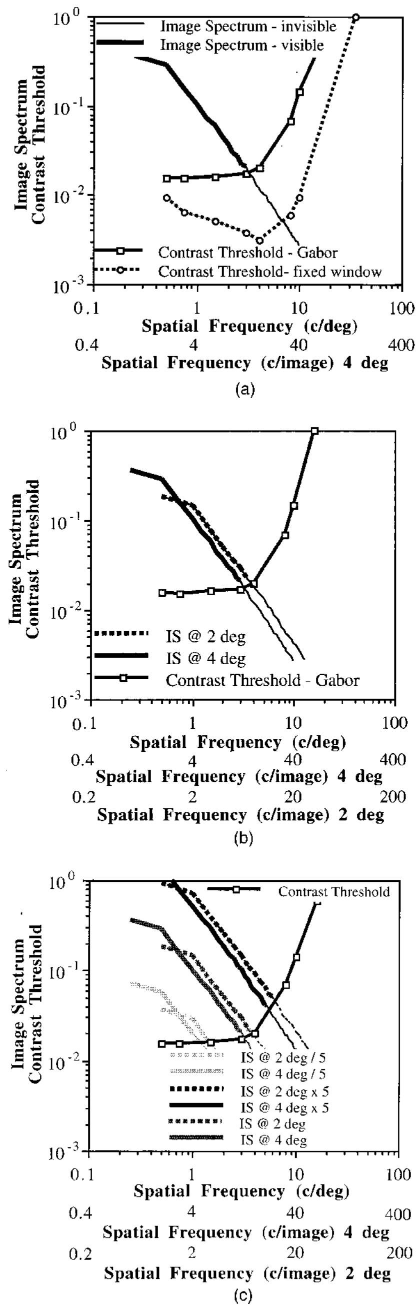

Fig. 1. (a) Schematic illustration of the interaction of image spatial-frequency content with the observer's CSF. The thick line represents a typical image spectrum (changing as $1/f$). The transformation of spatial frequencies from units of cycles per image to units of cycles per degree is determined by the image size of 4 deg. The part of the spectrum below the observer's CSF (detection threshold obtained with Gabor stimuli) will not be detectable, as illustrated by the change of the spectrum line from a thick to a thin line. The fixed window contrast threshold represents the CSF that was rejected by Peli's[11] study. As can be seen here, a single retinal frequency testing is sufficient to distinguish the two CSF's. (b) A change in observation distance, which causes the image to shrink to 2 deg on the observers' retina, shifts the corresponding image spectrum, IS, along a slope of -1.0. At the new distance lower object frequencies are removed by the observer's CSF, but essentially the same retinal frequencies are involved. (c) The additional spectral curves represent the spatial spectra of images with increased and decreased contrast that shift the intersection of the spectra with the threshold to higher and lower retinal frequencies, respectively, permitting testing of other parts of the CSF.

distances. Any information in the image that falls below the observer's threshold (i.e., below the point at which the contrast threshold curve intersects the image spectrum curve) is treated by the model as not visible to the observer. To account for this, the simulation should remove all that information. This is illustrated by the change of the spectrum line into a thin line at the values that are below threshold in Fig. 1(a). The operation illustrated in Fig. 1(a) is a linear filtering operation, applied globally to the whole image. The processing actually used in the study is spatially variable and is applied frequency band by frequency band to a nonlinear function of the image, resulting in a highly nonlinear operation. Note that in Fig. 1 the CSF is presented as a contrast threshold function. This emphasizes the way the CSF is actually being applied in our model as a threshold function and not as a linear filter (which is how it is typically being applied; see Ginsburg[1] and Lubin).[6]

If the original and the simulated images (obtained by removing all subthreshold components) are viewed from the simulated distance or farther away, they should be indistinguishable, because the information from the original that would be lacking as a result of the observer's visual response was removed in the simulation as well. However, if the original and the simulation are viewed from a closer distance, the difference in content between the original and the simulation should be visible. This requires also that the CSF (contrast threshold) used in the simulation indeed be representative of the observer's sensitivity in actually performing the discrimination task. If the CSF used in the simulation is incorrect and the observer's sensitivity, for example, is represented by the second CSF [dashed curve in Fig. 1(a)] the observer will be able to discriminate the simulation from the original at a much farther distance than that assumed in the simulation.

As the size of the object displayed on the observer's retina gets smaller when the distance of the object from the observer increases, its retinal spatial frequencies increase. It was previously thought by this author[13] and others[14] that this change results in a shift of the spectrum to the right along the spatial frequency axis [in Fig. 1(b)]. The spectrum in this case referred to the Fourier amplitude of the image radially averaged across orientation. However, as Brady and Field[15] pointed out, the spectrum actually shifts both to the right (higher frequencies) and down (lower contrast), sliding along the line with a slope of -1.0. Most natural images have a spatial-frequency amplitude spectrum that behaves approximately as $1/f$, which also has a slope of approximately -1.0 on this graph.[16–20] Thus a change in object size causes such a spectrum to "slide along itself" [Fig. 1(b), 2-deg spectrum]. As a result, the spectrum of the farther image intersects the CSF curve at essentially the same *retinal* frequencies. Only the mapping of the relevant *object* frequencies to retinal frequencies changes. Therefore the experiments by Peli[11] probed only a very limited range of retinal spatial frequencies in the contrast threshold function.

To examine the CSF at other frequencies, one needs to use images whose spectra intersect the CSF at other retinal frequencies. This was achieved in the current study by using higher- and lower-contrast versions of the same images, as illustrated schematically in Fig. 1(c). Changing the image contrast shifts its log spectrum vertically only up or down (for decrease and increase in contrast, respectively). As can be seen in Fig. 1(c), such a conversion results in images that intersect the CSF at different frequencies.

The simulations were tested by presenting the original image side by side with the simulation of its appearance from a certain distance. If the simulations are valid, the simulated image and the original should be indistinguishable from a distance equal to or farther than the distance assumed in the simulation.[11,13] The two images should be progressively easier to distinguish at distances shorter than the simulated distance.

For the following reasons, the analysis represented in Fig. 1 cannot replace the information we seek from the simulations and from direct testing of the simulation. The effects of contrast threshold on apparent contrast in the images are local, not global as represented in the figure. The effective contrast is not accurately represented by the (one-dimensional) radially averaged *amplitude* spectrum, because in the simulations we were working with local contrast, not amplitude,[8] and thus the simulation algorithm is not represented accurately by the essentially linear filtering depicted in the schematic of Fig. 1. In the experiments described here the general concept illustrated in Fig. 1 was tested directly, enabling us to probe the CSF over a wide range of frequencies and revalidating the use of such model for image simulations and other applications.

2. METHODS

A. Observers

Four observers were tested, although not all under all experimental conditions. The observers ranged in age from 25 to 30 years and had 20/20 corrected vision as determined by a Snellen chart. Three of the subjects were experienced psychophysical observers, and one of them, AL, had been a subject in the previous study. The fourth subject, JML, was a novice psychophysical observer and was not familiar with either the contrast sensitivity measures or the discrimination task.

B. Stimuli and Apparatus

Observers viewed image pairs from various distances and were asked to make a forced-choice distinction between the simulated and the original image. The observers indicated which of the two images appeared blurrier. The simulated images used to test each observer were calculated by using her or his individual CSF. Four different scenes each at five different contrasts were used in this experiment. For each image, three simulated views were generated, representing views from three different distances (106, 212, and 424 cm). For the three simulated observation distances, the images spanned visual angles of 4, 2, and 1 deg, respectively. The simulated distance and the corresponding span in degrees served to establish the relations between the subject's CSF expressed in cycles per degree and the image spatial content expressed in terms of cycles per image.

The CSF data used in the simulations were obtained for each subject individually. The CSF's were obtained with 1-octave Gabor patches and a simple detection task. Data were collected on a Vision Works system (Durham, N. H.), with an M21LV-65MAX monitor with DP104 phosphor operating at 117 Hz, noninterlaced. The stimuli were the same Gabor patches of 1-octave bandwidth in all cases (vertical orientation only).

The image pairs were presented on a 19-in. (48-cm) 1148×896-pixel, noninterlaced monochrome video monitor of a Sparc 10 Workstation (Sun Microsystems, Mountain View, Calif.). Linearity of the display response was obtained with an 8-bit lookup table.[21] The screen calibrated with the lookup table provided a linear response over a 2-log-unit range. The images were 256×256 pixels each and were presented side by side at the middle of the screen, separated by 128 pixels. The background luminance around the images was set to 40 cd/m^2, a value that was close to the average mean luminance of all images.

The four images were common images frequently used in image processing.[22] The original unprocessed images were also produced at varying contrasts.[23] This was achieved by subtracting the mean luminance level from the image, multiplying each pixel by the corresponding contrast (0.1, 0.3, 0.5, and 3.0), and adding the mean luminance back. The 300% contrast image was saturated wherever the dark or bright values exceeded the dynamic range of the display. Examples of the various contrast versions of one of the images and their simulated appearance from the three distances are presented in Fig. 2.

C. Simulations

To simulate their appearance from various distances, the images were processed assuming the corresponding visual angle. The details of the simulation method are given by Peli.[8] Briefly, the image is sectioned into a series of bandpass-filtered versions of 1-octave bandwidth and separated by one octave. For each section we calculated the corresponding local band-limited contrast for each point in the image. This was done by dividing the bandpass-filtered image, point by point, by the corresponding low-pass-filtered image for the corresponding scale.[8] This local band-limited contrast is different from the contrast expression used in other models. In other models the local amplitude is divided by the global luminance mean to derive a contrast expression. The global contrast is therefore a linear function of the amplitude, whereas the local band-limited contrast is spatially variable (nonlinear) function of the amplitude.

On the basis of the simulated distance, the spatial frequency in cycles per degree (c/deg) associated with the band-pass-filtered version was determined. Each spatial point at each frequency band can be tested against the appropriate threshold taken from the individual CSF to determine whether it will be visible. A suprathreshold point is left unchanged, and a subthreshold point is set to zero contrast. Note that the threshold is applied to the band-pass-filtered amplitude on the basis of the corresponding local band-limited contrast function values. The thresholded band-pass-filtered images are then combined to generate the simulated image.

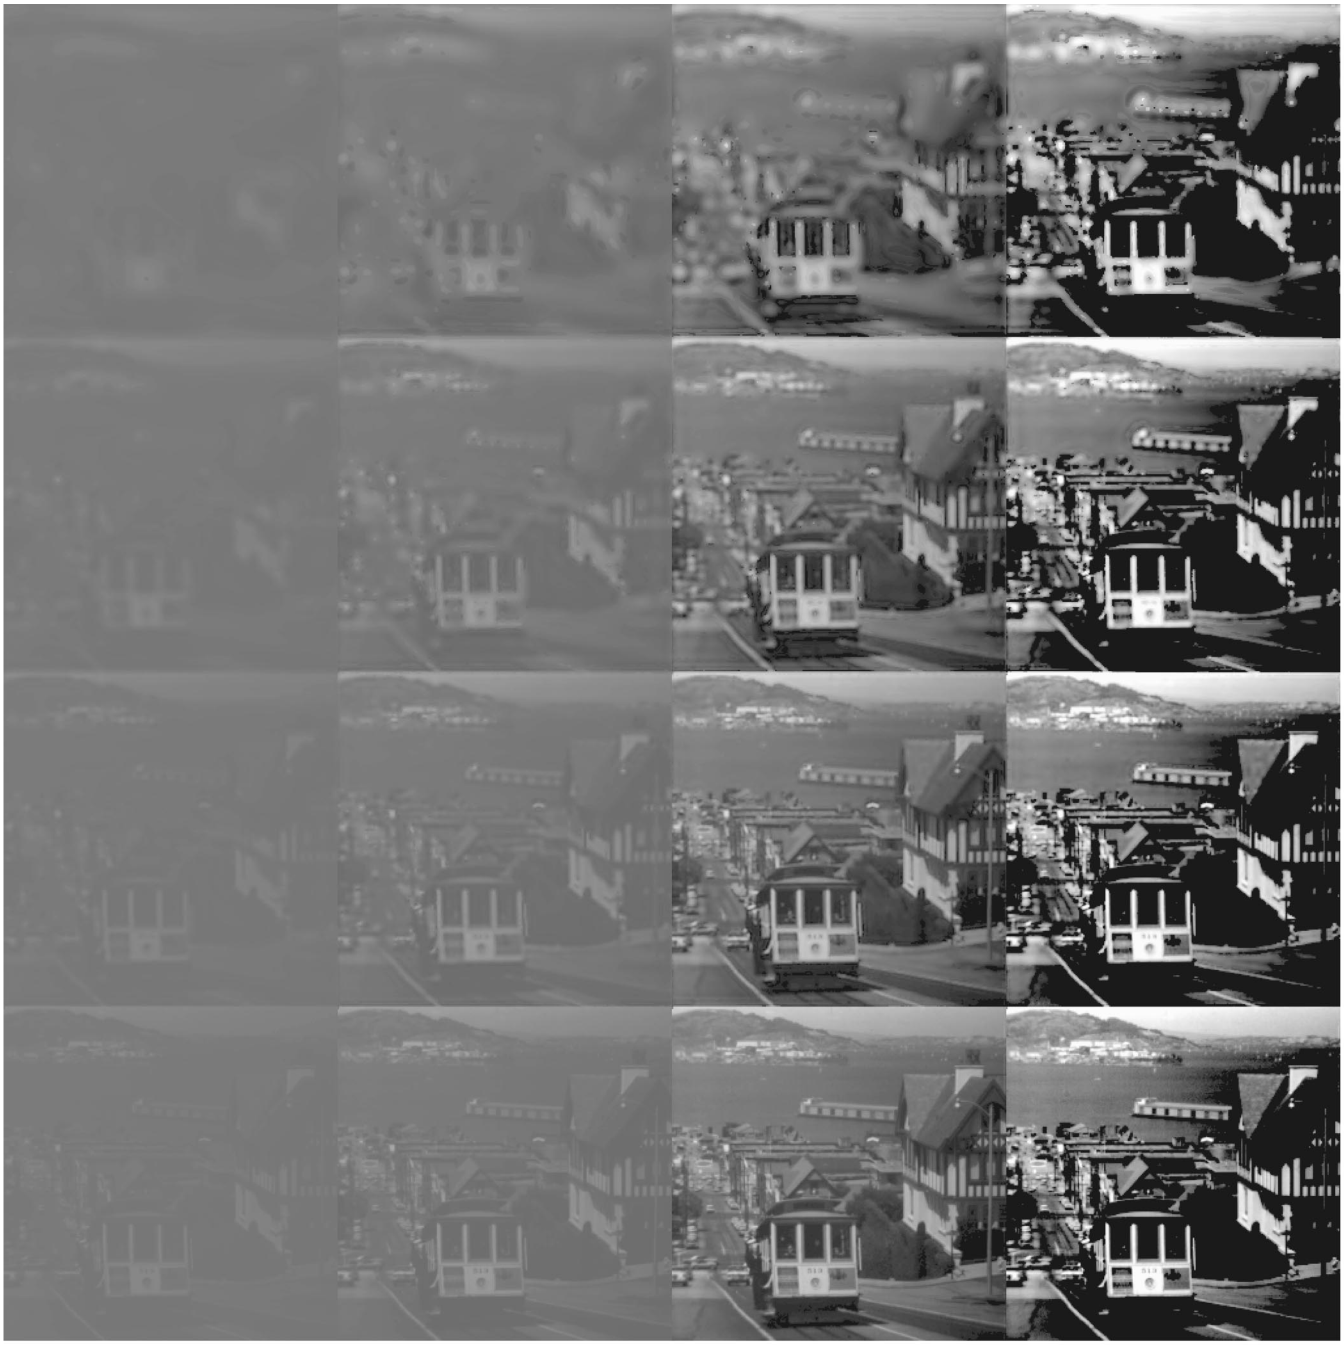

Fig. 2. Examples of the images used in the study. The original unprocessed versions at various contrast levels are shown in the bottom row. The columns from left to right represent images with 10%, 30%, 100%, and 300% contrast. The simulations of images spanning 1, 2, and 4 deg are shown in the first, second, and third row, respectively. The appearance of the simulations of the other scenes at 100% contrast can be found in Fig. 6(a) below.

D. Testing Procedure

CSF data were collected with the method of adjustment[24] (MOA). Six responses at each frequency were averaged, and the order of tested frequencies was randomized. The first experiment was conducted with simulations calculated by using individual CSF data measured from a fixed 2-m observation distance. The display size and resolution limited the range of frequencies measured from this distance to 0.5–16 c/deg. The CSF values needed for the simulations at frequencies outside this range were extrapolated by extending the low- and high-frequency limbs of the CSF linearly.[13] For reasons explained below, the contrast sensitivity was remeasured for three of the four subjects with the same system, stimuli, and procedure, but the observation distance was varied to permit extension of the tested frequency range. The shortest distance of 0.5 m transferred the lowest frequency tested from 0.5 to 0.125 c/deg. The three lowest frequencies were measured from this distance. The farther distances of 4 and 8 m permitted testing at frequencies as high as 24 c/deg (our observers could not detect the 32-c/deg Gabor stimuli at any contrast).

For the image discrimination task, observers were seated in a dimly lit room and adapted to the mean luminance of the display for 5 min before beginning the experiment. The observers indicated the location of the simulated image (right or left) by using the right and left buttons on a mouse. A new pair of images emerged abruptly 0.1 s after each response and remained on until the subject responded. The order of observation distances was randomized.

The subjects viewed the image pairs from nine distances, including ones shorter (53 cm) than the shortest simulated distance and longer (848 cm) than the longest simulated distance. Each image at each simulation distance was presented ten times at each viewing distance. The position of the simulated image relative to the original (right or left) was randomly selected for each presentation. The observers indicated which of the two images appeared blurrier. No feedback was given to the subject.

E. Data Analysis

From each observation distance the percent correct identification of the processed/simulated image was calculated for each simulated distance for the four images. The data for each simulated distance (percent correct out of 40 responses for each observation distance) was fitted with a Weibull psychometric function to determine threshold at a 75% correct level. The distance at which the subject performed at the 75% level was compared with the simulated distance. If the simulations and the CSF used in the simulation represent the subject perception correctly, the measured and simulated distance should be equal.

3. RESULTS

A. Image Discriminations with the Contrast Sensitivity Function Obtained from 2 m

If the simulations were veridical, the fitted Weibull curves should have crossed the 75% correct level at the simulated distance and thus all points in Fig. 3 should lie on the diagonal line. As can be seen in Fig. 3(a), the results of the first experiment were veridical only for the images in the 30–100% contrast range, even for the most practiced subject (AL, who had participated in a previous study employing a similar task[11]). For these moderate-

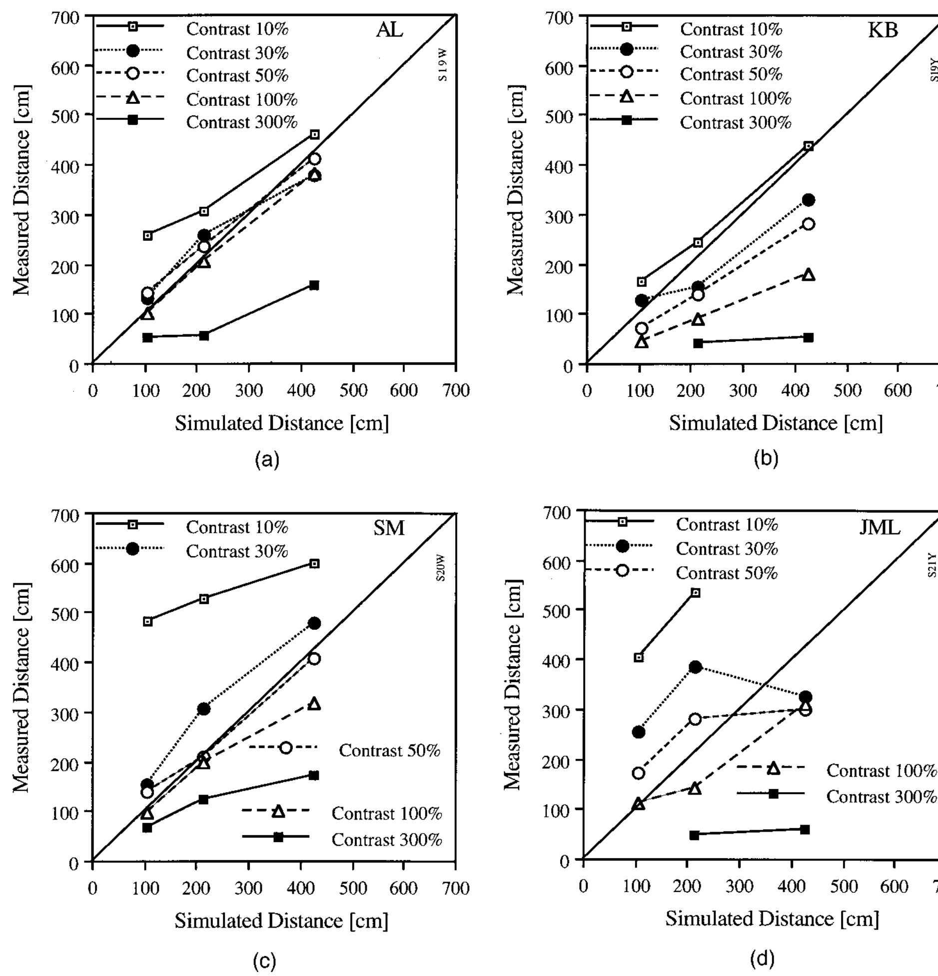

Fig. 3. Distances at which the simulated images were distinguished from the corresponding original images compared with the simulated observation distance. (a) For a well-practiced subject the data deviate from the prediction (diagonal line) only for the extreme contrast conditions corresponding to detection of low spatial frequencies (10% contrast) and high spatial frequencies (300% contrast). (b) For a novice subject the simulated images were distinguished from the original image at a distance shorter than the simulated distance (c) and (d). Similar results were obtained for two more subjects. For all subjects the deviations of different contrast lines from each other are regular and consistent.

contrast images the distance at which the original was distinguished from the simulation was very close to the simulated distance. The 10% contrast image was discriminated at distances larger than the simulated distances, indicating that the CSF values used in the simulations at low frequencies were too low. Stated otherwise, the thresholds implemented in the simulations were too high, removing more image features than appropriate and thus making the discrimination task easier.

The 300% image was discriminated at a shorter distance, indicating that the CSF values used for the simulations at the high frequencies were too high (thresholds too low). The results for a second subject (KB), who was well trained in psychophysical tests but was novice to this task, are shown in Fig. 3(b). For this subject, performance was overall poorer, requiring shorter observation distances to distinguish the simulation image from the originals. In addition, the results for the various contrast versions for this subject differ even more for the moderate-contrast versions as compared with the results for AL. The results for two more subjects [Figs. 3(c) and 3(d)] were similar to those of subject AL in that they were centered around the diagonal prediction line, but their variability was larger, i.e., of the same order as the results of subject KB. Note that in all cases the relative positions of the various lines on the graph were orderly and similar, indicating a consistent performance rather than just noisy data.

As previously mentioned,[11] these results illustrate that, using this methodology, one can reject values of the CSF data used for the simulation. The addition of the image contrast variable in this experiment enables us to test the CSF along a wider range of retinal frequencies than that tested by Peli.[11] We note that the sensitivities measured by the CSF procedure used at both the low and high ends of the range of frequencies were not representative of the observers' perception in the task. In particular, the data suggest that the individual CSF measured and used in the simulation underestimated the observer's sensitivity at low spatial frequencies and overestimated the sensitivity at high spatial frequencies. Since extrapolated CSF values were used at both ends of the frequency range in the simulation, further experiments were carried out to determine whether the deviation at low and high contrast was a result of an error introduced through the use of extrapolations instead of measurements of the CSF.

B. Image Discrimination with the Combined Contrast Sensitivity Function

As can be seen in Fig. 4, the CSF for the low frequencies taken at the shorter distance resulted in higher manifest sensitivity, as was predicted by the simulation results of the first experiment. The CSF at the high frequencies taken from longer distances of 4 and 8 m were almost overlapping. These results, at high frequencies, were substantially lower in sensitivity in comparison with the data measured and extrapolated from the 2-m measurements. These changes in sensitivity are also consistent with the results obtained in the simulations, suggesting that the contrast sensitivity of the observers in the task is better represented by the CSF values measured (at the corresponding distances) and not those extrapolated from the CSF obtained at 2 m. It should be noted that except for the 20- and 24-c/deg conditions, the new measurements in all other cases used the same physical stimuli used at the 2-m distance. Possible reasons for the different results are presented in Section 4.

To verify the effect of the CSF used in the simulation, the procedure of the previous experiment was repeated for two subjects with the CSF obtained by combining the data from the various observation distances. The short distance (0.5-m) CSF was used for the low spatial frequencies, the 2-m measurements for intermediate frequencies, and the 4-m measurements for high frequencies. The CSF at 32 c/deg used in the simulations was extrapolated from values at 8, 16, 20, and 24 c/deg. The simulations were recomputed by using the combined CSF functions presented in Fig. 4, and the testing was repeated. The results, shown in Fig. 5, clearly show a convergence of the data toward the diagonal line for subject AL. Subject KB shows a substantial convergence of the data from various contrast versions, and in addition this subject discriminated the images from a farther distance overall. This improvement may be accounted for by the increased familiarity with the task. For both subjects the deviations from the predicted distance of distinguish-

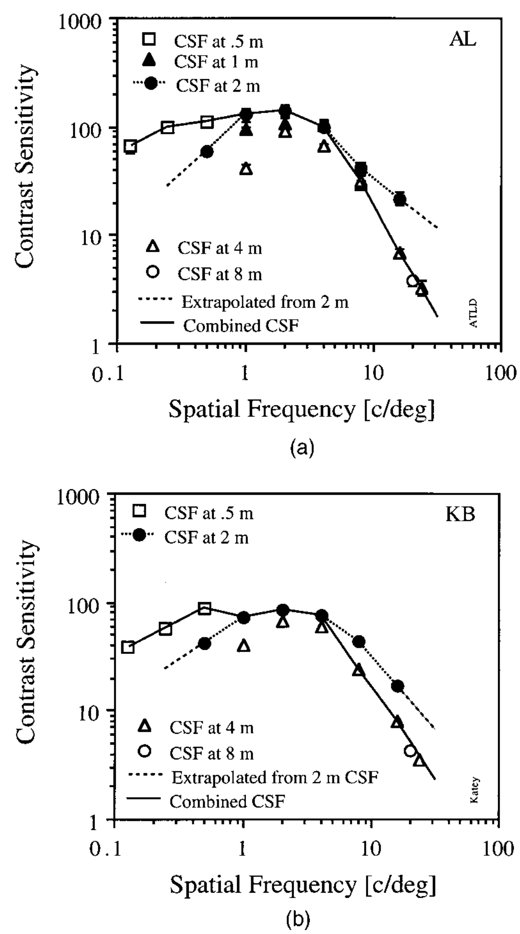

Fig. 4. CSF data measured for two subjects at different observation distances. The data collected at 2 m distance together with the illustrated extrapolations were used in the first experiment. The data shown by a solid line marked "combined CSF" was used in the simulations of the second experiment.

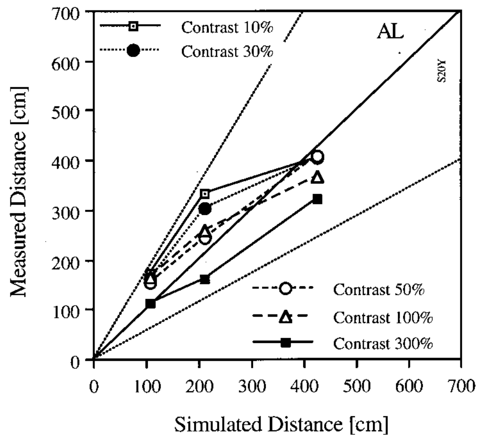

(a)

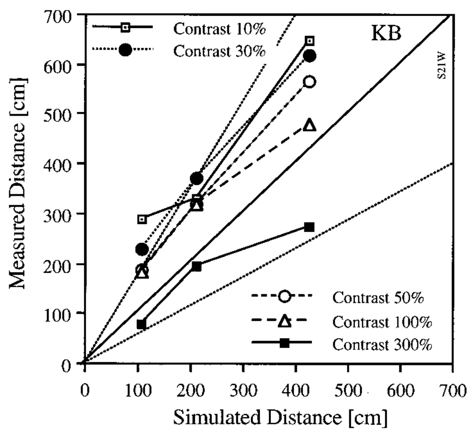

(b)

Fig. 5. Distances at which the simulated images were distinguished from the corresponding original images compared with the simulated observation distance, for the two of the subjects in Fig. 3. Here the simulations were computed with the combined CSF's obtained from different observation distances. (a) For the well-practiced subject the data with the combined CSF are now very close to the prediction represented by the diagonal solid line. (b) For the novice subject the practice gained in the task resulted in the simulated images here being distinguished from the original image at a distance farther than the simulated distance. In addition, the different contrast versions are detected closer to each other and closer to the prediction line than in Fig. 3(b). The dotted lines include all observation distances that deviate from the simulated distances by a factor of 2. As can be seen, all data points for one subject and most of the data for the other are included in this range.

ing the original from the simulation is reduced in comparison with the data of Fig. 3. In particular the values for the 10% and 300% contrast images converge toward the other values. The results for the 300% contrast image remain separated from the rest of the samples. Since the 300% contrast images tested the CSF at high spatial frequencies, this result indicates that the observers' perception in the task is represented by even lower sensitivity than that measured from the 4-m observation distance.

4. DISCUSSION

The results of these experiments verified the model proposed by Peli[8] again and justify its use to simulate the appearance of an image from different observation distances. The changes that occur with parameter changes are consistent and orderly. The simulated images are distinguishable from the original at distances close to the simulated distances (in all cases the error is less than that of doubling the observation distance: Fig. 5, area between the dotted lines). The size of the effects that occurs when the observer's distance from the display is doubled is small and of the magnitude of interest in image-quality metrics. Since we are able to simulate such effects accurately by using the vision model employed here, it stands to reason that such models could be employed successfully to calculate such differences in order to estimate image quality.[6,9] In addition, the current study has demonstrated that this method can be used to test the applicability of a specific empirically derived CSF to the performance of a discrimination task with complex images.

The vision model applied here[8] differs from many previous models in two respects. First, the CSF is applied here as a nonlinear threshold function and not as a linear filter. When the CSF is applied as a linear filter it is usually applied to the amplitude of the image. When the CSF is applied as a threshold function it is generally applied to the contrast, not the amplitude. The second difference is that here the threshold is applied to the local band-limited contrast, computed by normalizing the local luminance variations (band-pass-filtered amplitude) by the local luminance mean. In many other cases the thresholds have been applied to the globally normalized contrast (obtained by dividing the amplitude by the global luminance mean[25,26]), which is equivalent to operating in the amplitude domain rather than in the contrast domain. The latter difference between these two approaches in computing the simulations was illustrated in Peli's[11] Fig. 2.

The use of the local band-limited contrast (as described in Subsection 2.C) is now widely accepted.[6,9,20] However, most models continue to apply the CSF as a linear filter for predicting the appearance of complex visual images,[1,2,6,27] although the analysis of experimental results with simple patterns has frequently been based on the detection-threshold concept.[28,29]

The CSF values commonly presented and used as linear filter functions are computed as an inverse of the measured thresholds. The values obtained this way are larger than one (1.0), but filter functions cannot exceed the value of one. Thus the CSF values can be applied as filter values only after application of an arbitrary scaling factor. In most cases the CSF values are normalized to a value of 1.0 at the maximum sensitivity (at frequency of 2–4 c/deg).[1]

To illustrate the difference between the linear filtering approach and our nonlinear processing, I compared the nonlinear simulations used in the study [Fig. 6(a)] with simulations generated by using the linear approach [Fig. 6(b)]. The processing was applied band by band in both cases, with the same contrast threshold values used for both cases. The linear filter values were normalized to 1.0 at the maximal value. Note that any other possible normalization will result in lower filtering levels and will cause blurrier images than those shown in Fig. 6(b). As

(a)

Fig. 6. Continues on facing page.

can be seen from Fig. 6 the two processes are not equivalent. The simulated images generated by the linear filtering [Fig. 6(b)] are much blurrier than those used in the current study [Fig. 6(a)]. It is therefore clear that observers would have distinguished the linear filtering images at distances much larger than the simulated distances. The differences are so large that it is obvious that the results of the current study argue against the use of the linear filtering approach as a representation of image appearance with a given CSF.

What are the possible reasons for the differences between the CSF's obtained at different observation distances? The low-frequency end is simple to account for. The low-frequency Gabor patches used from a distance of 2 m were quite large, physically occupying a substantial part of the CRT screen. The edge of the screen (outside the active video area) is dark and creates a high-contrast feature that, when close to the patch, may mask its visibility.[30] Moving the observer closer to the screen reduces the physical size of the patches on the screen (for the same spatial frequencies) and thus increases their distance from the edge and reduces the masking effect. Indeed, for both subjects the detection threshold for the three lowest spatial frequencies was almost equal at 2 m

(b)

Fig. 6. Comparison of (a) the simulations (of 100% contrast versions of the images) used in this study with (b) the simulations obtained with linear filtering of the images by using the normalized CSF as the filter function. In the two cases the same contrast detection data was used and was applied band by band. The linearly filtered images are much blurrier than those used here. The linearly filtered simulations would be distinguishable at a distance much larger than the simulated observation distance and are therefore inadequate. For each scene each column and row represent the same simulations as the columns and rows in Fig. 2.

and 0.5 m (which were the same physical stimuli), suggesting that the reduction in sensitivity for these Gabor patches at low frequency is mostly a masking effect. This result suggests that the real CSF was even higher sensitivity at low frequency than represented by the combined CSF in Fig. 4.

The explanation for the change in CSF for high spatial frequencies with increase in distance is not as obvious. Although the high-spatial-frequency targets were physically small on the screen at 2 m distance, apparently there was sufficient resolution to represent the Gabor patch adequately (~8 pixels cycle). The answer to this puzzle emerged in a recent study of CRT artifacts.[31] We found that when high-frequency vertical gratings (as large as 10 pixels/cycle) were presented at high contrast, asymmetry in the CRT response resulted in a significant drop in local mean luminance (an effect that is not found for horizontal gratings). A similar drop in mean luminance of a CRT for a high-frequency, high-contrast pattern was previously reported by Mulligan and Stone.[32] Thus it is likely that when measurements were made from the 2-m distance, observers actually detected the change in local luminance rather than the contrast of the grating, which resulted in an apparent increase in sensitivity. Reducing the observation distance reduced the effect though probably did not eliminate it.

The methods of simulation and of testing the simulation by using the paradigm presented here are sensitive enough to be affected by the differences among CSF's obtained with different methods. As was shown here they are also sensitive enough to distinguish CSF's obtained at different distances. Thus this methodology can be used to determine the type of CSF data that more closely represents the appearance of images.

With the same method it may be possible to determine the shape of the CSF directly from simulation experiments by generating the simulation from an array of arbitrary threshold values rather than from measured CSF curves. Such a determination is independent of the specific stimuli used for the CSF measurement and may provide us with a CSF that should be used in conjunction with visual models. Discrimination of moving video segments can be used in a similar way to determine the spatiotemporal characteristics of the CSF that affects perception.

ACKNOWLEDGMENTS

This research has been supported in part by grants EY05957 and EY12890 from the National Institutes of Health and grant DE-FG 02-91ER61229 from the U.S. Department of Energy. I thank Angela Labianca for valuable technical help and Brian Sperry for programming support.

Send correspondence to Eli Peli, Schepens Eye Research Institute, Harvard Medical School, 20 Staniford Street, Boston, MA 02114. Phone, 617-912-2597; fax, 617-912-0111; e-mail, eli@vision.eri.harvard.edu.

REFERENCES AND NOTES

1. A. P. Ginsburg, "Visual information processing based on spatial filters constrained by biological data," Ph.D. dissertation (Cambridge University, Cambridge, UK, 1978).
2. B. L. Lundh, G. Derefeldt, S. Nyberg, and G. Lennerstrand, "Picture simulation of contrast sensitivity in organic and functional amblyopia," Acta Ophthalmol. **59**, 774–783 (1981).
3. D. Pelli, "What is low vision?" Videotape, Institute for Sensory Research, Syracuse University, Syracuse, N.Y., 1990.
4. L. N. Thibos and A. Bradley, "The limits of performance in central and peripheral vision," in *SID '91 Digest of Technical Papers* (Society for Information Display, Playa del Rey, Calif., 1991), Vol. XXII, pp. 301–303.
5. J. Larimer, "Designing tomorrow's displays," NASA Tech. Briefs **17**, 14–16 (1993).
6. J. Lubin, "A visual discrimination model for imaging system design and evaluation," in *Vision Models for Target Detection*, E. Peli, ed. (World Scientific, Singapore, 1995), Chap. 10, pp. 245–283.
7. E. Peli, R. B. Goldstein, G. M. Young, C. L. Trempe, and S. M. Buzney, "Image enhancement for the visually impaired: simulations and experimental results," Invest. Ophthalmol. Visual Sci. **32**, 2337–2350 (1991).
8. E. Peli, "Contrast in complex images," J. Opt. Soc. Am. A **7**, 2030–2040 (1990).
9. S. Daly, "The visual differences predictor: an algorithm for the assessment of image fidelity," in *Human Vision: Visual Processing, and Digital Display III*, B. E. Rogowitz, ed., Proc. SPIE **1666**, 2–15 (1992).
10. M. Duval-Destin, "A spatio-temporal complete description of contrast," in *SID'91 Digest of Technical Papers* (Society for Information Display, Playa del Rey, Calif., 1991), Vol. XXII, pp. 615–618.
11. E. Peli, "Test of a model of foveal vision by using simulations," J. Opt. Soc. Am. A **13**, 1131–1138 (1996).
12. E. Peli, L. Arend, G. Young, and R. Goldstein, "Contrast sensitivity to patch stimuli: effects of spatial bandwidth and temporal presentation," Spatial Vision **7**, 1–14 (1993).
13. E. Peli, "Simulating normal and low vision," in *Vision Models for Target Detection and Recognition*, E. Peli, ed. (World Scientific, Singapore, 1995), Chap. 3, pp. 63–87.
14. B. R. Stephens and M. S. Banks, "The development of contrast constancy," J. Exp. Child. Psychol. **40**, 528–547 (1985).
15. N. Brady and D. J. Field, "What's constant in contrast constancy? The effects of scaling on the perceived contrast of bandpass patterns," Vision Res. **35**, 739–756 (1995).
16. D. J. Field, "Relations between the statistics of natural images and the response properties of cortical cells," J. Opt. Soc. Am. A **4**, 2379–2394 (1987).
17. D. J. Tolhurst, Y. Tadmor, and T. Chao, "The amplitude spectra of natural images," Ophthalmic Physiol. Opt. **12**, 229–232 (1992).
18. D. L. Ruderman and W. Bialeck, "Statistics of natural images: scaling in the woods," Phys. Rev. Lett. **73**, 814–817 (1994).
19. Y. Tadmor and D. J. Tolhurst, "Discrimination of changes in the second-order statistics of natural and synthetic images," Vision Res. **34**, 541–554 (1994).
20. D. J. Tolhurst and Y. Tadmor, "Band-limited contrast in natural images explains the detectability of changes in the amplitude spectra," Vision Res. **37**, 3203–3215 (1997).
21. E. Peli, "Display nonlinearity in digital image processing for visual communications," Opt. Eng. **31**, 2374–2382 (1992).
22. These images were originally recorded with standard video cameras designed to display on a nonlinearized CRT. To enable a linear relationship between the displayed luminance levels and the numerical representation of the images, we presented the images using a linearizing (Gamma corrected) lookup table. To maintain the natural appearance and contrast range of the images, the original images were preprocessed to include the measured display Gamma function.[21]
23. In fact, it was the amplitude, not the contrast, of the images that was increased or decreased. This operation, in which the image mean value is subtracted and the remaining values are scaled up or down, is frequently referred to as contrast increase or decrease. As noted by Peli,[8] the changes in contrast are equivalent to changes in amplitude only where the local luminance is equal to the mean luminance. I will use the term contrast changes here to conform to previous usage, recognizing that in many places the differences were small. This distinction has no bearing on the results or the conclusions drawn here. The contrast of an image can be changed by a fixed factor for all frequencies and locations by using a band-by-band amplification within the context of the contrast metric developed in Ref. 8.
24. The CSF was also measured with a staircase procedure. Only the CSF measured with MOA methods was used in the simulation study. For the subjects who were well-trained psychophysics subjects, the results with MOA differed only slightly from the CSF obtained with the staircase procedure. The CSF data and the standard error of the

measurements were similar to data collected for these stimuli with different systems and with adaptive forced-choice procedures.[12] This was not the case for the novice subject. For this subject (JML) the staircase-procedure data was similar to the data from the other observers, but the MOA data showed substantially reduced sensitivity (as much as 0.5 log unit at middle and low frequencies), even when measured repeatedly. It is interesting to note that for this subject the MOA results provided a better prediction of the simulation performance than did the CSF obtained with the staircase procedure.

25. A. B. Watson, "The cortex transform: rapid computation of simulated neural images," Comput. Vision Graph. Image Process. **39**, 311–327 (1987).
26. H. R. Wilson, "Quantitative models for pattern detection and discrimination," in *Vision Models for Target Detection and Recognition*, E. Peli, ed. (World Scientific, Singapore, 1995), Chap. 1, pp. 3–15.
27. A. J. Ahumada, Jr., "Simplified vision models for image-quality assessment," in *SID'96 Digest of Technical Papers* (Society for Information Display, Santa Ana, Calif., 1996), Vol. XXVII, pp. 397–400.
28. F. W. Campbell and J. G. Robson, "Application of Fourier analysis to the visibility of gratings," J. Physiol. (London) **203**, 223–235 (1968).
29. M. A. Garcia-Perez and V. Sierra-Vazquez, "Visual processing in the joint spatial/spatial-frequency domain," in *Vision Models for Target Detection*, E. Peli, ed. (World Scientific, Singapore, 1995), Chap. 2, pp. 16–62.
30. L. Hainline, J. de Bie, I. Abramov, and C. Camenzuli, "Eye movement voting: a new technique for deriving spatial contrast sensitivity," Clin. Vision Sci. **1**, 33–44 (1987).
31. E. Peli and M. A. Garcia-Perez, "Artifacts of CRT displays in vision research and other critical applications," in *SID 2000 Digest of Technical Papers*, J. Morreale, ed. (Society for Information Display, San Jose, Calif., 2000), Vol. XXXI, pp. 396–399.
32. J. B. Mulligan and L. S. Stone, "Halftoning method for the generation of motion stimuli," J. Opt. Soc. Am. A **6**, 1217–1227 (1989).