Eli Peli, George M. Young, M. Estella Lee, and Clement L. Trempe, "Effects of Distortion Due to Image Enhancement on Face Recognition," in Noninvasive Assessment of the Visual System Technical Digest, 1992 (Optical Society of America, Washington, D.C., 1992), Vol. 1, pp. 18-21.

Effects of Distortion Due to Image Enhancement on Face Recognition

Eli Peli, George M. Young, M. Estella Lee, and Clement L. Trempe

Eye Research Institute, and Department of Ophthalmology, Harvard Medical School, Boston, Massachusetts

The loss of sensitivity at high spatial frequencies and the difficulties that many low-vision patients have in recognizing faces led us to test and demonstrate that (nonlinear) highpass filtering of images may improve recognition of faces (Peli et al, 1991a).

Our initial work was based on a linear model. The finite dynamic range available in the video display and the contamination of the enhanced image by high spatial frequency noise limited the model's usefulness. We proposed a method to address some limitations of the original model by considering the nonlinear response of the visual system (contrast constancy), and requiring enhancement of subthreshold spatial information only (Peli 1991). This modification was designed to increase the dynamic range available by decreasing the range previously used by the linear models to enhance visible details and by employing partial saturation. Results of preliminary testing suggest that only low levels of enhancement are possible without substantial saturation (Table 1). The benefit in dynamic range (reduction in saturation) that can be attained by attenuation of the low frequency is modest and may be effective only at moderate levels of enhancement (Peli 1991). Therefore, we concluded that the enhancement should be optimally tuned to the critical band of frequencies that are just undetected by the patient. For features at these frequencies, a limited level of enhancement may be sufficient to make them visible and thus improve recognition.

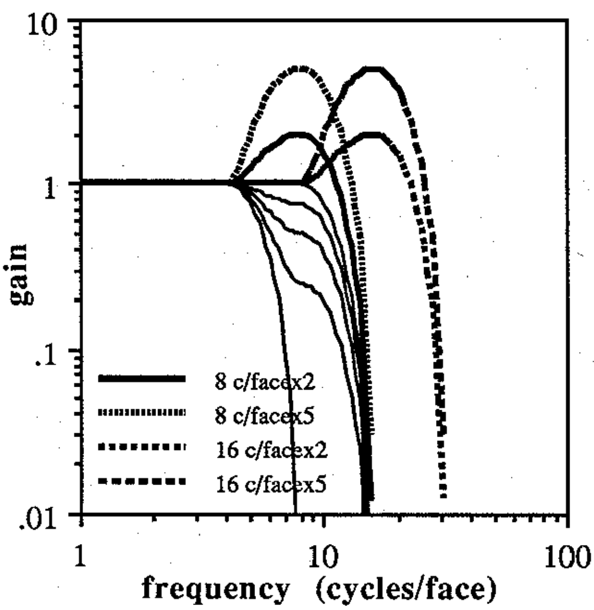

Fig. 1. The five low pass filters used in the previous study (thin lines) compared with the four bandpass filters used in this study to enhance the images.

We recently showed that the spatial frequency content at the band of frequencies above 4 cycles/face was critical for recognizing familiar faces (Peli et al, 1991b). Since face recognition was reduced greatly without this band, and a small fraction of the energy content in that band was sufficient to substantially increase recognition. We concluded that the enhancement of this band probably will aid face recognition for low-vision patients. The present study was carried out to determine whether the distortion of the image resulting from the enhancement and the associated saturations affects face recognition, and how such a reduction is related to the enhancement level. Face images were enhanced by amplifying a 1-octave-wide band of frequencies centered at 8 cycles/face or 16 cycles/face and removing all energy at higher frequencies.

Methods

Subjects. 16 adult subjects with good visual acuity at least in one eye (>20/40) were selected. Most were patients with uniocular age-related maculopathy (ARM) or spouses of patients. The other subjects were volunteers with normal visual acuity in the same age group. The mean age of the subjects was 62 years (range: 24-86 years). Of the total, 13 subjects were older than 50 years. Informed consent for participation in the study was obtained from each subject before testing.

Images. Photographs of 50 celebrities and 40 unfamiliar people were used. The celebrity photographs were expected to be familiar to most subjects. Transparencies of both sets of photographs were digitized at a resolution of 256 x 256 and at 256 gray levels. Illumination was adjusted to obtain good dynamic range and clear visibility of all images. All images were digitized under the same magnification and illumination conditions.

Images were filtered to enhance the bands of 8 cycles/face or 16 cycles/face by a factor of 2 or 5. The face images were filtered using a bank of 1-octave bandpass filters; when added in full magnitude, the filters summed to unity. To obtain the enhancement, the appropriate band was amplified by a factor of 2 or 5, and the higher bands were eliminated (Figs. 1 & 2).

Following filtration in the frequency domain via FFT, the images were transformed back to the space domain. In many cases this filtration resulted in values outside the display range (i.e., higher than 255 or negative values). The filtered images were saturated at both ends of the range rather than scaled back into that range. Rescaling would have reduced the amplitudes at all frequencies by the same amount, while clipping distorted only the information enhanced by the filtering. For two randomly selected face images from our set, we found between 3-50% saturation for most parameters (Table 1).

TABLE 1

The percent of pixels that were saturated or cut off due to enhancement by a 1-octave filter.

Amplification	2	5	10
8 cycles/face	5.5%	31%	50%
16 cycles/face	3.0%	18%	35%

Procedure. The images were presented to the subject sitting in a dimly lighted room on a 60-Hz, non-interlaced video monitor. The image sizes were adjusted to 4° x 4° on the display. Original (unfiltered) and filtered images were presented at random by the computer. The filtered (enhanced) version of the images was presented before the original version. Subjects indicated their level of confidence in recognizing a face as a celebrity on a scale of 1 to 6. A rating of 1 meant that the subject was positive that the face belonged to a celebrity; 2 indicated that the subject was quite sure but not positive that the face was a celebrity; 3 and 4 were used when features were difficult to discern (score of 3 meant that the subject had an inkling that the image was a celebrity; 4 signified that the image was not clear but was judged not to be that of a celebrity); 5 signified that the subject was quite sure that the face was not a celebrity, and 6 meant that the face was clearly visible but not recognizable as a celebrity. An even number of ratings was used to reduce the tendency of subjects to select the midpoint and to force a choice in each case. If subjects could not recognize a particular celebrity from the original, unfiltered image and rated it as 5 or 6, we reclassified that celebrity as a person unfamiliar to these

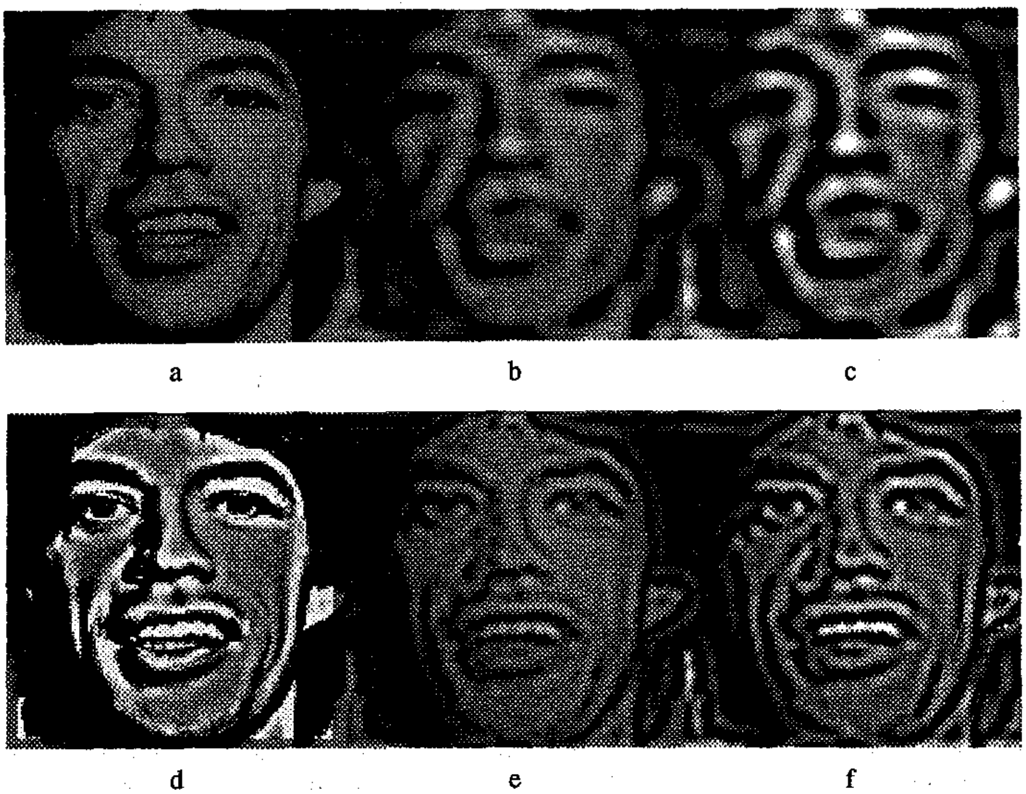

Fig 2. Examples of the images used: a) original unfiltered image, b) image enhanced by a factor of 2 at the 1-octave band of 8 cycles/face, c) image enhanced by a factor of 5 at the band of 8 cycles/face, d) image enhanced with the adaptive enhancement e) image enhanced by a factor of 2 at the band of 16 cycles/face, f) image enhanced by a factor of 5 at the band of 16 cycles/face.

subjects in our analysis of their responses.

These responses were used to calculate receiver operating curves (ROCs), plotting the probabilities of true celebrity vs false celebrity. Separate curves were calculated for original and filtered images. Because the same faces were presented in both forms, the responses for each face were assumed to be correlated requiring a correlated ROC analysis. The area under the ROC (A_z) was taken as a measure of recognition. If filtration reduced face recognition, the area under the ROC for the filtered images should be less than that for the original image. The level of correlation was used in determining the significance of the difference between the two areas.

Results

The performances of subjects in recognizing the unfiltered images varied substantially. Therefore, we normalized the data by calculating the ratio of the area under the ROC obtained from the filtered images to the area under the ROC obtained from the original, unfiltered images. This ratio is presented in Fig. 3 as a function of the 8 cycles/face or 16 cycles/face band used in the enhancement filter. The data also were is compared with the results of the previous study using low pass filtered images (Peli et al, 1991b). The ratio was expected to be close to 1 if filtration did not affect recognition, larger than 1 if filtration improved face recognition, and near zero if filtration substantially reduced recognition performance.

The results of our previous study showed that face images low pass filtered to 4 cycles/face (i.e., the images did not include 8 cycles/face) were difficult to recognize. However, we found that adding even 25% of the magnitude of the 8 cycles/face band significantly increased performance. Including the full magnitude of the 8 cycles/face band improved recognition almost to the level of recognition of the original, unfiltered image.

Increasing the amplitude of the 8 cycles/face band by a factor of 2 (200%) resulted in distortion of the image, which led to a significant reduction in face recognition (Fig. 3). The reduction was even more substantial with the amplification factor of 5 (500%) for the same band. With this level of amplification and the resulting distortions, most subjects performed almost at chance level. These results were somewhat surprising since we recognized most of the enhanced distorted images. The subjects recognized far fewer of the celebrities from the distorted images than from the unprocessed images. However, we were very familiar with the specific set of images under a variety of processing modifications, while the subjects were exposed to those images for the first time.

To date, we completed testing with the 16 cycles/face only for the amplification factor of 2.

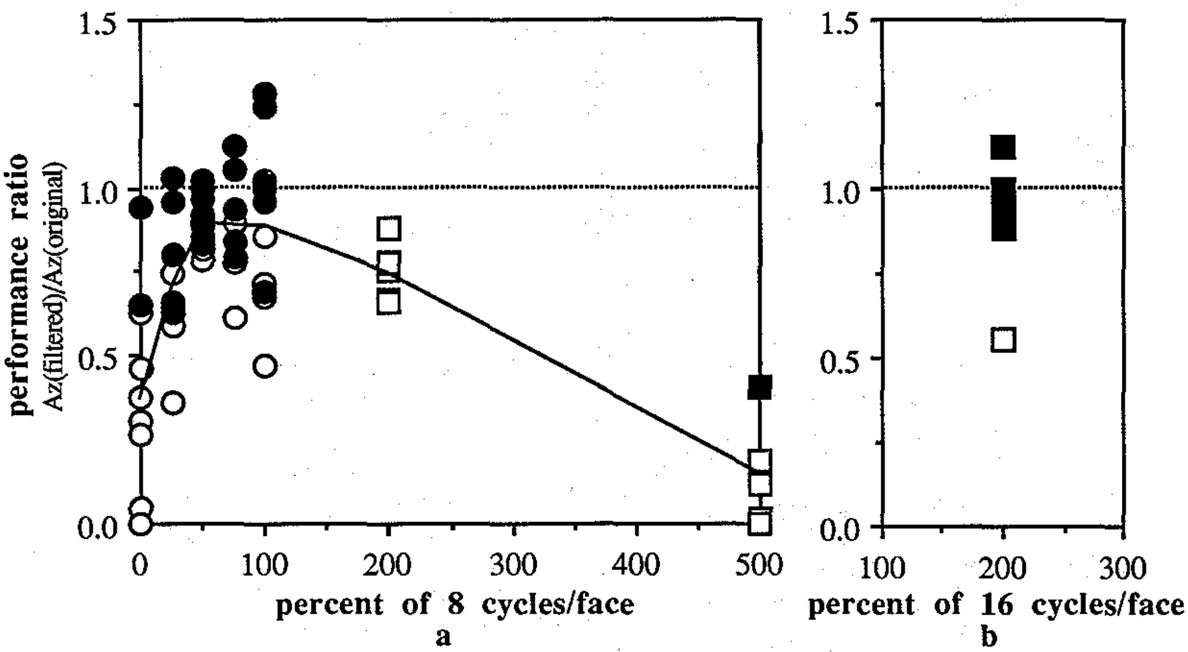

Fig. 3. Degradation in face recognition performance as a function of the percent of the content at (a) 8 cycles/face and (b) 16 cycles/face. The ratio of the area under the receiver operating curve (ROC) for the filtered facial images to the area under the ROC for the original, unfiltered facial images for each subject is represented by one data point. Filled symbols, cases in which areas under both conditions did not differ significantly. Open symbols, cases in which the difference was significant ($p<0.05$). Data from low pass filtering (circles) are compared with results of bandpass-filtering enhancement (squares). The solid line connects the mean values for each condition.

Unlike the case of the band of 8 cycles/face, the enhancement of the higher band resulted in less distortion due to reduced levels of saturation (Table 1). Consequently, the subjects did not show a significant reduction of recognition with this set of enhanced images (Fig. 3b).

Discussion

Using bandpass-filtered face images, Schuchard and Rubin (1989) compared face recognition performance of filtered images with center frequencies at 4.0, 11.3, and 32.0 cycles/face width. They found that the performance of normal observers did not depend on the center frequency of the bandpass-filtered image. Hayes et al (1986), using a paradigm similar to that of Schuchard and Rubin (1989), found decreased performance with decreased center frequency, down to a chance level at 3.2 cycles/face width. However, for higher filtration frequencies their data also did not show substantial decrement of recognition performance. The images in both studies were rescaled after filtration and, therefore, did not suffer from saturation-related distortions.

Our bandpass-filtered images differed from those used in these other studies in three ways, which may account for the different results. Our images retained all the energy at frequencies lower than the enhanced band. This is unlikely to cause a reduction in recognition, as it adds relevant information. Furthermore, similar levels of low-frequency information were available in our adaptive enhanced images that were found to improve patient recognition. Our enhancement filter was narrower (1 octave vs. 2 octaves in the other studies), resulting in more "ringing" artifact, which may have contributed to the distortion and the reduced recognition we found. In a preliminary study evaluating the preferred enhancement by low-vision patients, almost all preferred 2-octave enhanced images over the 1-octave enhancement. Visual inspection of such images shows that the appearance of the 2-octave filtered images is more natural. Our images had a significant level of distortion due to saturation. Masking effects due to the spurious frequencies generated by the saturations could result in reduced recognition. We will test this effect by comparing the recognition of similarly filtered low-contrast images, which should result in minimal saturation.

If image enhancement in the range of 8 cycles/face is needed to improve recognition, the enhancement should not reduce recognition by normal observers. The filters we tested in this study are clearly inadequate for this purpose. The adaptive enhancement images that were found useful in our previous study (Peli et al, 1991a) appear to be similar to images filtered with a 2-octave filter centered at a frequency of 16 cycles/face and with an amplification factor of 5. Enhancement at lower frequencies may be severely restricted by the available dynamic range of the display.

References

T. Hayes, M.C. Morrone, & D.C. Burr (1986). "Recognition of positive and negative band pass-filtered images," Perception **15**, 595-602.

E. Peli, R. B. Goldstein, G. M. Young, C. L. Trempe, & S. M. Buzney (1991a). "Image enhancement for the visually impaired: simulations and experimental results," Invest. Ophthalmol. Vis. Sci. 32, 2337-2350.

E. Peli, G. M. Young, R. B. Goldstein, & C. L. Trempe (1991b). "The critical spatial frequency for face recognition," in Technical Digest on Noninvasive Assessment of the Visual System. Vol 1, pp. 105-108.

E. Peli (1991). "Limitations of image enhancement for the visually impaired," Optom. Vis. Sci. In press.

R.A. Schuchard & G.S. Rubin (1989). "Face identification of band pass filtered faces by low vision observers," Invest. Ophthalmol. Vis. Sci. (ARVO Suppl.) **30**, 396.

Acknowledgements

Supported in part by NEI grant #EY05957 and a grant from the Teubert Charitable Foundation.