

Wideband enhancement of television images for people with visual impairments

Eli Peli

The Schepens Eye Research Institute, Harvard Medical School, Boston, Massachusetts 02114-2500

Jeonghoon Kim

The Schepens Eye Research Institute, Harvard Medical School, Boston, Massachusetts 02114-2500, and Department of Electronic Communication, Shinheung College, #117, Howon dong, Uijeongbu-City, Gyeonggi-do, 480-701, Korea

Yitzhak Yitzhaky

The Schepens Eye Research Institute, Harvard Medical School, Boston, Massachusetts 02114-2500, and Department of Electro-Optics Engineering, Ben Gurion University, Beer Sheva 84105, P.O. Box 653, Israel

Robert B. Goldstein and Russell L. Woods

The Schepens Eye Research Institute, Harvard Medical School, Boston, Massachusetts 02114-2500

Received October 3, 2003; revised manuscript received January 13, 2004; accepted January 23, 2004

Wideband enhancement was implemented by detecting visually relevant edge and bar features in an image to produce a bipolar contour map. The addition of these contours to the original image resulted in increased local contrast of these features and an increase in the spatial bandwidth of the image. Testing with static television images revealed that visually impaired patients ($n = 35$) could distinguish the enhanced images and preferred them over the original images (and degraded images). Most patients preferred a moderate level of wideband enhancement, since they preferred natural-looking images and rejected visible artifacts of the enhancement. Comparison of the enhanced images with the originals revealed that the improvement in the perceived image quality was significant for only 22% of the patients. Possible reasons for the limited increase in perceived image quality are discussed, and improvements are suggested. © 2004 Optical Society of America

OCIS codes: 100.2980, 110.3000, 120.2040, 330.3790, 330.4060, 330.5510.

1. INTRODUCTION

A growing number of people suffer from visual impairments. These impairments and the resulting disabilities greatly affect the quality of life of many older, otherwise healthy people. The rehabilitation needs of people with visual impairments cover a wide scope of activities including reading, face recognition, independent mobility, attending to daily activities, and watching television (TV). Traditionally, vision rehabilitation has been aimed at improving mobility and reading skills.

The incorporation of computerized image enhancement to improve video images for people with visual impairments was first proposed by Peli and Peli.¹ Although image enhancement may be used in portable mobility devices,^{2,3} in the near future we see the main value of image enhancement in providing people with visual impairments with access to the growing volume of video images presented on stationary monitors. TV is an important means of obtaining information and sharing in our culture. Since TV is primarily a visual medium, people with visual impairments do not have full access to it. Yet, most do watch TV with their families and prefer watching TV to other activities.⁴⁻⁶ TV use by people with visual

impairments has increased over the years, and they watch TV nearly as much as, or more than, normally sighted people.^{4,6,7} It is clear that video access will become even more important, serving a wide variety of activities on the Internet. Video imaging has become a major method of obtaining services (including shopping and banking). Access to such services may be even more important to a person who is elderly and visually impaired (who is frequently home bound) than it is to the rest of the population. Image enhancement that would work for TV could serve without further modification in these new applications. Our goal is to develop image-enhancement techniques to assist people with impairments.

We have demonstrated that narrowband contrast enhancement (Adaptive Enhancement¹) of images significantly and substantially increased face recognition for visually impaired patients.⁸ Real-time processing of live color video, using the Adaptive Enhancement algorithm,¹ was made possible with the development of the DigiVision CE-2000 device.⁹ A pilot study using this device found increased recognition of details in the videos and almost uniform (95%) preference for individually selected enhancement.¹⁰ A different study, using a face-recog-

nition task (static images), found that individually selected enhancement improved recognition but not more than uniform enhancement.¹¹ Another more elaborate study of live video enhanced with the DigiVision device, using fixed enhancement parameters and individually selected viewing distance,¹² found a statistically significant improvement in performance, but the effect was small, and only 20% of the subjects in that study indicated a preference for the enhanced images. However, in an additional study,¹³ in which we continuously tracked the perceived quality of motion video viewed with and without enhancement, using both individually selected and modified enhancement parameters, we showed that subjects significantly preferred the enhanced images over the original images. In that study, subjects significantly preferred all of the enhancement parameter choices. The conditions under which the subjects individually selected the enhancement parameters resulted in the largest effect, although it was not significantly greater than with the other enhancement options.¹³

Despite considerable success in these efforts to assist people with central visual field loss (CFL), in which narrowband contrast enhancement was applied to compensate for contrast sensitivity loss, it remains clear that poor image perception in the visually impaired cannot be fully accounted for by the loss of contrast sensitivity and cannot be fully compensated by such filtering. Other aspects of pattern perception, such as phase discrimination and feature localization, specifically with the peripheral retina, need to be considered if we are to more fully understand the visual effects of CFL and to design better image enhancement.

A variety of evidence suggests that the near-peripheral vision used by patients with CFL differs from the central vision afforded by the fovea in ways that cannot be explained by loss of contrast sensitivity alone. A number of psychophysical measures of visual performance scale well with cortical magnification, e.g., two-point separation, Snellen acuity,¹⁴ and grating contrast sensitivity.¹⁵ These functions can be restored to foveal levels by magnification based on the cortical magnification, M .¹⁶ Some other visual functions fall off with eccentricity more rapidly than predicted by M , e.g., vernier acuity,¹⁷ grating orientation sensitivity,¹⁵ pattern symmetry, Landolt-C acuity,¹⁸ and identification of numerals.¹⁹ The addition of contrast scaling can equate foveal and peripheral vision for numerals²⁰ and faces.²¹ The exact nature of these sensory losses in peripheral vision is not well understood.

Phase discrimination of narrowband stimuli has been found to be worse in peripheral vision than in central vision (foveal),²²⁻²⁴ and such loss of phase discrimination was not compensated fully by M .²⁴ Conversely, Morrone and Burr,²⁵ using a wideband stimulus (multiple harmonics), found no difference in the phase-discrimination performance between the fovea and periphery. These findings suggest that performance is improved if the target stimuli are wideband in nature rather than the typical repetitive compound grating pattern, which is only slightly more than 1 octave wide.²³ Phase and localization are frequently loosely equated, but this is true only for pure sine waves.²⁵ It is possible that the reported poor phase discrimination in the periphery actually represented po-

sition uncertainty.²⁵ A similar conclusion was derived by Hess and Hayes,²⁶ who investigated the coding of spatial position. While alignment task accuracy in the fovea had been found to be determined only by the test patch envelope size,²⁷ in the periphery there was an interaction with patch bandwidth.²⁶ Thus with wideband stimuli, localization in the periphery was found to be superior to the performance with narrowband.²⁶

The literature reviewed above suggests that the retinal periphery performs better in these tasks with wideband stimuli than with narrowband stimuli. The enhancement methods that we have used previously resulted in the enhancement of a band of frequencies approximately 1 to 2 octaves wide.¹¹ In fact, that study¹¹ showed that patients clearly preferred the 2-octave-wide enhancement to the 1-octave-wide version. If this is the case, then an enhancement method implementing a wideband approach should be tested.

One approach to wideband enhancement²⁸ that we have tested here involves the addition of "bar" and "edge"-type features^{25,29} to the image. These bar and edge features can be added in ways that enhance the visibility of the original feature and increase the bandwidth of the processed images. To achieve such an enhancement, we created (computed) a line drawing (outline, cartoon, or feature map) of the main visual features in the image. Then we superimposed the line drawing on the original image. Within the framework of edge and bar feature enhancement, various enhancement versions were examined. The basic enhancement method can be described in two main steps. In the first step, the locations of features (outlines of main objects in the image) were produced by using a visual-model-based feature-detection algorithm.³⁰ In the second step the outlines were combined with the image, producing its enhanced version. Such enhancement was suitable for our goal since it used a visual model for the feature detection, and the number of the visible feature outlines in a real TV image was a small fraction of the image size. This should make it possible to transmit such feature information with the TV signal.

Here we report the development of a wideband image-enhancement method³¹ and the wideband algorithm employed in a study that evaluated the perceived benefit to 35 people with visual impairments when they were viewing static TV [National Television System Committee (NTSC)-format] images.

2. METHODS

We start by describing the wideband enhancement method and then the experimental design and the two procedures that evaluated patient preference and patient perceived image quality. Next we describe the TV images, the apparatus for display and data collection, the data analysis, and, finally, the visually impaired patients.

A. Wideband Enhancement Method

The wideband image-enhancement method consists of locating visually relevant features in the image (edges and bars) and enhancing the contrast of the pixels of such features.²⁸ The edge-detection algorithm used here was a dual-polarity edge detector based on a vision model³⁰

(Fig. 1). This algorithm marks “edge” features with dual-polarity pairs of bright and dark lines with the bright line on the bright side of the edge and the dark line on the dark side of the edge. Figures 2(c) and 2(e) show features detected with this algorithm. Thin “bar” features are represented with a single, appropriate polarity line at the location of the bar. The feature outlines detected by the algorithm may be used to enhance the visibility of the fea-

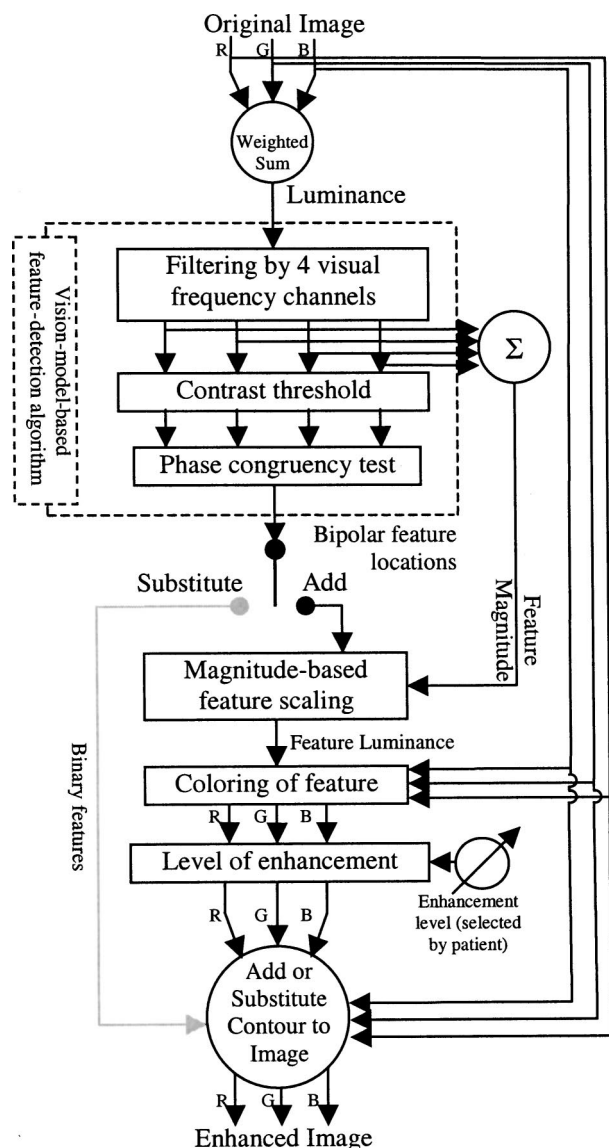


Fig. 1. The image-enhancement algorithm was based on a feature-detection algorithm³⁰ (shown within the dashed lines). The RGB image was converted to luminance, and a bipolar feature representation was generated. An intermediate computation of the bandpass-filtered version of the image, at the feature locations, was used to estimate the magnitude of the underlying feature. This magnitude was used to scale the feature pixel value to be added to (subtracted from, for dark features) the original pixel if the feature was above the threshold of all four filters (phase congruency). The scaled feature values maintained the RGB ratio of the original underlying pixel and thus maintained about the same hue. An individually selected level multiplied the feature pixels' magnitude before adding them to the original image. For binary enhancement (used during pilot experiments), the contours replaced the original pixels (substitution).

tures they underlie in several ways. Bright and dark lines can replace (substitute for) the original pixels' values at their corresponding locations, or they can be added to (subtracted from, for dark lines) the original pixels' values. In both cases the outline magnitudes can be fixed or variable. In color images the outlines can be rendered as black and white (gray) features or they can maintain or approximate the hue (ratios of RGB values) of the underlying original pixels. The simplest application of this approach is to replace each feature pixel with a maximum-contrast pixel of the appropriate polarity (255 for a bright feature pixel and 0 for a dark feature pixel).²⁸

Various versions of this enhancement algorithm were tested in pilot experiments by using small numbers of patients (Appendix A). The visually impaired patients, like normally sighted observers, preferred to see images that appeared as similar as possible to natural images that are enhanced without distortions. Particularly indicated as disturbing were high brightness of partial edges that represented finer (possibly less important) features in the image. Therefore we developed a variant of the enhancement in which colored edges were *added* to the original image. These edges had brightness proportionally scaled to the strength of the features [Figs. 2(d) and 2(f)]. Brightness addition rather than substitution was found to be superior, as it permitted a more uniform contrast enhancement across areas with differing local luminance values. In the case of substitution, the value of bright pixels that afforded some enhancement in a bright section of the image caused too much enhancement in darker sections of the image and vice versa for the values of dark pixels.

Outlines detected in the image were added to the original image at their locations but were scaled in magnitude according to the strength of the feature at the location. The strength of the outline was determined from an intermediate stage of the edge-detection algorithm. At that stage the image was filtered through a multiscale set of visual channels that act as a bank of bandpass filters. The filters were 1 octave wide and separated by 1 octave in their center frequency. A feature was present at a location if the contrast polarity in all four filters was the same (phase congruency step in Fig. 1). The filtered outputs at the feature locations were summed from the four filters to derive an estimate of the feature strength (ranging from -1 to $+1$) at each location. The feature strength was then multiplied by a scale factor (enhancement level) resulting in the final magnitude of the edge or bar pixel to be added at that location. The enhancement level significantly affected feature appearance, ranging from hardly noticeable to levels where the highest screen brightness (255) was assigned to most of the identified bright feature pixels. For the feature strength range of -1 to $+1$, in generating various levels of enhancement, scale multiplication factors of the enhancement process were computed in (arbitrary) steps of 32, starting at 31 (31, 63, ..., 3199). Ten of these 102 scale factors were selected to simulate a continuous change in the image appearance as the level was progressively changed and also so that differences between adjacent levels were noticeable to a normally sighted viewer. Four levels of degraded images³² and the original (unenhanced) image



Fig. 2. (a) Original image. This TV image was particularly sharp and had high signal-to-noise ratio. (b) Degraded image processed with the Adaptive Enhancement algorithm¹ with $K = 0.37$. (c) The bipolar edge and bar features of various strengths detected from the original are shown at scale factor 255 (level 9 in Table 1); they were then added to the original image to create the enhanced image in (d). (e) Shown at scale factor 3199 (level 15 in Table 1) are the bipolar edges of various strengths detected from the original that were then added to the original image to create the enhanced image in (f). The selected median enhancement level 7 from 35 patients in procedure 1 was not clearly visible in print even though it was clearly visible on the TV monitor. The high enhancement level shown in (e) and (f) is equivalent to enhancement done with binary substitution.

were used in addition to these 10 enhanced levels for a total of 15 levels (Table 1). The scale factors used ranged from 63, in which the lines added at the strongest edges

were slightly noticeable for normally sighted viewers, to 3199, where most of the edge pixel values were saturated (to 255 or 0). The setting of 3199 returned us to the re-

sult for the basic, full-substitution wideband enhancement algorithm. The four degraded levels were selected so that they created an almost continuous change in image appearance as the level was sequentially changed.

The color of a feature pixel was a scaled product of the R, G, and B values of the original underlying pixel at that location to maintain the approximate color (hue) of the feature. The R, G, and B values of each feature pixel were added to the original pixel. The R, G, and B values were limited to 255. Thus the color of bright pixels was similar, but not identical, to the original color. In such pixels the color became desaturated, and the hue was somewhat modified as well. Very bright (saturated) pixels were represented as pure white.

To create the degraded images, the Adaptive Enhancement algorithm¹ was applied with the contrast-enhancement parameter, $K < 1.0$. Four sets of degraded images were made with this algorithm by using values of $K = 0.2, 0.37, 0.50,$ and 0.70 and then were used in procedure 1 (Table 1). The other parameters of the algorithm were $L = 2.7$ and window size = 10. The original image is equivalent to the image processed with $K = 1.0$ and, for the wideband enhancement algorithm, with a scale factor equal to zero. In procedure 2 (see below), we used level 2 of Table 1 ($K = 0.37$) as the “degraded images.”

B. General Methods

Two procedures were used to examine patients’ responses to the wideband enhanced images. In the first procedure, for multiple images, patients moved a mouse on a graphics tablet to dynamically adjust the level from a set of 15 levels (Table 1) until they selected the level that appeared to be the best (*preferred*). Using this result, in the second procedure each patient compared images processed by using his or her individually selected level (me-

dian selection in procedure 1) with the original image and rated appearance of the processed images on a scale³³ of *perceived image quality*.

Patients were asked the size of their home TV and how close they usually sat to it. Then they were seated so that the visual angle subtended by our 27-in. TV set approximated their home TV arrangement. For two patients this distance was reduced because the patients could not see the image change as the enhancement level was varied by moving the mouse. The average distance of patients to the TV was 39 ± 15 in., which is much closer than the standard viewing distance of 105 inches.³⁴ In the dimly lit room, illuminated by recessed overhead incandescent lights, illuminance at the monitor surface was approximately 1 ft-candle.

Because many of the patients were elderly and had little or no computer experience, some needed a practice session in the use of the graphics tablet and mouse. This was conducted before the actual study session, with other images.

C. Procedure 1: Preferred Wideband Enhancement Level

In procedure 1, patients actively changed the enhancement level of the displayed image. An image drawn from ten still images (each shown twice) was displayed on the TV screen. By moving the mouse up and down on the blank graphics tablet, patients progressively changed which of the 15 precalculated levels (described in Subsection 2.A) of the image was displayed.

The patients were asked to find the level where they “liked the picture the best, where it was clearest to them, and where they got the most detail from the picture.” Once a patient found an image that looked the best, he or she recorded that setting by clicking on a mouse button. For each trial, the active region of the graphics tablet was randomly shifted vertically so that the patients were unable to associate a fixed mechanical position with their choices. The individually selected enhancement level was determined as the rounded median³⁵ of the patient’s selections from the 20 presentations. For 18 of the patients, procedure 1 was repeated after procedure 2 to assess the consistency of the responses (repeatability).

D. Procedure 2: Comparative Image Quality

The individually selected enhancement level from procedure 1 was used in procedure 2. Four versions of 50 images (a total of 200 images) were shown to patients in a randomized sequence. Shown were 50 images from each set of (1) originals, (2) images processed with the individually selected enhancement level (based on procedure 1), (3) images processed with a second wideband enhancement level, and (4) degraded images, created with the Adaptive Enhancement algorithm.

This second enhancement level was chosen to supply another wideband enhancement that had an enhanced appearance that was significantly different from that of the individually selected level. This second enhancement level tested whether the patients were responding as expected. In particular, since the second enhancement level also contained wideband enhancement, if patients were responding to the presence of the wideband

Table 1. The 15 Levels Used in Procedure 1^a

Level	Scale Factor or K
1	0.2
2	0.37
3	0.50
4	0.70
5	Original
6	63
7	127
8	191
9	255
10	319
11	511
12	767
13	1023
14	1535
15	3199

^aThe four lower levels were images processed with the Adaptive Enhancement algorithm¹ to produce degraded images ($K = 0.2-0.7$). The fifth level was the original image, and the other ten images were processed with the wideband algorithm with the indicated scale factors (63–3199).

enhancement (i.e., detecting the enhancement) rather than reporting image quality, the two wideband enhancements would be expected to be rated similarly. Also, we predicted that the second wideband enhancement level would have a lower perceived image quality than the individually selected level. This second enhancement level was chosen to be several levels above the individually selected enhancement. However, for patients who selected a high level of enhancement this was not possible, and a lower level of enhancement was used for the second enhancement. If the original image (level 5) was selected by the patient in procedure 1, two moderate enhancement levels (7 and 9) were used.

By moving the mouse on the graphics tablet, the patients rated the quality of each test image as compared with the original image. The patients were asked to rate the image as “better,” “slightly better,” “typical,” “slightly worse,” or “worse” than the original images. These words were printed in a large font on the graphics tablet with “worse” at the bottom of the tablet (closer to the patient). Before the computer accepted their rating, the patients were forced to view the original image at least once for comparison by moving the mouse to a designated section of the tablet (right side) marked by a black stripe. Once the original image was viewed, the patients were allowed to rate the test image. Patients were able to view the original image for comparison as many times as desired.

E. Image Acquisition and Categorization

Single video frames (static images) were randomly grabbed from cable TV channels in Boston, Massachusetts, during one day (26 June, 2000) by randomly changing the channel selection and then grabbing whatever image was on the screen. The images were captured by using a Video Toaster (NewTek, San Antonio, Texas). Captured as $480 \times 720 \times 3$ RGB bitmaps, the images were processed in that format by using Matlab (Releases 12 and 13; Math Works, Natick, Massachusetts) programs and then converted to a NewTek proprietary format for presentation on a TV monitor with SpeedRazor (In:Sync, Bethesda, Maryland). Of the 200 digitized images acquired, 127 judged by two normally sighted observers to contain little or no apparent motion due to differences between the two interlaced fields were selected.

It was hypothesized that the wideband enhancement might be of more benefit for some types of images than others. Therefore these randomly selected static TV images were categorized into several subcategories. Four normally sighted observers (aged 25–65 yr) independently categorized the 127 images. The observers were instructed to categorize on the basis of the important information in each image. Important information was considered to be that required for a viewer to understand the major or critical elements in the image. Five major categories were employed: Face, Figure, Text, Busy Scene, and Other. “Other” was used only by two observers, each for one image. Some of these major categories were subcategorized to make a total of ten categories. Text was subdivided into Partial and Full. Face and Figure were subdivided into Single and Multiple. Figure was also subdivided into Real-World (human, animal, and car) or Cartoon. Each image could belong to more than

one category (e.g., it might be necessary to see a face and two automobiles to understand the image). Some images were assigned to two major categories (e.g., Face and Figure). To maintain nonoverlapping groups for data analysis, five observers (two of whom participated in the original categorization) assigned these multiple-category images to one of the major groups on the basis of the main information content.

Of the original 127 images, 19 were not reliably categorized (i.e., no category was indicated by three or more observers). These 19 and 8 arbitrarily chosen others were removed so that the study set consisted of 100 images. These 100 images were divided, and patients viewed only one of the two sets of 50 images.

F. Apparatus

All processing, experiment control, and analysis were done with an Intel-based personal computer running Windows NT 4.0 (Service Pack 6). Images were displayed on a 27-in. (diagonal) Sony Trinitron NTSC-format TV monitor by using the Video Toaster image-processing system under control of programs written in Microsoft Visual Basic and Matlab. In procedure 1, patients moved the mouse position over a 12-in. SummaSketch III (GTO CalComp, Inc., Columbia, Maryland) graphics tablet device to change the image presented. The same tablet was used to rate the images in procedure 2. In both procedures, pressing the mouse button indicated the final decision. A different printed page was placed over the tablet for the two procedures.

G. Data Analysis

Since the image-processing levels used in procedure 1 (Table 1) were ordered but the perceptual intervals were not equal intervals, the rounded median³⁵ preferred enhancement level and the group distributions were used for most analyses.

Data from procedure 2 were analyzed by using a signal-detection approach.³⁶ The Rokit program³⁷ was used to determine the area under the fitted receiver-operating-characteristics (ROC) curve (A_z).³⁸ Paired comparisons were made between responses to the original images and the processed images. As there were three sets of processed images for each patient, three ROC curves were determined (see, e.g., Fig. 6 below) that represented the difference in perceived image quality between the original and that form of image manipulation (processing).

In ROC analysis a detector’s (e.g., patient’s) responses to “noise” presentations and to “noise-plus-signal” presentations are compared. In our study, the original images were treated as the noise presentations, and the processed images were treated as the noise-plus-signal presentations. Patients were asked to report perceived image quality, so that they could be considered image-quality detectors. As can be seen in Fig. 3, our raw data consisted of multiple frequency distributions along the perceived image-quality dimension (for simplicity, Fig. 3 shows data for only three of the four test image sets). When the perceived image quality of the processed images was better than the original images (level-9 image set in Fig. 3), A_z was greater than 0.5 [Fig. 6(a) below]. For the degraded image set, subjects’ perceived image-




Fig. 3. The scores represented by these distributions of the perceived image-quality scores of test images (for the purposes of illustration, bins are 0.5 unit wide, but it is important to note that the ROC analysis does not involve binning). This patient clearly preferred the individually selected (level 9) enhancement (and thus has distributions that were clearly separated). These distributions were used to construct two of the ROC curves shown in Fig. 6(a) below. For simplicity, the second wideband enhancement image set is not shown.

Table 2. Group Characteristics^a

Group	<i>n</i>	Age [yr] Median (range)	Visual Acuity [Log MAR] Median (range)	Documented CFL	Definition
A	35	70.0 (19.2–86.0)	0.89 (0.52–2.00)	27	Patients who completed procedure 1
B	18	68.6 (27.3–86.0)	0.95 (0.74–1.30)	14	Subset of Group A who repeated procedure 1
C	25	69.1 (19.2–86.0)	0.88 (0.52–1.30)	20	Subset of Group A who evaluated an alternate image set in procedure 1
D	23	69.1 (19.2–86.0)	0.86 (0.66–2.00)	17	Subset of Group A who completed procedure 2
E	5	60.7 (19.2–70.0)	0.94 (0.73–1.13)	3	Responders. Subset of Group D who significantly liked the enhanced images

^aAs not all patients finished all parts of the study, different groups were involved in each comparison; *n* is the number of patients in the group. Documented CFL is the number of patients in that group who had central visual field loss. Patients without documented CFL, i.e., without a specific record of scotoma, were considered to have CFL as indicated by their substantially reduced visual acuity and other clinical information.

quality distributions were always worse than those of the original images, resulting in $A_z < 0.5$. As our ROC analysis was of perceived image quality—not of enhancement detection, as might be done in another application—the traditional labels of the axes of the ROC figure (e.g., true-positive rate, or “hit” rate) do not apply directly to our situation. In our analysis, the true-positive-rate dimension was the proportion of the *processed* image set with a higher perceived image quality, whereas the false-positive rate (“false-alarm” rate) dimension was the proportion of the *original* image set with a higher perceived image quality (higher being relative to the criterion used for that point on the ROC curve).

Although the graphics tablet is a continuous response measure, for some patients the responses were multimodal, a consequence of the large-font guide words on the graphics tablet (i.e., many patients did not interpolate well between the five words). The data shown in Fig. 3 have a slight tendency toward this multimodal response pattern. Also, often the response distributions were not normally distributed. Even so, in most cases the Rokit program appeared to give a reasonable fit to our data

(e.g., Fig. 6 below). The Rokit program provides 95% confidence limits for each A_z ,^{37,38} and where appropriate we report these. The confidence intervals were used to determine the significance of the responses of individual patients to a particular type of image processing (i.e., A_z was considered significantly different from 0.5 when the 95% confidence interval did not include 0.5). All variances are reported as standard error of the mean (SEM).

Since the image-processing levels used in procedure 1 (Table 1) were ordinal, nonparametric statistical tests were usually used for these comparisons. A_z data distributions from procedure 2 were found to be approximately normally distributed, so parametric statistical tests were used for these comparisons.

H. Subjects

Most patients had central retinal dysfunction such as from age-related macular degeneration. Patients were at least 18 yr of age, able to follow the instructions in English, and not suffering from a condition—such as arthritis—that would inhibit their ability to control the

computer mouse. Inclusion binocular visual acuity range was 20/60 to 20/2000 (0.48 to 2.0 LogMAR).

Single-letter visual acuity was measured by using a BVAT Model 22-4850 (Mentor O&O, Norwell, Massachusetts). Visual fields were measured by using a Bausch & Lomb Autoplot Tangent Screen (Bausch & Lomb, Rochester, New York) to establish CFL. Visual fields were measured monocularly by using a 6-mm white target at 1 m while patients wore their habitual distance correction (e.g., glasses). Patients always viewed the TV images with both eyes open. Some of the patients did not undergo the visual field tests but had a clear diagnosis of macular lesions accounting for their visual acuity loss and thus were presumed to have CFL as well.

Of the 43 total patients who were referred for the study, 8 did not meet the inclusion criteria. The remaining

($n = 35$) completed procedure 1 (Group A of Table 2). Because of patient clinical schedules and physical condition (e.g., age-related stamina), not all were able to complete both procedure 1 and procedure 2. Some patients were able to repeat procedure 1 a second time following procedure 2. Table 2 shows the characteristics and numbers of patients who completed the various portions of the experiment.

3. RESULTS

Patients' preference for the wideband enhancement was evaluated in three ways: (1) through informal interviews; (2) by their preferred levels of enhancement, measured in procedure 1; and (3) by their image-quality responses, measured in procedure 2. In interviews, most patients reported noticing at least some of the modifications to the images that were presented, and many reported liking some level of the wideband enhancement. In procedure 1, 35 patients (Group A, Table 2) demonstrated that they could detect the wideband enhancement and usually preferred a modest amount of enhancement (none preferred degraded images, levels 1–4). The median preferred level (Fig. 4) was 7 (25% quartile, level 6; 75% quartile, level 8), which was significantly different from the original images (level 5) (Wilcoxon signed-ranks test, $Z_{34} = 4.9$, $p < 0.001$). Although the preferred enhancement level was modest, patients clearly preferred wideband enhancement over the original images and over ones that were intentionally degraded.

Ten patients of Group A (those not in Group C) used an alternate image set in procedure 1. The median level of enhancement selected by these ten patients was 6.5, which did not significantly differ (Mann-Whitney test, $Z_{34} = 1.00$, $p = 0.31$) from the median level of patients in Group C (median = 7, $n = 25$). This illustrates that the preferred level of wideband enhancement was not dependent on a particular image set.

To assess the repeatability of our results, 18 patients (Group B) repeated procedure 1 after completing procedure 2. Figure 5 shows the difference between the median levels selected in the two repetitions of procedure 1 as a function of the median level selected in the first session. Most patients selected the same median level in the two sessions. The median level of their first session was 7, and the median of the repeat session was 6, a difference that was not significant (Wilcoxon signed-ranks test, $Z_{17} = 1.10$, $p = 0.27$). The individual median levels on the two procedure-1 sessions were correlated (Spearman $r = 0.69$, $p = 0.002$). Whether such ordinal data can be analyzed by using parametric statistics is debatable (see, e.g., Barbieto and Simpson³⁹). There was no significant difference between the mean levels chosen on the two repetitions (7.3 and 6.9: t -test, $t_{17} = 1.07$, $p = 0.30$). The repeatability coefficient⁴⁰ was 3.0 units.

In procedure 2, 23 patients (Group D) viewed 50 images with their individually selected levels of enhancement to determine their perceived image quality. As discussed in Subsection 2.G, the results of these measurements were converted to an ROC curve, an associated A_z (area under the ROC curve), and the asymmetric 95% confidence interval (CI) for each processed image set. For example,




Fig. 4. Group A patients shown as two subsets. The patients in Group A who did not complete procedure 2 (Group A–D) selected slightly higher levels of wideband enhancement than the patients in Group D (who completed procedure 2) ($p = 0.02$). Most patients preferred a moderate level of enhancement. Note: No patient preferred any of the degraded levels (levels 1–4).

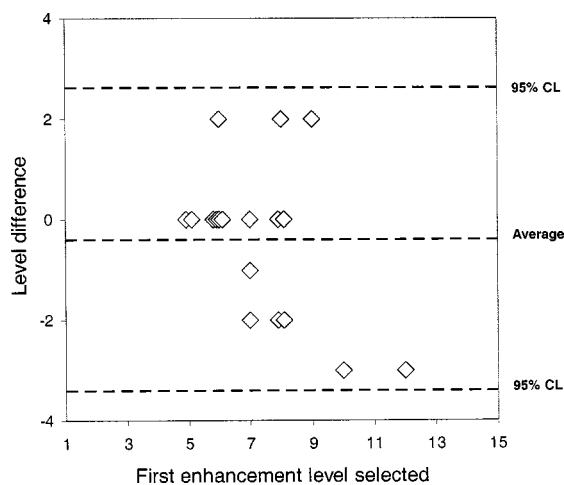

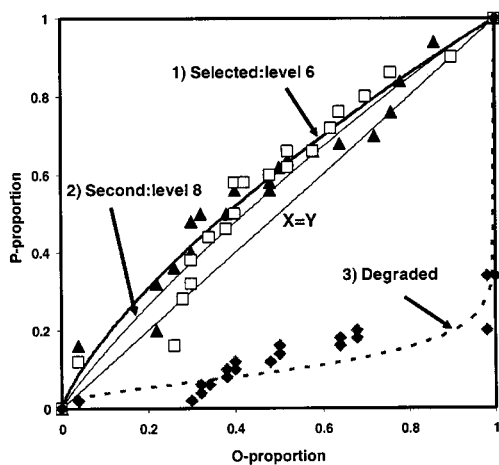


Fig. 5. Difference between enhancement levels selected on the two repetitions of procedure 1 (group B). Half the patients selected the same level. There was a slight tendency to select a lower enhancement level on the repeat, especially if a high level had been selected on the first. Dashed lines show the mean and 95% confidence limits (CL). Note: Overlapped symbols were shifted slightly horizontally to make them visible.



(a)



(b)

Fig. 6. ROC data and fitted curves for two patients. P-proportion is the proportion of the processed images with higher perceived quality, and O-proportion is the proportion of the original images with higher perceived quality. The thick solid curves are the fits to the solid triangular symbols (individually selected enhancement level), and the thin curves are the fits to the open square symbols (second enhancement levels). The dashed lines at the right of (a) and hugging the lower right corner of (b) are the fits to the solid diamond symbols (the degraded images). (a) A 43-yr-old patient (visual acuity 20/250) who clearly favored the wideband enhancement. Here the individually selected enhancement level was 9. The second enhancement level was 7. This patient clearly rejected the degraded image and significantly favored the enhanced images. (b) A more typical example in which A_z was only slightly larger than 0.5. This 69-yr-old patient (visual acuity 20/180) had an individually selected enhancement level of 6. The second enhancement level was 8. The degraded level was clearly rejected. Two ROC curves shown in (a) are constructed from the scores represented by the distributions shown in Fig. 3.

Fig. 6 shows results for two patients. One patient [Fig. 6(a)] clearly favored the enhancement [individually selected level 9, $A_z = 0.86$ (95% CI 0.77–0.92); second enhancement level 7, $A_z = 0.78$ (95% CI 0.68–0.86)]. The preference of the other patient [Fig. 6(b)] for the enhanced levels was not significant [individually selected level 6,

$A_z = 0.58$ (95% CI 0.47–0.69); second enhancement level 8, $A_z = 0.56$ (95% CI 0.44 to 0.66)].

If the image qualities were judged to be not different from the quality of the original images, A_z would equal 0.50. For the 19 patients who had a preferred enhancement greater than level 5 (original), the individually selected wideband enhancement was reported, on average, as having slightly better image quality than the original images ($A_z = 0.57 \pm 0.026$; t -test, $t_{18} = 2.77$, $p = 0.012$). As shown in Fig. 4, the procedure 1 enhancement levels selected by the 23 participants in Group D were significantly lower than those selected by the 12 people in Group A who did not go on to procedure 2 (Mann–Whitney $Z = 2.32$, $p = 0.02$). Although this difference might explain our less-than-optimal results, such explanation is not supported by the fact that perceived image quality (as measured by A_z) was not significantly correlated with the patients' individually selected enhancement level (Spearman $r = 0.06$, $p = 0.81$). Patients in Group D were not significantly younger (Mann–Whitney $Z = 1.55$, $p = 0.12$), did not have better visual acuity (Mann–Whitney $Z = 1.48$, $p = 0.14$), were no less likely to have documented CFL (Fisher exact test, $p = 0.30$), and were no more likely to be female (Fisher exact test, $p = 0.13$) than the patients who did not complete procedure 2. For the individually selected enhancement level, 5 of the 23 Group D patients (22%) had, for their individually selected enhancement level, an A_z significantly greater than 0.5 (i.e., $A_z \geq 0.68$, with the 95% confidence interval excluding 0.5: Group E). Three other patients approached this level of significance with the lower bound of their confidence interval between 0.47 and 0.50 but were not included in Group E. Only one patient from Group D had an A_z that was significantly below 0.5 for his or her chosen enhancement level. The five patients in Group E (shown in Fig. 7) were not significantly younger (Mann–Whitney $Z = 1.57$, $p = 0.12$), did not have worse visual acuity (Mann–Whitney $Z = 0.60$, $p = 0.55$), were no less likely to have documented CFL (Fisher exact test, $p = 0.20$), and were no less likely to be female (Fisher exact test, $p = 0.16$) than the other 18 patients in Group D.

Although the patients indicated a preference for a particular wideband enhancement in procedure 1, most did

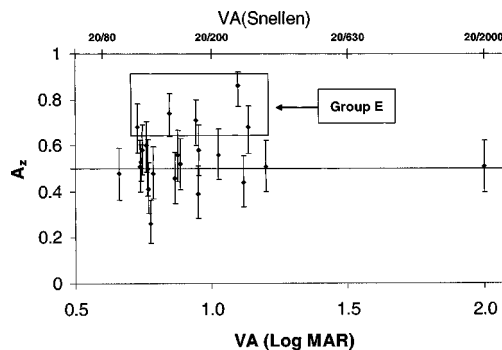


Fig. 7. In procedure 2 (Group D), perceived image quality with the individually selected enhancement, as measured by using A_z , was not correlated with visual acuity. Error bars show the asymmetric 95% confidence intervals.³⁷ For five patients (box), the lower bound of the A_z confidence interval was greater than 0.5, and those patients were grouped for additional analyses (Group E).

Table 3. Average Face Width in 44 Images of 4 Subcategories^a

Category	Number of Images	Face Width
Single face	19	8.6°
Multiple faces	9	4.4°
Single figure	7	4.1°
Multiple figures	9	2.5°

^a Categories such as "Text" that had only a small number of images or did not have faces were not included). Face width was the ear-to-ear visual angle computed for the average observation distance of 38 in. For each image the face widths of all faces in that image were first averaged. These average face widths were then averaged to obtain the reported value.




Fig. 8. Average A_z for four image categories from the patients in Group D and Group E. For Group D, while all show a mean A_z more than 0.5 (dotted line), the multiple-face category had the highest perceived image quality, and it was the only one that was significantly different from 0.5 (original). For Group E, the five patients showed A_z values significantly higher than 0.5 for all four image categories. The error bars represent SEM.

not find the quality of individually selected enhancement images to be much better than the original images. In procedure 2, patients also viewed a second wideband enhancement along with intentionally degraded images. These control conditions allowed us to investigate the validity of our psychophysical method. If our method was flawed (i.e., failed to find a real difference), we might expect that the patients would not report a difference in image quality for the other two image sets. The patients in Group D did indicate that *compared with the original images*, the second wideband enhancement set had slightly, but not significantly, worse image quality ($A_z = 0.44 \pm 0.044$; t -test, $t_{22} = 1.37$, $p = 0.18$) and that the degraded images had much worse image quality ($A_z = 0.13 \pm 0.023$; t -test, $t_{22} = 16.5$, $p < 0.0001$). Also, the second wideband enhancement images had worse perceived image quality than the individually selected enhancement images (paired t -test, $t_{22} = 3.52$, $p = 0.002$).

Preference for wideband enhancements might be related to the patient's visual acuity. In particular, we expected that patients with worse visual acuity would prefer higher levels of enhancement, both as a stronger compensation for visual loss and because they would be less likely to notice the artifacts associated with the enhancement. For procedure 1 there was no significant correlation between the median preferred enhancement level and visual acuity (Spearman $r = 0.04$, $p = 0.82$). Simi-

larly, for procedure 2 there was no significant correlation between A_z and visual acuity (Spearman $r = 0.07$, $p = 0.75$) (Fig. 7). There was a modest negative correlation between sitting distance and visual acuity ($r = -0.36$, $p = 0.044$). This suggests that patients with worse visual acuity tended to sit closer to the TV and thus reduced or eliminated the effect of the difference in visual acuity and its possible effect on the levels selected in procedure 1.

We expected that preference for wideband enhancement might be related to image content. As an example, an image such as a cartoon that already has very well-defined edges might not have an improved appearance with our wideband enhancement. For procedure 1, 25 patients made their selection from a single image set (Group C). There were small, but not statistically significant, differences between the selected wideband enhancement levels for the ten images in the set (Friedman two-way ANOVA, $\chi^2_9 = 16.2$, $p = 0.063$). There was a tendency for higher settings for two images (a "cartoon" and a "page-of-text" image), and lower settings for one (noisy) single-face image. This tendency appeared to be related to the amount of noise in the original image. The two images that received higher settings had relatively low noise in comparison with the lowest-scored image. The lowest-scored image contained an NTSC TV artifact called "cross color" and substantial noise. This implies that enhancement of noisy images and images containing these artifacts might produce results opposite to those desired.

Similarly, we investigated the possibility that the results of procedure 2 depended on image category (image content). For this analysis we chose only the four subcategories that had the largest number of images (see Table 3). The numbers of images in the four subcategories were not equal, because the images were randomly captured. As shown in Fig. 8, for Group D the enhanced multiple-faces images had the highest preference, average A_z (0.61 ± 0.039). There were small, but not significant, differences among the four image categories (repeated-measures ANOVA, $F_3 = 2.51$, $p = 0.066$). Only the multiple-faces category had a perceived image quality (A_z) significantly better than the original images (t -test, $t_{22} = 2.73$, $p = 0.012$). The five patients in Group E (statistically significantly higher individual A_z for the full set of images) were slightly different from the rest of Group D, as they preferred the wideband enhanced images to the original images for all four image categories (t -test, $t_4 \geq 3.78$, $p \leq 0.02$), and again there were small, but not significant, differences among the four image categories (repeated-measures ANOVA, $F_3 = 3.0$, $p = 0.073$) (Fig. 8).

4. DISCUSSION

The patients participating in our study sat very close to the TV (average 39 in.), as they do at home. This approach provides an increased size of the retinal image of the screen and helps compensate for the patients' reduced contrast sensitivity and reduced visual acuity. Most of these patients reported to us that from such a short distance, they had sufficiently good resolution to follow the

visual aspects of the programs they watched and therefore enjoyed their TV viewing experience. This short viewing distance, we presume, also permitted them to detect artifacts of the enhancement and facilitated distinguishing the degraded images.⁴¹ Many of the patients remarked that they liked TV images to have the appearance of natural images and that this was especially true for images of the human face. Image enhancement of any type necessarily distorts the image.^{42,43} Such distortions at some level, even if they are helpful in resolving some image details, might be expected to be rejected by viewers of TV. In particular, we noted during the pilot studies that patients rejected the wideband enhancement whenever the distortion or the noise was clearly noticeable to them. The concept of enhancement is therefore necessarily limited to moderate levels of enhancement, in which case the distortion may either go unnoticed or actually appear to improve the image as perceived by a patient with visual impairment. Our wideband enhancement was therefore adjusted to permit such moderate levels of enhancement. Indeed, most patients favored images enhanced with only moderate scale factors in procedure 1 (Fig. 4) and generally showed less favor to higher levels of enhancement (often presented as the second enhancement image set in procedure 2).

The magnitude of feature enhancement applied in the study was scaled with the strength of the underlying feature (edge or bar). With this processing, minor features (real or noise) were only slightly enhanced and probably were not visible to the patients. On the other hand, high-contrast features (that might have been visible to the patients even without enhancement) were strongly enhanced, creating a distortion that might have been a cause of some of the rejections. In future evaluations it might be preferable to apply a nonmonotonic rule of enhancement that enhances only moderate-contrast features and leaves the high-contrast features unenhanced or only slightly enhanced. Such enhancement of only "critical" features was proposed for the spatial-frequency range enhanced in narrowband enhancement.⁴³

During the tests, patients frequently remarked that they saw occurrences of edge points in places that edges were not expected. These were perceived as noise and detracted from image quality. This kind of noise is a common result of the processing of broadcast (even cable-quality) video. We experimented with a number of approaches for cleaning this noise following edge detection. This required significant computational effort and was only moderately successful. When we developed this wideband enhancement technique we expected that such noise would occur but hoped that owing to its random nature the noise pixels would be averaged out in a motion video sequence (temporal averaging). If that were the case, the noise problem would have been an artifact only of testing with static images. We have generated a few video sequences that were processed with the wideband enhancement frame by frame (as well as field by field). To our disappointment, we found that the high-contrast single-pixel noise was not sufficiently reduced in such sequences and remained quite disturbing. Given the patients' sensitivity and aversion to that noise, and since temporal averaging of it did not reduce it on the screen or

visually at the 30 frames/s rate, we believe that successful use of wideband enhancement would require elimination or substantial reduction of such noise in the enhanced images. One possible approach is to reduce or eliminate the level of pixel noise by applying a filter to the original image before enhancement, thereby cutting off content at spatial frequencies too high to be seen at any contrast by the patients (e.g., above 8 or 10 cycles/deg.⁸).

That 5 of 23 individual patients (22%) had a statistically significant image-quality improvement for the wideband enhancement was disappointing when compared with 42% (16 of 38) of patients who had a statistically significant improvement in recognizing face images enhanced with the Adaptive Enhancement algorithm in a previous study.⁸ In that earlier study the faces subtended 4 deg. Faces are very important features on TV, and difficulties with face recognition as well as recognition of facial expressions are common complaints of visually impaired patients. Many of the TV images in our set included faces. We note that the whole group of patients showed significant preference for enhancement for the multiple-face category (Fig. 8). As Table 3 shows, the face width in the multiple-face category spanned 4.4 deg at the average viewing distance, whereas many of the faces in the single-face category were twice as large on average. Thus it appears that in both studies the benefit patients derived from enhancement was maximal for faces spanning ~ 4 deg of visual angle, which is what they seem to span in multiple-face TV images at the short viewing distance. Such images are ubiquitous in TV programming, and therefore improvement in the visibility of such faces might be appreciated by visually impaired viewers. These results indicate a possible interaction between the level of enhancement and the spatial-frequency spectrum of the image. The wideband nature of our enhancement was explicitly designed to reduce such dependency, and thus it is surprising to see that such an effect was possibly maintained.

Only 4 of 35 patients preferred the original image to some wideband enhancement in procedure 1, none of them selected degraded images, and few selected high levels of enhancement (Fig. 4). This illustrates that the differences between the enhancement levels were visible and that low-to-moderate enhancement levels were largely preferred. In view of this, it was somewhat surprising that only a few patients judged the perceived image quality in procedure 2 to be improved, even using the individually selected enhancement level from procedure 1. While most patients did find that the wideband enhanced images had a better image quality, for most patients that improvement was modest. The image-quality improvement was statistically significant for only 22% of individual patients. The reasons for this apparent discrepancy between the results of the two test procedures need further investigation. It is possible that our results were affected by the inability of a number of people who preferred high levels of wideband enhancement to complete procedure 2 (Fig. 4). However, there was no significant correlation between the preferred enhancement and the perceived image quality of those same individually selected enhancements. We hypothesize that the apparent difference in the outcomes for the two procedures oc-

curred because the two procedures measured two different perceptual dimensions. Procedure 1 required that patients report the best appearance, whereas procedure 2 required that patients say how much better it was than the original. So, patients could see an improvement due to wideband enhancement, but that improvement was not much better (to most).

Another explanation that might account for the patients' failure to report substantial improvements in perceived image quality when viewing levels of enhancement that they individually selected in procedure 1 is that it might be an expression of the type of adaptation to sharpening (and blur) reported recently. Webster *et al.*⁴⁴ demonstrated that a strong adaptation to image sharpening was achieved by image processing. After observers adapted to a sharp image for just 2 min, they judged other sharp images to be substantially less sharp than they did before adaptation. This effect might account for the low appreciation of enhanced images in procedure 2, a paradigm that was similar to the paradigm applied by Webster *et al.*⁴⁴ They⁴⁴ concluded that the visual responses are continuously calibrated to compensate for variation in sensitivity with spatial scale. If this is the case, then any image enhancement will lead to adaptation, reducing the perceived benefit after a short-term use. Such effect, however, might be counteracted by leaving a part of the image (a frame or a margin) unprocessed, permitting the initial calibration to be maintained.

In summary, we developed a wideband image-enhancement method with the potential to improve the appearance of images for people with visual impairment due to central vision loss. Patients demonstrated a preference for the enhanced images (procedure 1), but the improvement in perceived image quality was limited (procedure 2). It is possible that an alternative wideband algorithm would produce better results, though we did explore the possibilities quite extensively (as detailed in Appendix A). Also, our patients had limited experience with the wideband enhancement (typically less than 2 h), so it is possible that they would note greater benefit if allowed to experience wideband enhanced images for a longer period (e.g., weeks of viewing at home). In that case, the wideband enhanced images might come to be perceived as more natural. However, we suspect that the failure to find a substantial improvement in perceived image quality may be a consequence of adaptation to the enhancement of the sort described by Webster *et al.*⁴⁴ If so, it has implications for all forms of image enhancement.

APPENDIX A: WIDEBAND ENHANCEMENT VARIANTS

1. Pilot Enhancement Experiments

Pilot experiments were conducted to evaluate such variables as number of images, precise instructions to the patients, variants of the enhancement algorithms, and the practice environment. The patients were 11 men and 4 women, with a median age of 73 years. It took a long time for the patients to respond to 100 images presented in the four varieties, and was too fatiguing. Therefore we reduced the number of images to 50.

In an early pilot of the procedure, two patients were told to rate the quality of each image on a scale from better to worse as compared with the usual TV image quality they see at home. That comparison appeared to be too vague since the viewing conditions were not identical, despite the attempt to match the visual angle. The procedure was therefore modified to the comparison with the original image.

In one pilot experiment six patients completed procedure 1 twice. The first time, the patients were asked the usual question regarding "liking" the images. The second time the patients were asked to choose the level in which they maximized the "amount of detail" they were able to see in the image. This was done to test the hypothesis that although the more enhanced images conveyed more information, patients did not like an image quality that was not natural. There was no significant difference between the "image quality" and the "detail seen" question in the enhancement levels chosen (Wilcoxon signed rank test: $Z_5 = 1.09$, $p = 0.28$).

Most of the pilot phase was devoted to testing variants of the enhancement algorithm. These are described in the following sections.

2. Binary Edge Substitution

In the basic application of the wideband enhancement approach, the pixels of edges detected by the edge-detection algorithm³⁰ were *replaced* by a full-brightness binary representation of the features (0 for a dark bar or dark side of edge and 255 for a bright bar). Substitution was accomplished by using the "blue-screen" technique. In this technique, the edge-detected image (on a blue background) was used as the foreground image and the original was used as the background image. Thus the resulting composite was the edge image with the (blue) pixels that were not part of the edges that were replaced by the original image.

Binary edge substitution was considered first because it provided the maximal enhancement by direct application of the wideband enhancement concept.²⁸ Also, because all feature values were fixed, it would allow a very economical way of coding such features for transmission with the video signal. Our patients were easily able to see these very-high-contrast edges, but they complained about the brightness of the light features as being excessive although not about the darkness of the dark line.

3. Low-Brightness Edges

Since only the brightness associated with the bright features was noted as disturbing, we reduced the brightness by dividing the pixel values of the white edges by a factor of 2. However, such lower enhancement also provides less effective enhancement, particularly in brighter areas of the image. Indeed, for these bright areas, substituting a lower brightness contour could produce a locally dark area in the image, which is the reverse of what was desired (though it still provides for enhancement of edge visibility). A total of five patients viewed variants of this enhancement method, and they were less bothered by them than by the binary edge substitution.

We also attempted to reduce the disturbing brightness by removing the accelerating gamma function effect that

is applied to images displayed on TV monitors.⁴⁵ We did this by raising the normalized image pixel values to the power of 1/2.2. The number of “disturbing” white edges was indeed reduced—however, not sufficiently to justify the additional processing and some other artifacts that were noted.

4. Black Edges

Initial anecdotal responses from patients indicated that the bright features and flickering (dynamic noise associated with the edge points presented on an interlaced video display) were annoying and were rejected even if they helped in discerning important edges in the image. We therefore tried several variants of a method in which black lines were substituted into the image instead of the bright lines. This resulted in double-thickness black lines at edges and only dark bars. Morphological transformations were applied to attempt to close the gaps in the found edges. The morphological transformations applied were BRIDGE (a “closing” operation), SHRINK (an “erosion” operation), CLEAN (noise removal), and DILATE.^{46,47} Although the flickering was significantly reduced, patients disliked these images as well.

5. Addition of Edges and Determination of Gradient Strength

The substitution approach (Fig. 1) replaced each detected edge with a single level of enhancement signal. Thus very minor edges in the image and even false edge detections (noise pixels) were represented by the same pixel value. The increased visibility of minor edges was perceived as bothersome noise by the patients. In addition, the apparent brightness of the bright feature pixels was particularly noticeable and bothersome in darker sections of the images where the contrast was very high. Therefore we switched from “substituting for” the features to “adding” feature pixels, scaled with the strength of the underlying edge. With addition, the total brightness was lower at darker image sections and brighter at bright areas. The scaling made the brightness of the added value proportional to the strength of the edge at that location. This resulted in minor edges being less enhanced and thus less visible and bothersome.

6. Colored Edges

In an effort to make the edges appear more natural, we used colored edges rather than white or black edges. The detected features were assigned colored values according to the color of the original pixels at their locations. Colored outlines gave a more natural appearance to the enhanced image. In initial pilot experiments, patients disliked the use of constant color brightness of the outlines, and therefore it was combined with the edge-strength-based brightness. This resulted in the wideband enhancement method used in the study and described in more detail in Subsection 2.A.

ACKNOWLEDGMENTS

This study was supported in part by National Institutes of Health grants EY05957 and EY12890 to E. Peli and by a postdoctoral fellowship program grant from Korea Sci-

ence and Engineering Foundation (KOSEF) to J. Kim. We thank Doris Apfelbaum, Avni Vora, and Bridget Hendricks for help in data collection.

Corresponding author Eli Peli can be reached at The Schepens Eye Research Institute, 20 Staniford Street, Boston, Massachusetts 02114-2500; phone, 1-617-912-2597; fax, 1-617-912-0111; e-mail, eli@vision.eri.harvard.edu.

REFERENCES AND NOTES

1. E. Peli and T. Peli, “Image enhancement for the visually impaired,” *Opt. Eng.* **23**, 47–51 (1984).
2. E. Peli, “Head mounted display as a low vision aid,” in *Proceedings of the Second International Conference on Virtual Reality and Persons with Disabilities* (Center on Disabilities, California State University, Northridge, Calif., 1994), pp. 115–122.
3. E. Peli, “Simple 1-D enhancement for head-mounted low vision aid,” *Visual Impairment Res.* **1**, 3–10 (1999).
4. M. Berkowitz, L. G. Hiatt, P. de Toledo, J. Shapiro, and M. Lurie, *Characteristics, Activities and Needs of People with Limitation in Reading Print* (American Foundation for the Blind, New York, 1979).
5. E. Josephson, *The Social Life of Blind People: a Leisure Activities Study*, Research Series No. 19 (American Foundation for the Blind, New York, 1961).
6. J. Packer and C. Kirchner, “Who’s watching? A profile of the blind and visually impaired audience for television and video” (American Foundation for the Blind, August 25, 1997); retrieved 2003, http://www.afb.org/info_document_view.asp?documentid=1232.
7. B. J. Cronin and S. R. King, “The development of descriptive video service,” *J. Visual Impairment Blindness*, 503–506 (1990); http://www.afb.org/jvib_toc.asp.
8. 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).
9. R. G. Hier, G. W. Schmidt, R. S. Miller, and S. E. DeForest, “Real-time locally adaptive contrast enhancement: a practical key to overcoming display and human-visual-system limitations,” in *1993 SID International Symposium Digest of Technical Papers*, J. Morreale, ed. (Palisades Institute for Research Services, Inc., Seattle, Wash., 1993), pp. 491–494.
10. E. Peli, E. M. Fine, and K. Pisano, “Video enhancement of text and movies for the visually impaired,” in *Low Vision: Research and New Developments in Rehabilitation*, A. C. Kooijman, P. L. Looijestijn, J. A. Welling, and G. J. van der Wildt, eds. (IOS Press, Amsterdam, 1994), pp. 191–198.
11. 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).
12. E. Fine, E. Peli, and N. Brady, “Video enhancement improves performance of persons with moderate visual loss,” in *Proceedings of the International Conference on Low Vision, “Vision 96”* (Organización Nacional de Ciegos Españoles, Madrid, Spain, 1996), pp. 85–92.
13. E. Peli, “Perceived quality of video enhanced for the visually impaired,” in *Vision Science and Its Applications*, Vol. 1 of 1999 OSA Technical Digest Series (Optical Society of America, Washington, D.C., 1999), pp. 46–48.
14. J. Marchant, “Sampling theory of the human visual sense,” *J. Opt. Soc. Am.* **55**, 1291–1296 (1965).
15. E. Peli, J. Yang, and R. B. Goldstein, “Image invariance with changes in size: the role of peripheral contrast thresholds,” *J. Opt. Soc. Am. A* **8**, 1762–1774 (1991).
16. V. Virsu, R. Nasanen, and K. Osmoviita, “The cortical magnification and peripheral vision,” *J. Opt. Soc. Am. A* **4**, 1568–1578 (1987).

17. G. Westheimer, "The spatial grain of the perifoveal visual field," *Vision Res.* **22**, 157–162 (1982).
18. J. Millodot, C. A. Johnson, A. Lamont, and H. W. Leibowitz, "Effects of dioptics on peripheral acuity," *Vision Res.* **15**, 1357–1362 (1975).
19. H. Strasburger, L. O. Harvey, Jr., and I. Rentschler, "Contrast thresholds for identification of numeric characters in direct and eccentric view," *Percept. Psychophys.* **49**, 495–508 (1991).
20. H. Strasburger, I. Rentschler, and J. L. O. Harvey, "Cortical magnification theory fails to predict visual recognition," *Eur. J. Neurosci.* **6**, 1583–1587 (1994).
21. P. Mäkelä, R. Näsänen, J. Rovamo, and D. Melmoth, "Identification of facial images in peripheral vision," *Vision Res.* **41**, 599–610 (2001).
22. P. J. Bennett and M. S. Banks, "Sensitivity loss in odd-symmetric mechanisms and phase anomalies in peripheral vision," *Nature* **326**, 873–876 (1987).
23. P. J. Bennett and M. S. Banks, "The effects of contrast, spatial scale, and orientation on foveal and peripheral phase discrimination," *Vision Res.* **31**, 1759–1786 (1991).
24. C. M. E. Stephenson, A. J. Knapp, and O. J. Braddick, "Discrimination of spatial phase shows a qualitative difference between foveal and peripheral processing," *Vision Res.* **31**, 11315–11326 (1991).
25. M. C. Morrone and D. C. Burr, "Discrimination of spatial phase in central and peripheral vision," *Vision Res.* **29**, 433–445 (1989).
26. R. F. Hess and A. Hayes, "The coding of spatial position by the human visual system: effects of spatial scale and retinal eccentricity," *Vision Res.* **34**, 625–643 (1994).
27. A. Toet and J. J. Koendrink, "Effects of blur and eccentricity on differential displacement discrimination," *Vision Res.* **28**, 535–553 (1988).
28. E. Peli, "Wideband image enhancement for the visually impaired (abstract)," *Invest. Ophthalmol. Visual Sci. Suppl.* **39**, s398 (1998).
29. M. A. Georgeson and T. C. A. Freeman, "Perceived location of bars and edges in one dimensional images: computational models and human vision," *Vision Res.* **37**, 127–142 (1997).
30. E. Peli, "Feature detection algorithm based on a visual system model," *Proc. IEEE* **90**, 78–93 (2002).
31. E. Peli, "Wideband image enhancement," U.S. patent 6,611,618 (August 26, 2003).
32. Blurred images were used (1) to prevent "floor" effects, (2) to determine whether degraded images were visible to low-vision patients, and (3) to test the validity of our psychophysical methods.
33. The graphics tablet collected continuous values, but there were large words on the tablet indicating a rating scale to the patient.
34. M. Ardito, "Preferred viewing distance and display parameters," in *MOSAIC Handbook* (O'Reilly, Sebastopol, Calif., 1996), pp. 165–181.
35. J. Freund and B. Perles, "A new look at quartiles of ungrouped data," *Am. Stat.* **41**, 200–203 (1987).
36. N. A. Macmillan and C. D. Creelman, *Detection Theory: A User's Guide* (Cambridge U. Press, Cambridge, UK, 1991).
37. C. E. Metz, P. Wang, and H. B. Kronman, "A new approach for testing the significance of differences between ROC curves measured from correlated data," in *Proceedings of VIII Conference on Information Processing in Medical Imaging* (Nijhoff, The Hague, 1984), pp. 432–445.
38. J. A. Hanley and B. J. McNeil, "The meaning and use of the area under a receiver operating characteristic (ROC) curve," *Radiology* **143**, 29–36 (1982).
39. R. Barbeito and T. L. Simpson, "Should level of measurement considerations affect the choice of statistic?" *Optom. Vision Sci.* **68**, 236–242 (1991).
40. J. M. Bland and D. G. Altman, "Statistical methods for assessing agreement between two methods of clinical measurement," *Lancet* (8476) **1**, 307–310 (1986).
41. E. Peli, "Test of a model of foveal vision by using simulations," *J. Opt. Soc. Am. A* **13**, 1131–1138 (1996).
42. E. Peli, G. Young, E. Lee, and C. Trempe, "Effects of distortions due to image enhancement on face recognition," in *Noninvasive Assessment of the Visual System*, Vol. 1 of 1992 OSA Technical Digest Series (Optical Society of America, Washington, D.C., 1992), pp. 18–21.
43. E. Peli, "Limitations of image enhancement for the visually impaired," *Optom. Vision Sci.* **69**, 15–24 (1992).
44. M. A. Webster, M. A. Georgeson, and S. M. Webster, "Neural adjustments to image blur," *Nat. Neurosci.* **5**, 839–840 (2002).
45. E. Peli, "Display nonlinearity in digital image processing for visual communications," *Opt. Eng.* **31**, 2374–2382 (1992).
46. R. C. Gonzalez and R. E. Woods, *Digital Image Processing* (Addison-Wesley, Reading, Mass., 1992).
47. J. Russ, *The Image Processing Handbook*, 2nd ed. (CRC Press, Boca Raton, Fla., 1995), p. 674.