

# Test of a model of foveal vision by using simulations

Eli Peli

*Schepens Eye Research Institute, Harvard Medical School, Boston, Massachusetts 02114, and  
New England Eye Center, Tufts University School of Medicine, Boston, Massachusetts 02111*

Received July 20, 1995; revised manuscript received December 8, 1995; accepted December 19, 1995

The appearance of four different images from three different distances was simulated by using the individual contrast sensitivity functions (CSF's) of normally sighted observers. The simulations were generated by using the observers' CSF's as a threshold in a pyramidal vision model of band-limited local contrast [J. Opt. Soc. Am. A **7**, 2030 (1990)]. Simulations based on CSF's obtained in an orientation discrimination task underestimated the observer's sensitivity in discriminating the images. Simulations based on CSF's obtained in a detection task provided a good estimate of observer's performance. The testing method was shown to be sensitive enough to be affected by the high-frequency residual, which is frequently ignored in visual models and simulations. An image-dependence effect found when the high-frequency residual was present was eliminated when the residual artifact was removed. The simulations based on the pyramidal vision model accurately predicted the distance at which they were discriminated from the original image, and thus this model may also serve as the basis for image-quality metrics. The testing method developed can also be used to determine the type of CSF that best represents observer performance in a task. © 1996 Optical Society of America

## 1. INTRODUCTION

One frequent application of visual models has been the generation of simulations. 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 an effort to illustrate the effects of visual disability,<sup>1,2</sup> changes in observation distance,<sup>3</sup> and the perception of images falling on the peripheral retina.<sup>4</sup> More recently, a model of an observer's visual system has been integrated with the physical simulation of new display systems in a computer-aided design tool.<sup>5,6</sup> This combination should permit visual effects of the display perceived by the observer to be considered as part of the design of novel display and in the design of complicated display systems.

A family of visual models characterized by a pyramidal structure of bandpass-filtered versions of the image have been used in basic research and various applications.<sup>7</sup> One recent model of spatial vision was developed that was used to calculate local band-limited contrast in complex images.<sup>8</sup> This contrast measure, together with observers' contrast sensitivity functions (CSF's), expressed as thresholds, was used to simulate the appearance of images to observers, with many of the nonlinearities inherent in the visual system taken into account. With CSF's of low-vision patients, the simulation was applied to the design of image-enhancement algorithms for the visually impaired.<sup>9</sup> Further studies demonstrated that the enhancement derived with the simulation was similar to that chosen by low-vision patients when they were allowed to set the enhancement parameters themselves.<sup>10</sup> The local band-limited contrast model was also used to simulate the appearance of images presented to the peripheral retina<sup>11</sup> with use of the CSF measured at various retinal eccentricities. A similar model has been used for

simulations by Duval-Destin.<sup>12</sup> Others have applied the same concept of local band-limited contrast with small variations<sup>6,12,13</sup> and found it useful in comparing image quality<sup>13</sup> and in other applications of visual models.<sup>6</sup>

The current study was designed to test the validity of the central (foveal) visual model by using simulations of the appearance of complex images generated with the model. This study also sought to determine what CSF data should be used in this or any other vision model of this type. Peli *et al.*<sup>14</sup> have shown that CSF data collected with different stimuli (i.e., fixed aperture versus fixed number of cycles) may differ substantially. Although methodological changes can account for the variability of CSF data in the literature, we do not yet know which is the appropriate method for determining CSF's that represent image perception in the context of pyramidal multiscale vision models.

## 2. EVALUATING THE MODEL BY USING SIMULATIONS

Because the CSF data are given in terms of cycles per degree (c/deg) and the local band-limited contrast model is implemented in terms of cycles per image (c/image), the size of the image in degrees or the observation distance must be defined in the simulation process. To evaluate the validity of the simulations generated by using the model, one can view the simulated image and the original side by side from different viewing distances. The original image should be indistinguishable from the simulated image when both are viewed from a distance farther than that assumed in computing the simulation. However, as the observer moves closer to the two images, the simulation should be easier to discriminate from the original (the reasons for these assertions are explained in the next paragraph). Within the context of a nonlinear-threshold

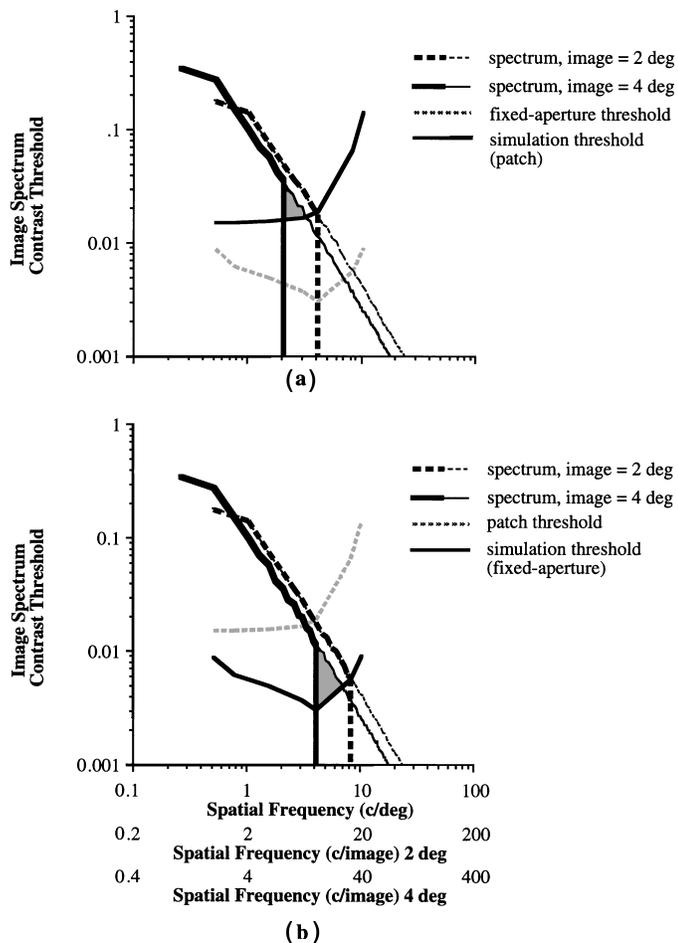


Fig. 1. Relationships among spatial frequency spectra of images and contrast thresholds. Spatial frequency is expressed in cycles per degree and cycles per image for different image sizes. Thin solid-curve spectrum, 2-deg image; thin dashed-curve spectrum, 4-deg image. (a) Simulation with patch CSF, (b) simulation with fixed-aperture CSF. Medium-thick solid curve, CSF used for simulation. Contrast below that curve is below the simulated subject's contrast threshold. Therefore I removed image components to the right of the point where the threshold curve intersects the 2-deg image spectrum (thick dashed curve). At the 2-deg distance the removed components are below threshold, and thus the original image and the simulation should appear identical. When both are moved to the 4-deg distance a portion of the removed components (shaded area) will be above threshold and be visible if the CSF used for the simulation is an accurate description of the viewer's visual system. When tested, viewers can see the difference at 4 deg between the patch simulation and the original, indicating that the viewer's threshold curve lies below the shaded area of Fig. 1(a) and above the shaded area of Fig. 1(b). Note the shift of the spectrum under change of observation distance (see text for explanation).

vision model, such an evaluation process avoids the difficulty of the double-pass problem frequently cited<sup>15</sup> as a limitation on evaluating simulations created with a linear model.<sup>16</sup>

Figure 1 depicts schematically the relationship between the observer's measured contrast thresholds and the amplitude spectra of an image at two viewing distances, with two different CSF's. The patch-CSF data were obtained by using 1-octave Gabor patches as stimuli, and the fixed-aperture CSF data were obtained with a 4-deg, fixed aperture of sinusoidal gratings.<sup>14</sup> The radially averaged amplitude spectrum of the image represents the contrast

at each frequency. The same analysis applies to the simulations generated both with the patch CSF [Fig. 1(a)] and with the fixed-aperture CSF [Fig. 1(b)]. The dashed curves in each panel represent the radially averaged amplitude spectra of the real (thin-dashed curve) and the simulated (thick-dashed curve) images when they subtend 2 deg on the retina. The solid curves represent the spectra of the same images when they subtend 4 deg. Note that the spectrum of the smaller image is shifted by a factor of 2 up along the frequency axis as well as down along the contrast axis. This relationship, pointed out recently by Brady and Field,<sup>17</sup> indeed represents the correct effect of image minification (or change in observation distance) on the spectrum. Previous similar treatments in the literature erroneously asserted that the spectrum is simply shifted along the spatial frequency axis.<sup>16,18</sup> It should be noted that under this transformation the image spectrum is shifted along a diagonal line with slope of  $-1$  on the log-log graph. Most natural images have been shown to have a spatial frequency spectrum that can be approximately described as  $1/f$ , where  $f$  is spatial frequency.<sup>19,20</sup> Such an image spectrum under minification or magnification transformation will be shifted along a line parallel to itself and will be seen to approximately slide along its own with the transformation, as is illustrated in Fig. 1.

Any information in the 2-deg image that falls below the observer's threshold (i.e., to the right of 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 illustrate this, the simulation should (as is shown) remove all that information (thick dashed curves). If the original and the simulated images are viewed from the simulated distance or farther (subtending 2 deg or less), they should be indistinguishable, because the same information from the original that would be lacking, owing to the visual response, was removed in the simulation as well. However, if the original and the simulation are viewed from a closer distance (e.g., subtending 4 deg), the difference in content between the original and the simulation (shaded area) should be visible.

Figure 1 is useful only to illustrate the logic of the experiments described below. The analysis it represents cannot replace the information we seek from the simulations and from direct testing. The effects of contrast threshold on apparent contrast in the images are local, not global. Thus the effective contrast is not accurately represented by the radially averaged amplitude spectrum, because in the simulations, we were working with local contrast, not amplitude,<sup>8</sup> and the simulation algorithm is not represented accurately by the linear filtering depicted in the schematic.

The experiment represented in Fig. 1 was implemented with CSF data averaged from 14 observers. Pilot observations of such simulations<sup>21</sup> led to the conclusions illustrated in Fig. 1: The loss of detail depicted by the shaded area was visible in the case of the patch-CSF simulation [Fig. 1(a)] but not the fixed-aperture CSF simulation [Fig. 1(b)], indicating that the viewer's actual threshold lies above the shaded area in Fig. 1(b), and below the shaded area in Fig. 1(a). Therefore the patch-CSF simulation was deemed to represent the visibility of detail to a normal observer more closely than the fixed-aperture CSF. Both the model and the CSF data used in the simu-

lation can be detected by formally measuring the observations described above. Such testing is described below.

The simulations used here implemented the local-band-limited contrast measure,<sup>8</sup> in which the bandpass-filtered versions of the images are normalized by the local luminance mean calculated as a low-pass-filtered version of the image at each scale. Although this approach has been adopted by some recent models,<sup>6,13</sup> many previous models have used a simpler measure of band-limited contrast<sup>22</sup>. In these earlier models the bandpass-filtered amplitude image was used as contrast. Although not always explicitly, these models did normalize the amplitude to obtain a contrast measure by dividing the

amplitude either by 128 (arbitrarily as the mean of the 8-bit representation range) or by the mean of the image.<sup>1</sup> Frequently this was done in the image domain by normalizing the whole image before filtering. Peli<sup>8</sup> has demonstrated that the different approaches resulted in substantial differences among the contrast measures and argued that these differences would be important for analyzing and simulating image perception. The effect of the local normalization for the simulations of natural images used here is illustrated by calculating one set of the simulation images, also with use of the global normalization by image mean (Fig. 2, far-right column), in addition to the local normalizations used in the study.



Fig. 2. Illustrations of the appearance of the original (far-left column) and of simulated images for three simulated observation distances as noted: when the image spans 4 deg (second column), for a span of 2 deg (middle column), and for a span of 1 deg (fourth and fifth columns). The photographic and printing process prohibit direct evaluation of the effect; however, it is possible to appreciate that only small changes are effected by all simulations. The simulation for a span of 1 deg, using a model without local normalization by mean luminance (far-right column) is presented for comparison. In this case all bandpassed-filtered versions were normalized by the global image mean. The changes between the local and global simulations are of similar magnitude, as are the differences between the 1- and 2-deg simulations. The original images presented here are without the high-frequency residual (i.e., the originals used in experiment 3).

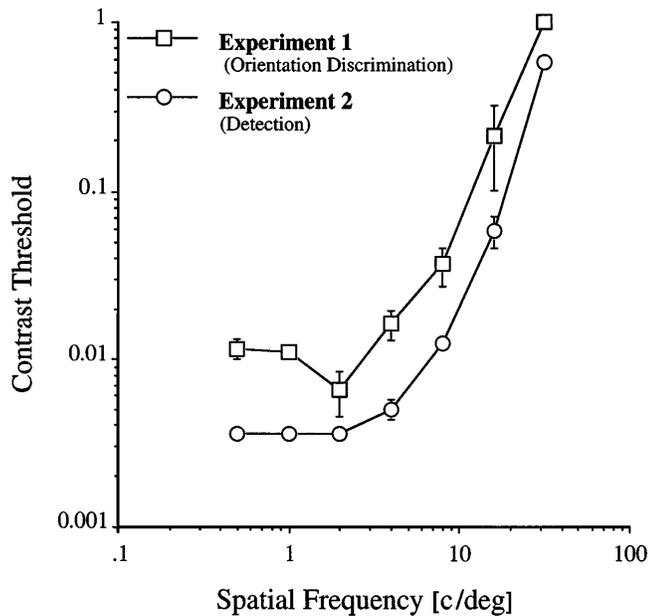


Fig. 3. CSF (contrast threshold) measured with a 1-octave Gabor patch for orientation discrimination and detection. Data shown are means and standard errors of the mean of subjects in experiments 1 and 2.

### 3. METHODS

The simulations were tested by presenting the original image side by side with the simulation. 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. The two images should be progressively easier to distinguish at distances shorter than the simulated distance.

Observers viewed image pairs from various distances and were asked to make a forced-choice distinction between the simulated and the original image. The simulated images used to test each observer were calculated by using her or his CSF. Four different images were used in this experiment. For each image, three simulated views were generated representing views from three different distances (Fig. 2). For the three simulated observation distances [40, 80, and 160 in. (100, 200, and 400 cm, respectively) at approximately 1, 2, and 4 m], the images spanned visual angles of 4, 2, and 1 deg, respectively. The simulated distance and the corresponding span in degrees serves to establish the proper relations between the subject's CSF expressed in cycles per degree and the image spatial content expressed in terms of cycles per image. The subjects viewed the image pairs from six distances, including shorter [20 in. (50 cm)] than the shortest simulated distance and longer [300 in. (7.52 m)] than the longest simulated distance. Each image at each simulation distance was presented 10 times at each viewing distance for experiments 1 and 2, and 40 times for experiment 3. The position of the simulated image relative to the original (right or left) was randomly selected for each presentation. From each observation distance the percent-correct identification of the processed/simulated image was calculated. The data for each simulated distance were fitted with a Weibull psychometric func-

tion to determine threshold at a 75%-correct level of performance.

The CSF data used in the simulations were obtained for each subject individually. In the first experiment, simulated images were produced with CSF's obtained by using 1-octave Gabor patches, where the observer's task was to discriminate gratings of horizontal orientation from those of vertical orientation.<sup>14</sup> The second experiment used CSF's of each subject obtained with a simple detection task and the same stimuli used in the discrimination task. The detection task yielded higher sensitivity (lower thresholds) than the orientation discrimination task (Fig. 3).

The image pairs were presented on a 19-in. (48-cm), 60-Hz, noninterlaced monochrome video monitor (U.S. Pixel, Framingham, Mass.) with use of an Adage image-display system. The spatial inhomogeneity across the screen was 5% at mean luminance. Linearity of the display response was obtained by using a 10-bit lookup table. The calibrated screen provided a linear response over 3 log units. The images were  $128 \times 128$  pixels each and were presented at the middle of the screen, separated by 128 pixels. The four images used are common images used in image processing. These images were originally recorded with standard video cameras designed to display on a nonlinearized CRT.<sup>23</sup> To make possible a linear relationship between the displayed luminance levels and the numerical representation of the images, we presented the images by using a linearizing (Gamma corrected) lookup table. To maintain the natural appearance and contrast range of the images, we pre-processed the original images to include the measured display Gamma function.<sup>23</sup> Observers were seated in a dark room and adapted to the mean luminance of the display ( $37.5 \text{ cd/m}^2$ ) for 5 min before beginning the experiment. Location of the simulated image (right or left) was indicated by using the right and left buttons on a graphic bit pad. A new pair of images emerged after each response and remained on until the subject responded.

The CSF's for the first experiment were measured on the same display system by using the procedure described by Peli *et al.*<sup>14</sup> A two-alternative spatial forced-choice procedure required the subject to decide whether the displayed patch was horizontal or vertical in orientation. The CSF data for experiments 2 and 3 were collected on a Vision Works system (Durham, N.H.) by using a M21LV-65MAX monitor with DP104 phosphor operating at 60 Hz, noninterlaced. A staircase with two practice reversals and four collected reversals was used for each of the seven interwoven frequencies separated by 1 octave between 0.5 and 32 c/deg. The stimuli were the same Gabor patches of 1-octave bandwidth as in experiment 1 (vertical orientation only). The task in this case, however, was a simple detection task compared with the orientation discrimination task of experiment 1.

### 4. RESULTS

#### A. Experiment 1: Orientation Discrimination CSF

Three observers participated in this experiment, and their results were similar. Data from one subject are shown in Fig. 4. These results demonstrate that the simulations generated by using the CSF obtained with

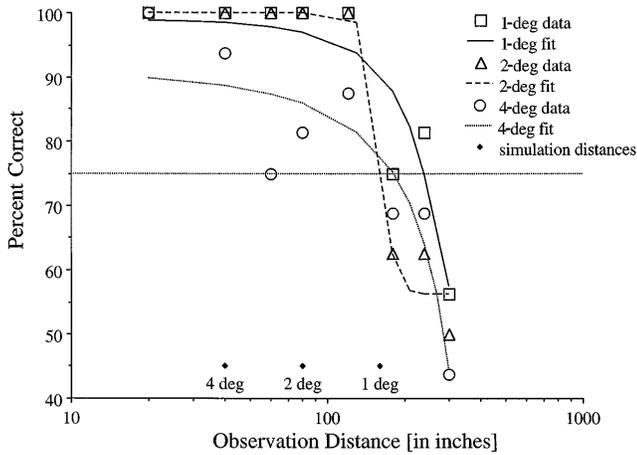


Fig. 4. Results of testing central-vision simulation by using the CSF based on discrimination of the orientation of 1-octave Gabor patches. The data and the psychometric function fits indicate that the subject could distinguish the simulation from the original at distances larger than the distances assumed in the simulations (see diamond inserts at bottom of graph).

a discrimination-of-orientation task of 1-octave stimuli did not support the hypothesis. If the simulations were veridical, the fitted curves should have crossed the 75%-correct level at the simulated distance (marked by diamonds at the bottom of the figure). It is apparent that the observers could detect the changes at distances farther from the screen than the distance assumed in the simulation. These results illustrate that, using this methodology, one can reject the simulations generated by using the orientation discrimination CSF. Furthermore, the fact that the images were distinguished at farther distances indicates that the sensitivity measured by the orientation discrimination task is lower than the sensitivity exhibited by the subjects in the simulation discrimination task.

**B. Experiment 2: Detection Contrast Sensitivity Function**

The results with use of the CSF's obtained in the detection task were closer to the predictions. The results from one subject are illustrated in Fig. 5. The results for all four subjects (one of whom participated in experiment 1) are presented in Table 1. Statistical analysis of these results (*t*-test, *df* = 3) failed to reject the hypothesis that the 75%-correct response occurred at the simulated distances (*p* > 0.18 for both the 1- and 2-deg simulations). The analysis for the 4-deg simulations is not included for reasons explained in experiment 3. Although this result appears to support the validity of the simulations, the statistical power available with only four subjects and one data point per subject per simulated distance (based on 10 repetitions × 4 images × 6 viewing-distance responses) may be too low to provide a definitive answer. More data for each subject are available if the responses for each image are analyzed separately, as discussed below.

**C. Image Independence**

During the experiments two subjects remarked that the originals of the two face images [Lena and high-school girl (HSGR)] were easier to distinguish from the simulations than were the originals of the scenery images [cable car

(CBCR) and Boat]. The data of experiment 2, when we separated out the responses for each image (Fig. 6), appear to support the subjects' observations. For all simu-

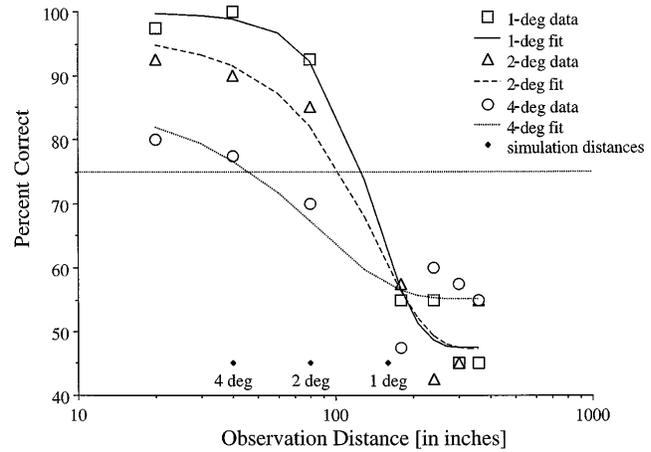


Fig. 5. Results of simulation testing with the CSF based on detection of 1-octave Gabor patches. Here the subject could distinguish the simulation from the original approximately at the distance assumed in the simulations (see diamond inserts at bottom of graph).

**Table 1. Discrimination Threshold Distances (in Inches) Averaged from Four Subjects by Using the Detection CSF's with High-Frequency Residual (Experiment 2)<sup>a</sup>**

Subject	Simulation		
	1-deg	2-deg	4-deg
1	127	105	46
2	233	124	66
3	175	64	20
4	205	151	105
Average	185	111	59
Prediction	160	80	40

<sup>a</sup>In all cases results from all four images were used. Compare the predictions with the averaged results.

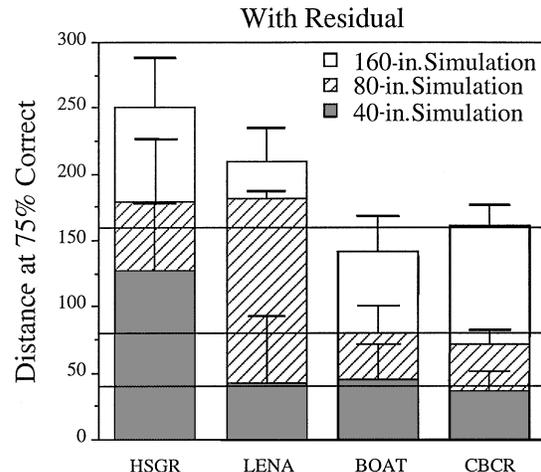


Fig. 6. Distance thresholds for discriminating the simulations from the original calculated separately for the four different images (averaged over the four observers). The results for the 4-deg (40-in.) and 2-deg (80-in.) simulations are superimposed over the results for the 1-deg (160-in.) simulation.

lated distances the face images could be distinguished from the originals at larger distances. A 4 (images)  $\times$  3 (simulated distances) analysis of variance performed on the data of experiment 2 showed a significant main effect of image ( $F = 8.54$ ,  $p = 0.0053$ ) and a significant main effect of simulation distance ( $F = 14.08$ ,  $p = 0.0054$ ), as may be expected. There was no significant interaction between these variables. *Post hoc* analyses found significant differences between the mean for the face images and the mean for scene images combined [ $t(3) = 4.8$ ,  $p = 0.017$ ], in agreement with the observers' impressions. No significant difference was found for any of the other possible groupings [ $t(3) = 1.3$ , not significant (n.s.) for both comparisons].

These results led me to conclude, at first, that some image attribute that is not being considered by this completely subject-centered model might be responsible for the pattern of data.<sup>16</sup> However, the model tested has no place for image dependence except possibly for the effects of masking, which can be implemented in such models explicitly.<sup>6,13</sup> In searching for an explanation for the image-dependence effect, I realized that the simulations as implemented for experiments 1 and 2 removed the so-called high-frequency residual. This is the spatial frequency content in the corners of the spatial frequency domain associated with frequencies outside the circle of maximal radius (64 cycles/picture in this case of  $128 \times 128$  images). This high-frequency residual is usually considered negligible<sup>8,24</sup> and is discarded, since it falls outside the detectable range of frequencies. For the 4-deg simulation [40-in. (1 m) simulated distance] the high-frequency residual contained information at frequencies corresponding to 16–22 c/deg, well within the visible range. Thus the loss of energy at the high-frequency residual for this simulation may have been detectable. To determine the effect of this residual on the previous results, I repeated the above testing with the same subjects and images with only one difference: the original images were processed to remove the high-frequency residual before being used in the testing.

#### D. Experiment 3: Effect of the High-Frequency Residual

When the high-frequency residual was removed from the original images, subjects could no longer distinguish the 4-deg simulations from the original at any of the tested distances (Table 2). This finding demonstrates that in experiment 2 (using the unprocessed originals) observers were able to discriminate the original from the simulation by detecting the residual and not the simulation effect, which was minimal for this image resolution. (The 4-deg simulation was not distinguished in each image even when the data were analyzed separately). Further analysis was carried out, therefore, only for the 1- and 2-deg simulations.

The results of experiment 3 collapsed over all images were only slightly different from the results of experiment 2 for the 1- and 2-deg simulations, with a larger difference for the 2-deg simulation than for the 1-deg case. This outcome was to be expected, since in these cases the residual falls at a higher (probably invisible) range of frequencies (in terms of cycles per degree), and more so for the 1-deg case (64–90 c/deg) and the just-visible frequencies for the 2-deg case (32–44 c/deg). The results for

both of these simulated distances, however, were closer to the predictions than those obtained in experiment 2 (see Table 1). Statistical analyses here, too, fail to reject the hypothesis that the measured discrimination distances were different from the predictions [ $t(3) < 1.0$ , n.s.] for both the 1- and 2-deg simulations.

When the data of experiment 3 (without the high-frequency residual) were analyzed separately for each of the four images (Fig. 7), the image-dependent difference disappeared. Examples of the data for one subject are illustrated in Fig. 8. The results presented in Fig. 7 also showed a good match between the predicted and the measured threshold distance for all four images at both simulation distances. A 4 (images)  $\times$  2 (simulated distances)  $\times$  2 (experiments) analysis of variance revealed significant main effects of image ( $p = 0.006$ ) and simulation distance ( $p = 0.009$ ) and an interaction between the image and the experiment ( $p = 0.013$ ). The interaction verified that the image effect was only in experiment 2 (with the high-frequency residual).

When the data were analyzed separately for each of the four images, the results of experiment 2 (with the high-frequency residual) reject the hypothesis that the predicted distance at which those simulations could be

**Table 2. Discrimination Threshold Distances (in inches) Averaged from Four Subjects by Using the Detection CSF's without High-Frequency Residual (Experiment 3)<sup>a</sup>**

Subject	Simulation		
	1-deg	2-deg	4-deg
1	125	48	17
2	184	50	-627
3	173	116	16
4	212	67	11
Average	174	70	-195
Prediction	160	80	—

<sup>a</sup>In all cases results from all four images were used. Compare the predictions with the averaged results.

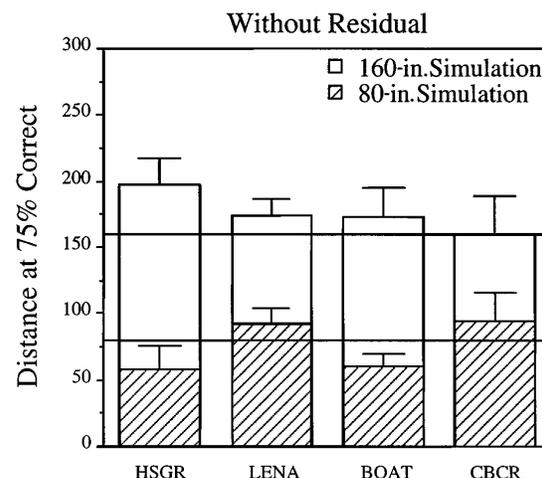


Fig. 7. Distance thresholds for discriminating the simulations from the original calculated separately for the four different images (averaged over the four observers). The results for the 2-deg (80-in.) simulation are superimposed over the results for the 1-deg (160-in.) simulation. Note the excellent agreement with the predictions of the results obtained without the high-frequency residual.

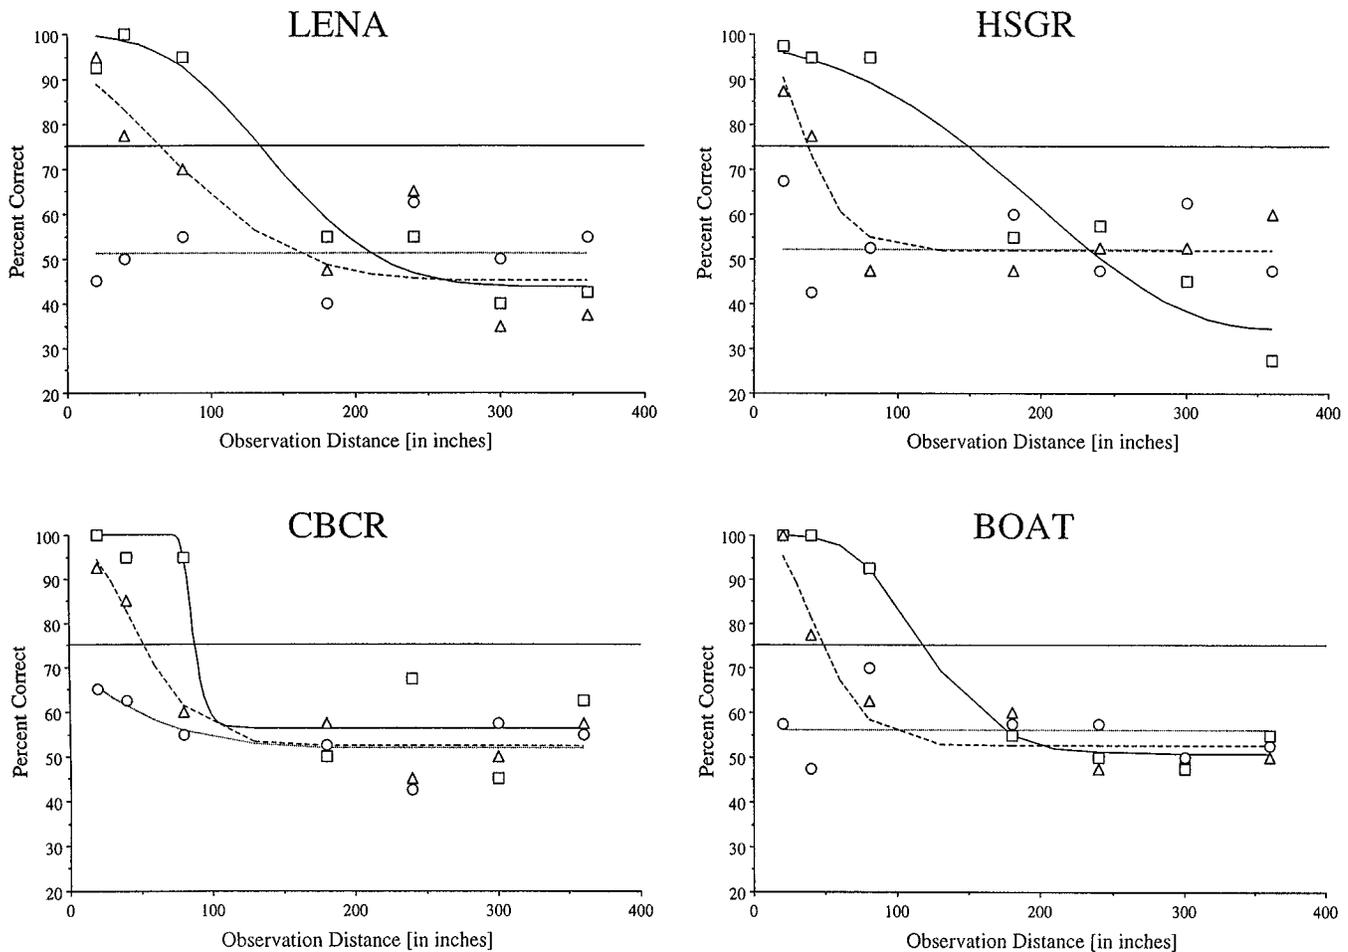


Fig. 8. Example of the discrimination data analyzed separately for each image for the experiment without the high-frequency residual for the same subject as in Fig. 5.

discriminated by the observers was found in the measurements [ $t(31) = -3.095$ ,  $p = 0.004$ ], in contrast to the result obtained with the grouped image data. The same analysis carried out for experiment 3 (without the high-frequency residual) failed to reject the same hypothesis [ $t(31) \leq 1.0$ , n.s.]. This illustrates both that the simulations of experiment 3 are valid and, since the variability in the two cases was similar (see error bars in Figs. 6 and 7), that the statistical power needed to reject that hypothesis is available with the amount of data collected, when each image data is analyzed separately.

## 5. DISCUSSION

The results of these experiments demonstrate that the model proposed by Peli<sup>8</sup> and used to simulate the appearance of an image from different observation distances is valid. The simulated images are distinguishable from the original at distances close to or less than the predicted distances but not farther than the simulated distances. The differences between the images simulated here were quite small (Fig. 2). These sizes of effects, as occur when the observer's distance from the display is doubled, are 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 successfully employed to calculate such differences for estimation of

image quality.<sup>6,13</sup> It is also important to note that the difference between the simulations with the local-band-limited contrast model and those with global normalization are of similar magnitude. Thus simulating with the simplified global model could not result in predictions as good as those demonstrated here.

The method for testing the simulation by using this paradigm is sensitive enough to be affected by the differences between CSF's obtained with different methods and the high-frequency residual. This methodology can be used to determine the type of CSF data that more closely represents the appearance of images. Using this method, I was able to reject my own postulation that the CSF based on discrimination of orientation is a better representation of image perception than is the typical detection CSF.<sup>14</sup> With the same method it may be possible to determine the shape of the CSF directly from simulation experiments by generating the simulation from synthetic, not measured, CSF curves. Such determination is independent of the specific stimuli used for the CSF measurement and may be used to decide whether more than one such stimulus could provide us with CSF's 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 affect perception.

Most similar models used in image-quality assessment have implemented a masking stage<sup>6,13</sup> to account for the

effect of contrast in neighboring channels. The simulation model used here did not include such a stage, yet it performed well in the tests presented. This, however, should not be taken as a contraindication to the use of a masking stage in other applications. In the test used here a masking stage was not necessary. The masking effects of concern are governed by a suprathreshold masker. The suprathreshold information in both the original and the simulated images had to be very similar and thus had similar effects on the appearance of the two images. Implementing masking in the simulation would run into the double-pass problem,<sup>15,16</sup> as the suprathreshold content would mask information first in the simulations and then have the same effect in the observer's visual system. However, the masking effect is needed in calculating image quality and in most other applications of these models.

The image-dependence effect found when the high-frequency residual was present serves as a demonstration of how masking affects the appearance of images. The face images had very little contrast at the high frequencies ( $>16$  c/deg), whereas the scene images had much more content at these high frequencies. Following the simulation, the difference is even larger because the face image's high frequency falls below threshold. Therefore, the high-frequency image content (at frequencies just below the high-frequency residual) in the scene images could mask the high-frequency residual more effectively than that of the face images. This is presumably the reason for the subjects' ability to discriminate the latter at a larger distance.

Other models implementing similar approaches<sup>6,13,22</sup> have used oriented bandpass filters in place of the simple bandpass filters used here. The use of oriented filters is required both on the basis of the current understanding of visual channels and as a necessary step in implementing the quadrature phase filters employed by these models.<sup>25</sup> It should be noted that for the purpose of the simulations as implemented here, the use of oriented filters would have no significant effect because the content in all orientations and bands is added together. Oriented filters are necessary in this context only if the oblique effects,<sup>26</sup> representing different contrast sensitivity for horizontal vertical and oblique gratings, is to be captured by the simulation. Similarly, quadrature phase filters may be needed for the full model implementation, but for the purpose of the simulation they would be needed only if patterns for which these filters would represent a different sensitivity could be found.

## ACKNOWLEDGMENTS

This research was supported in part by grants 05957 and 10285 from the National Eye Institute. I thank Angela Labianca and Elisabeth Fine for valuable technical help and Robert Goldstein and Brian Sperry for programming support.

## REFERENCES

1. B. L. Lundh, G. Derfeldt, S. Nyberg, and G. Lennerstrand, "Picture simulation of contrast sensitivity in organic and functional amblyopia," *Acta Ophthalmol.* **59**, 774–783 (1981).
2. D. Pelli, "What is low vision?" Videotape presentation, Syracuse University, Syracuse, N.Y., 1990.
3. A. P. Ginsburg, "Visual information processing based on spatial filters constrained by biological data," Ph.D. dissertation (Cambridge University, Cambridge, 1978).
4. L. N. Thibos and A. Bradley, "The limits of performance in central and peripheral vision," in *1991 SID International Symposium*, Vol. 22 of 1991 SID Digest of Technical Papers (Society for Information Display, Playa del Rey, Calif., 1991), pp. 301–303.
5. J. Larimer, "Designing tomorrow's displays," *NASA Tech. Briefs* **17**(4), 14–16 (1993).
6. J. Lubin, "A visual discrimination model for imaging system design and evaluation," in *Vision Models for Target Detection and Recognition*, E. Peli, ed. (World Scientific, Singapore, 1995), pp. 245–283.
7. E. Peli, ed., *Vision Models for Target Detection and Recognition* (World Scientific, Singapore, 1995).
8. E. Peli, "Contrast in complex images," *J. Opt. Soc. Am. A* **7**, 2030–2040 (1990).
9. 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. Vis. Sci.* **32**, 2337–2350 (1991).
10. E. Peli, E. Lee, C. L. Trempe, and S. Buzney, "Image enhancement for the visually impaired: the effects of enhancement on face recognition," *J. Opt. Soc. Am. A* **11**, 1929–1939 (1994).
11. E. Peli, J. Yang, and R. Goldstein, "Image invariance with changes in size: the role of peripheral contrast thresholds," *J. Opt. Soc. Am. A* **8**, 1762–1774 (1991).
12. M. Duval-Destin, "A spatio-temporal complete description of contrast," in *1991 SID International Symposium*, Vol. 22 of 1991 SID Digest of Technical Papers (Society for Information Display, Playa del Rey, Calif., 1991), pp. 615–618.
13. 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).
14. E. Peli, L. Arend, G. Young, and R. Goldstein, "Contrast sensitivity to patch stimuli: effects of spatial bandwidth and temporal presentation," *Spatial Vis.* **7**, 1–14 (1993).
15. C. W. Tyler, "Is the illusory triangle physical or imaginary?" *Perception* **6**, 603–604 (1977).
16. E. Peli, "Simulating normal and low vision," in *Vision Models for Target Detection and Recognition*, E. Peli, ed. (World Scientific, Singapore, 1995), pp. 63–87.
17. 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).
18. B. R. Stephens and M. S. Banks, "The development of contrast constancy," *J. Exp. Child Psychol.* **40**, 528–547 (1985).
19. 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).
20. D. J. Tolhurst, Y. Tadmor, and T. Chao, "The amplitude spectra of natural images," *Ophthalmic Physiol. Opt.* **12**, 229–232 (1992).
21. E. Peli, R. B. Goldstein, G. M. Young, and L. E. Arend, "Contrast sensitivity functions for analysis and simulation of visual perception," in *Noninvasive Assessment of the Visual System*, Vol. 3 of 1990 OSA Technical Digest Series (Optical Society of America, Washington, D.C., 1990), pp. 126–129.
22. A. B. Watson, "The cortex transform: rapid computation of simulated neural images," *Comput. Vision, Graphics, Image Process.* **39**, 311–327 (1987).
23. E. Peli, "Display nonlinearity in digital image processing for visual communications," *Opt. Eng.* **31**, 2374–2382 (1992).
24. A. B. Watson, "Efficiency of a model human image code," *J. Opt. Soc. Am. A* **4**, 2401–2417 (1987).
25. E. Peli, "Hilbert transform pairs mechanisms," *Invest. Ophthalmol. Vis. Sci. Suppl.* **30**, 110 (1989).
26. F. W. Campbell, J. J. Kulikowski, and J. Levinson, "The effect of orientation on the visual resolution of gratings," *J. Physiol. (London)* **187**, 427–436 (1966).