

Visualizing visual impairments

M.A. Hogervorst[#] & Wim van Damme^{*}

[#]TNO Human Factors
Kampweg 5
3769 DE Soesterberg

^{*}Sensis
Javastraat 104
6524 MJ Nijmegen

Contents

Visualizing visual impairments	1
Contents	2
Summary	3
1. General introduction	4
2. Fixation independent simulation	5
2.1. Introduction	5
2.2. General method	6
2.3. Main experiment	9
2.4. Contrast reduction	11
2.5. Peripheral vision	14
2.6. Image size	15
2.7. Conclusions & discussion	16
3. Fixation dependent simulation	18
3.1. Introduction	18
3.2. Acuity in the periphery	19
3.3. Perimeter data vs. acuity	21
3.4. Simulating eccentric viewing	23
4. Simulating Glare	24
4.1. Model	24
4.2. Results	25
5. Conclusion	26
Acknowledgements	27
References	27

Summary

Purpose. We developed and evaluated methods to simulate the limitations of visually impaired persons. This should give unimpaired persons insight into the problems faced by persons with low vision. Two types of simulation were developed: i) fixation independent, which shows which elements remain visible when they are inspected (related to object recognition) and ii) fixation dependent, showing which elements remain visible with fixation in a given part of the field (related to search).

Methods. In experiments with impaired and unimpaired subjects we determined the extent to which image could be degraded before the difference with the original was visible (discrimination thresholds). For the development of the fixation independent simulation we tested the band limited contrast model (Peli, 1996, JOSA) as a simulation tool. In a series of three successive steps the amount of just detectable blur was determined, followed by contrast levels for removing medium and low spatial frequency components from the image. We determined the relationship between these thresholds and contrast sensitivity to Landolt-C symbols (orientation discrimination). To develop a fixation dependent simulation we determined Landolt-C acuity of unimpaired subjects in the periphery and performed a pilot experiment to establish the link between perimeter data and peripheral blur threshold.

Results. The main result is that the relationship between the blur threshold and acuity is largely independent of the cause of reduced vision (visual impairment, reduced contrast or eccentric viewing). We did not find a match between the contrast levels for removing SF components from the image and contrast thresholds for Landolt-C patterns. We also developed a simulation of disability glare, based on well established CIE equations. In combination with data from cataract patients the disability glare of these patients can be simulated.

Conclusions. Given the latter finding it is unclear how the band limited contrast model used by Peli and colleagues can be used to accurately model visibility of elements in natural images. Instead, we propose a model in which the local blur adapts to the local contrast. We also show how the relationship between acuity and blur can be used in combination with the visual field simulation developed by Perry & Geisler (2002) to implement the fixation dependent simulation (in real time).

1. General introduction

To initiate research on problems caused by low vision and more importantly, find possible solutions to these problems, several institutions for the visually impaired joined their forces in the 'InZicht' Foundation. The aim of this foundation is to encourage multidisciplinary scientific research on low vision. The current study is part of this research program and focuses on the simulation of visual impairments.

Due to new technologies such as cash dispensers, internet and mobile phones visual information has become more and more important. Meanwhile, a large proportion (about a quarter of a million people in the Netherlands) could be registered as 'blind' or 'partially sighted', and this number is growing every year with the increase in life expectancy. This is due to the fact that most visual impairments are acquired late in life; about 2% are under the age of 16, 10% between 16 and 59, and 88% are over 60 years old. It is important that the needs of the visually impaired are recognised and that their needs are taken in consideration when designing new products, infrastructure, architecture etc.

The idea of visualizing the visual limitations of visually impaired persons originated from the colour deficiency simulator developed by TNO (Walraven & Alferdinck, 1997; Walraven, 2000). This simulator uses knowledge concerning the sensitivities of the human colour receptors and the spectral filtering of the eye lens and macular pigment to simulate severe dichromatic colour vision (1.25 % of total population), anomalous trichromatism, and (less severe forms of) deficient colour vision (3% of the total population). This simulation can be used as a diagnostic design tool, but also provides the means for adjusting the colours to the individual needs of a colour-deficient viewer. Figure 1 shows an example simulation for someone with protanopia.



Figure 1. Colour deficiency simulator simulation. Figure 1a shows the original image and Figure 1b shows the simulation for a person with protanopia ("red-weakness").

A similar simulation tool would be useful for other kinds of visual impairments. To make the simulation as useful as possible we focus on the most common impairments, i.e. macula degeneration, cataract, glaucoma and (diabetic) retinopathy. We also investigated the effect of myopia and visual limitations of unimpaired subjects due to large viewing distance, eccentric viewing and low contrast. One reason for including the second group is that this enables the use of (more easily available) unimpaired subjects. On the other hand, this study also gives insight in the

limitations of unimpaired subjects and shows how sensitivity to simple test patterns relates to the perception of natural images.

The term "simulation" may lead to misunderstanding, since it is ill defined. It may suggest that we seek to visualize how impaired subjects perceive the world. We think that such an undertaking is very difficult, if at all possible. Therefore we will instead try to visualize the *visual limitations* of impaired subjects. The method consists of deforming an image until the difference with the original becomes visible. It comes down to removing visual information that is invisible to the impaired subject. The unimpaired subject that is viewing the simulation is left with the same information, and the simulation therefore gives a good impression of the information that is available to the impaired subject. It is well known that viewing strategies can be and are developed by impaired subjects to make optimal use of the remaining visual information. However, this is not subject of our study.

In contrast to so called "artist impressions" of the perceived scene by impaired subjects we look for a validated way of simulating. This means that we investigate which information can be removed from an image before it can be detected. The degree of image degradation is well defined and reflects the information lost by the impaired subject. As such, this simulation is suitable for evaluating designs, environments and regulations. We are aware of the fact that other simulation methods exist (e.g. the simulating glasses used in Dutch centres of rehabilitation) and believe that they all contribute to improving insight into the perceptual world of persons with low vision (have their own applications).

2. Fixation independent simulation

2.1. Introduction

Peli (1990) devised a simulation method in which a measure of local band-limited contrast is used. In the transformation the image is divided in a range of (one octave wide) frequency bands (by using a Laplacian pyramid, see e.g. Burt & Adelson, 1983). A local contrast measure (akin to Weber Contrast) is derived for each spatial frequency component and point in the image. For each frequency band, the contrast is defined as the ratio of the band pass-filtered image at that frequency to the low pass image filtered to an octave below the same frequency (local mean luminance). Spatial frequency components with contrast lower than the threshold contrast are removed from the image, while those with contrast higher than the threshold remain unaltered. The transformation is based on two findings: i) gratings with contrasts lower than the threshold are not perceived and ii) apparent contrast of supra-threshold gratings is relatively independent of retinal eccentricity and spatial frequency (contrast constancy), while the contrast threshold changes significantly (Cannon, 1985). Peli (1995) argues that these non-linear properties make it more suitable to visualize limitations of the visual system than linear transformation methods (e.g. Ginsburg, 1975). The model relates the image transformation to the Contrast Sensitivity Function (CSF) of an individual subject.

Peli (1996) used this transformation method to simulate the effect of distance. In the simulation CSF's were used based on detection thresholds and (orientation) discrimination thresholds of (one-octave wide) Gabor patches. The threshold distance for discriminating the transformed image from the original coincided with the one predicted by the model based on the CSF for detection of Gabor patches. Still, the possibility exists that in these experiments the subjects based their judgments on part of the CSF curve, e.g. the high frequency cut-off. In a more critical test of the model Peli (2001) varied image contrast (10%-300%). The model based on the CSF for detection gives reasonable predictions of the threshold distance, although considerable differences exist from the simulation distance. When contrast thresholds were used obtained at a single distance threshold distance coincide with simulation distance for intermediate contrast levels, but not for low or high contrast

levels. Better predictions were obtained when a CSF was used based on a combination of curves obtained at various distances.

In the current study we used a similar transformation method to simulate the effect of distance and visual impairment. Instead of starting with a predefined CSF we sought to deduce the appropriate curve from experiments. The transformation can be thought of as a series of steps in which more and more information is removed from the image. We transformed the image in three steps. In the first step the amount of just noticeable blur was determined. This corresponds to the removal of high spatial frequencies from the image. In the second and third steps elements with medium and low spatial frequencies were removed with local contrasts below some threshold contrast. In the experiments the contrast threshold levels for medium and lower spatial frequencies are determined. Importantly, with this method the simulation parameters are deduced from observer experiments and are not fixed in advance. The first step (blur) has the largest impact while further steps refine the transformation method¹.

Like Peli we considered it useful to establish the relationship between sensitivity to standard test patterns and sensitivity to elements in a complex image. We used Landolt-C test patterns instead of Gabor patterns (used by Peli and colleagues) since these are used more regularly in ophthalmologic tests. The idea is that with the results of such tests it should be possible to deduce an appropriate simulation.

2.2. General method

2.2.1. STIMULI

The transformation method is based on the local band limited contrast model developed by Peli (1990). As mentioned above, Peli and colleagues created simulations under the assumption that the contrast thresholds for removing certain SFs from an image coincide with contrast thresholds of Gabor patches (orientation discrimination *or* detection: both were tested; detection was found to result in better predictions). We used a different approach in which the contrast thresholds of the SF-bands were determined independently. This method allowed us to construct the appropriate contrast sensitivity curve from measurements instead of relying on a predefined sensitivity curve.

The image transformation can be regarded as image degradation. The main effect comes from the removal of whole SF-bands. More subtle changes are introduced by removing only low contrast SF-components (with lower SFs).

For each SF-component and location a contrast measure (the local band limited contrast) can be obtained as follows. The image is divided into (one octave wide) spatial frequency (SF) bands by use of the Laplacian Pyramid model (Burt & Adelson, 1983; Simoncelli, 2004). Local contrast in each band is calculated by dividing the SF-band (the fluctuations) by a local average, which is the sum of all bands with lower SFs (i.e. a blurred image containing SFs lower than the SFs of the band).

We degraded the image in three successive steps. In the first step we determined how many SF bands could be removed from the image (which is equivalent to determining which high SFs can be removed). In the second and third steps elements with lower SFs were removed from the image. The threshold image from the previous step was used as the reference image for the next step.

¹ This resembles a local Taylor approximation of a curve.

In this step the image was not divided into separate SF-bands, but was blurred using a blur kernel with width σ . This is equivalent to removing high SF-bands, but the removal is not restricted to an integer number of bands (e.g. the result may be equivalent to the removal of 1.4 bands from the image). The same result can be obtained by first increasing the size of the image, then removing an integer number (e.g. 2) of SF-bands, and decreasing the size to its original value. This method was used in the second and third steps, except that the first (resp. second) remaining SF-bands were transformed before the image was reduced to its original size. In the second step SF-components were removed from the highest remaining SF-band with contrast below a given contrast value. In the experiment we determined the contrast threshold ($CB1$) corresponding to a 75% correct score for discriminating the transformed image from the reference. In the third step SF-components were removed from the second SF-band with contrast below a given contrast value. In the experiments we determined the threshold contrast value $CB2$. In successive steps threshold values of σ , $CB1$ (contrast level high SF-band) and $CB2$ (contrast level medium SF-band) are determined. In the first step the reference image was the original; in the second step the reference image corresponded to the threshold result of the first step (the original blurred by σ); in the third step the reference image corresponded to the threshold result of the second step (i.e. the blurred original in which high SFs with contrast below $CB1$ were removed).

In the first step the amount of just detectable blur is determined. Close inspection of the transformed images using Gaussian blur shows that this method also affects lower SFs (i.e. the average luminance in larger areas). To prevent subjects from using this artefact we used a filter that does not affect the lower SFs as much, with a steeper cut-off in the SF domain. We used a difference-of-Gaussians filter (DOG) with a profile defined by 2 times a Gaussian with width σ minus a Gaussian with width $\sqrt{2}$ times σ ($2 * \text{Gauss}(\sigma) - \text{Gauss}(\sqrt{2} \sigma)$).

The squared values of the pixel values are taken as luminance value (i.e. a gamma of 2 is assumed for the camera), and the transformations are performed on these luminance values. When displaying the images, the transformation from pixel values to luminance values by the monitor was taken into account. At the sides the images gradually faded into the background (see Figure 2).

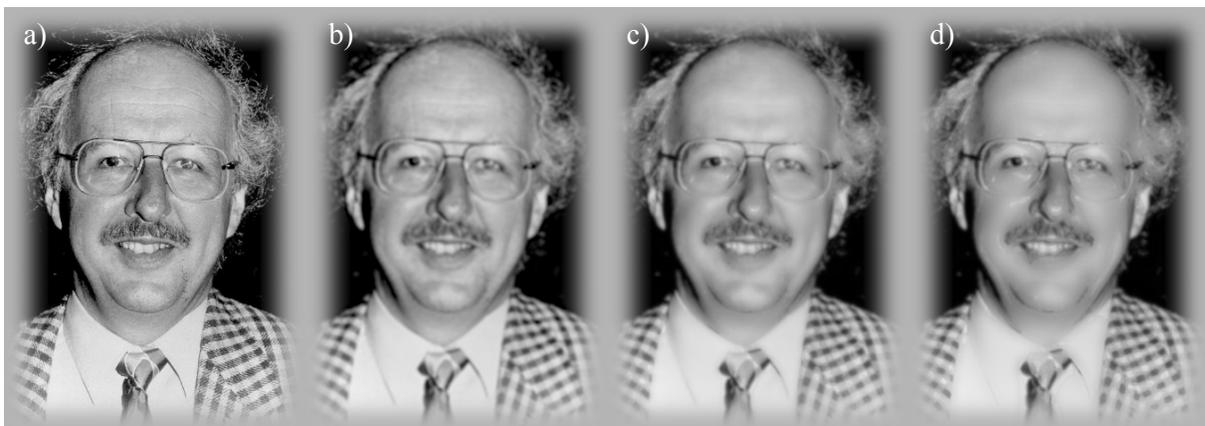


Figure 2. In the main experiment an image was deteriorated more and more with each successive step. In each step the threshold amount of deterioration was determined using a 2AFC method in which the reference and test image were shown side by side. The image shows the original image (a), the result of the first step, i.e. blurring (b), the result of the second step, i.e. removing low contrast medium SF components (c), and the result of the second step, i.e. removing low SF components with low contrast (d). The reference is always the result of the previous step.



Figure 3. The images used in the experiments, referred to as face (a), stairs (b) and forest (c).

Figure 3 shows the images used in the experiments. Various images were tested to determine the effect of image content on the simulation parameters. A face is included since correct interpretation of faces is important for social interaction. The images also differ in the spatial frequency content (e.g. the trees-image contains more high SFs than the other images).

2.2.2. EXPERIMENTAL SETUP

The images were displayed on a computer monitor (Philips Brilliance low emission 2110), size 37.9 x 28.4 cm or 1024 x 768 pixels at a frame rate of 75Hz. Calibrated luminance values were used ranging from 0.51 to 174 cd/m^2 . In front of the CRT monitor (40 cm from the monitor) a slide projection screen was placed with an aperture the size of the display. The projection screen was illuminated by 2 slide projectors, in such a way that a large uniform background around the monitor was created with a luminance that matched the background luminance (medium grey) of the monitor.

2.2.3. PROCEDURE

First the contrast sensitivity was determined of the subject using the method described below. The distance for the main experiment was set at a distance that scaled inversely with the acuity of the subject, such that the distance for subject with an acuity of 1.0 was set at 700 cm (e.g. distance for someone with acuity of 2 is 1400 cm). This (relatively) large distance was chosen such that the pixels could not be resolved by the subject. Also, this assured that the blur threshold was similar for all subjects in terms of pixels, and therefore the number of details that could be discerned in each image was similar for the subjects.

The Quest adaptive staircase method (Watson & Pelli, 1983) with a total of 60 trials was used to determine each threshold, corresponding to a 75% correct score. On each trial the reference and test images appeared side to side, with the order chosen at random. The task to the subject was to indicate the reference image by pressing the mouse button once or twice (reference left or right). The images remained on the monitor until the subject responded. Auditory feedback was given.

The images that were used were pre-computed. In the first step the blur threshold (σ) was determined, and successive σ values differed by 0.06 log units (factor 1.15). In the second and third steps contrast threshold were determined, with contrast levels differing by 0.1 log units (factor 0.79). The measurement program was written in Matlab using the using the Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997).

2.2.4. CONTRAST SENSITIVITY MEASUREMENTS

Contrast sensitivity of each subject was measured using Landolt-C test symbols, using the same setup as used for measuring simulation parameters (sensitivity to changes in natural images, see above). Instead of using Matlab software the stimuli were controlled by a Cambridge stimulus generator (VSGcard 2/3-20MB, maximum luminance of 53 cd/m²), with a frame rate of 75Hz. This system allows luminance values to be specified from a 12 bit resolution range. Again, the stimuli were seen through the whole in a slide projection screen, illuminated by 4 slide projectors, to generate a large uniform background with a luminance that matches that of the CRT.

A contrast sensitivity curve was obtained in three steps. In the first step the minimum angle of resolution (MAR, i.e. 1/acuity) was determined using high contrast Landolt-C patterns (black against a white background), corresponding to the threshold gap size of the Landolt-C symbol. In the next step we determined the contrast threshold at 3.16 times (= sqrt(10)) the MAR, and finally the contrast threshold was determined at 10 times the MAR. Measurements were obtained at similar distances as used for measuring simulation parameters (main experiment) (i.e. at about MAR x 700 cm, see 2.2). At each trial one out of four different Landolt-C symbols was shown with the gap to the right, bottom, left, or up. The subject indicated the orientation of the gap by using the joystick. The test pattern remained visible until the subject responded. Auditory feedback was given. A Quest adaptive staircase method (Watson & Pelli, 1983) was used with 80 trials per staircase to determine the threshold values.

The curve connecting the threshold combinations of (gap size, contrast) represents the contrast sensitivity curve of the subject (in which the contrast value corresponding to acuity was set to 100%).

2.3. Main experiment

The goal of the main experiment was to determine whether and how the simulation parameters (i.e. sensitivity to changes in natural images) can be predicted on the basis of contrast sensitivity (to standard Landolt-C test patterns) of the subjects, and to determine whether this depends on the type of impairment. We measured contrast sensitivity and simulation parameters for various visually impaired persons, as well as unimpaired persons under various conditions (normal, low contrast, eccentric viewing).

2.3.1. SUBJECTS

Ten visually impaired subjects were tested with a range of (combinations of) visual impairments, incl. macula degeneration, diabetes mellitus, cataract, glaucoma. One of the subjects was tested using only one eye. The rest was tested with both eyes. Age ranges from 58 to 95.

Seven unimpaired subjects were tested (or corrected to normal).

Six subjects were tested with myopia (nearsightedness). These subjects normally wore glasses (i.e. correction). Before starting the experiment the glasses were removed and the subjects adjusted to this situation before starting the experiments.

Five unimpaired subjects performed the same experiments with images that were reduced in contrast by a factor of 0.5, 0.25 and 0.125. Contrast sensitivity for these subjects (with Landolt-C symbols) was also measured with test patterns in which contrast was reduced by the same factors.

3 unimpaired subjects performed the same experiments with a fixation 2, 4 and 8 deg in the periphery from a distance of 400 cm (and in one case from 500 cm). The fixation point was fixed to the projection screen above the PC monitor. Scaled versions of the face image were used (scaled by a factor of 70% in each dimension). Only the blur threshold was measured.

2.3.2. RESULTS

Blur threshold

Figure 4a shows the blur threshold versus the resolution threshold, as measured with the Landolt-C test patterns, for different groups of subjects. The blur threshold increases with an increase in the MAR (threshold gap size of the Landolt-C), i.e. the amount of just detectable blur increases with decreasing acuity. Correlation coefficient using all data is 0.92. The relationship is approximately proportional: the blur threshold is about 2 times smaller than the MAR, indicated by the solid line in the graph.

The average ratio between MAR and blur threshold is 1.9 (impaired), 2.1 (normal+myopia), 1.6 (reduced contrast) and 2.1 (eccentric). Statistical tests (one sided student-t test) show that the values of the ratio of groups 2 and 4 are significantly different from the ratio value of group 3. This indicates that the ratio of the reduced contrast data is lower than the ratio of the other data sets. We'll come back to this later. The results also show that the ratio values are not significantly different from 2, except for the data of group 3 (reduced contrast data set).

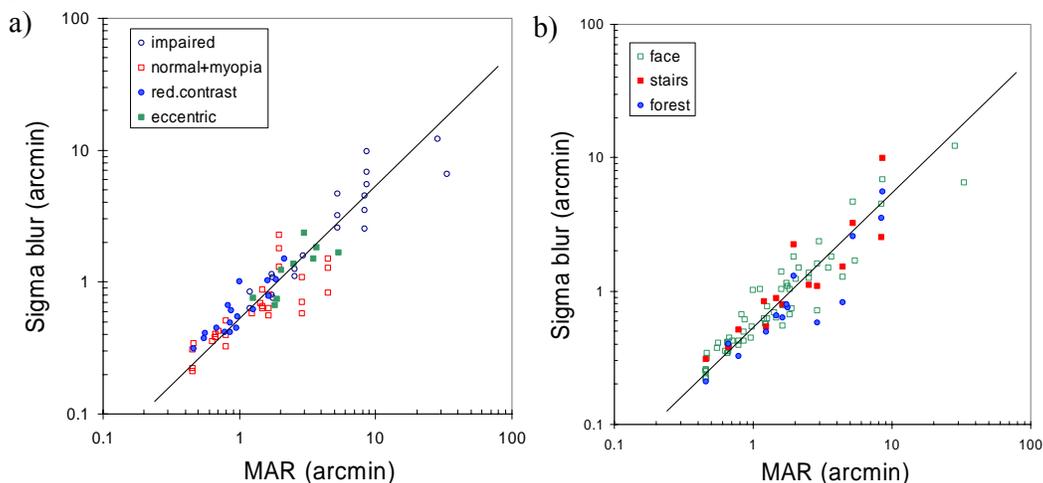


Figure 4. Blur threshold (sigma of blur kernel) versus minimum angle of resolution (= 1/acuity) for the various groups of subjects (a) and for the various types of images (b).

Figure 4b shows the same data, now divided according to image type. The average ratio between MAR and blur threshold is 1.9 (face), 1.8 (stairs) and 2.4 (forest). The ratio's of face and stairs are not significantly different, while the ratio for forest is significantly higher than the other two. Also, only the ratio for forest deviates significantly from 2. This indicates that the blur threshold for the forest image is significantly lower than for the other images. This is likely related to the fact that this image contains more power in the high spatial frequency range.

The main result of this experiment is that the relationship between the blur threshold and the acuity (or MAR) is largely independent of the cause of the reduced vision and of image content.

Contrast threshold

Figure 5 shows the threshold contrast values for removing low-contrast SFs from the image for high (CB1) and medium SFs (CB2) versus contrast thresholds for medium (CT1) and large (CT2)

Landolt-C test patterns. At first sight, these parameters appear unrelated. Using linear regression (in log-log coordinates) we found no significant correlation between CB1 and CT1 for all data, as well as for the two groups individually. We found just (i.e. $p = 0.05$) a significant (positive) correlation between CB2 and CT2 for normal+myopia. This finding matches the expectation that higher contrast sensitivity, as indicated by lower contrast thresholds for Landolt-C patterns, yields higher sensitivity to changes in natural images. The fact that the results do not show a correlation may be partly related to the difficulty of the task. Whereas the difference between the reference and the test image is easy to explain for the first step (blur), it is difficult to explain what characterizes the test image for the 2nd and 3rd steps.

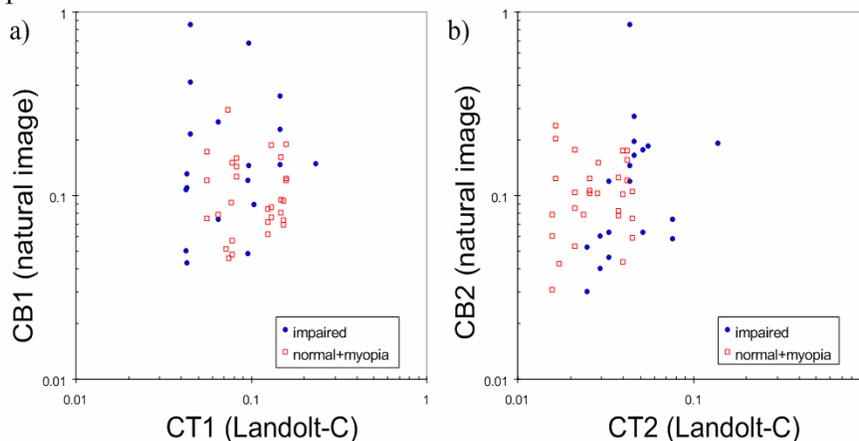


Figure 5. Threshold for removing low contrast SF components from the image versus contrast thresholds for discriminating the orientation of Landolt-C test patterns. Figure 5a shows the contrast threshold for the highest remaining SF-band (CB1) versus the contrast threshold for medium size Landolt-Cs (CT1; 3.2 times the acuity limit), and 4b shows the contrast threshold for a medium SF-band (CB2) versus the contrast threshold for large size Landolt-Cs (CT2; 10 times the acuity limit).

Contrast threshold of simple test patterns (Landolt-Cs) decrease with increasing pattern size: the (geometric) averages of CT1 and CT2 are 9 and 3% resp. In contrast, the threshold contrast levels for band 1 and 2 are about the same, i.e. 11 and 9% resp. This finding also indicates that the relationship between contrast levels for removing SF-components from natural images and contrast thresholds for discriminating simple test patterns is weak.

Another interesting finding is that the contrast thresholds (CT1 and CT2) of visually impaired subjects are comparable in magnitude to the values of normal+myopic subjects (although the values of CT2 are somewhat higher for impaired subjects). This can be attributed to the fact that the contrast thresholds are obtained at test pattern sizes that scale with the minimum angle of resolution of each subject.

2.4. Contrast reduction

As pointed out by Peli (2001), contrast reduction supplies a critical test of the model. Therefore we take a closer look at the effect of contrast reduction on the thresholds. Peli (1996, 2001) uses the assumption that the contrast levels used for removing low-contrast elements from the image are the same as contrast thresholds for detecting Gabor patches. Similarly, the hypothesis could be tested that the contrast levels are directly related to contrast thresholds for discriminating Landolt-C patterns. Since successive SF-bands are separated by an octave, one expects that CB1 is in a simple way related to the contrast threshold at two times the resolution limit. Also, one may expect that CB2 is related to the contrast threshold at four times the resolution limit. Since contrast thresholds

were measured at 3.2 and 10 times the resolution limit and not at 2 and 4 times the resolution limit, we first deduce these threshold values from the available contrast sensitivity data of each subject.

2.4.1. SUBJECTS

Subjects MAH, NB, MvdH and NvD (unimpaired) participated in the experiment. In all cases vision was corrected to normal.

2.4.2. METHOD

We obtained thresholds for the original (face) image, as well as for versions in which the contrast was reduced by 50, 25 and 12.5%. These measurements were repeated three times². For each subject and contrast reduction one set of contrast sensitivity measurements (MAR, CT1, CT2) was obtained. For the reduced contrast values this meant that the MAR was measured for a grey test pattern on a white background, i.e. with matching contrast. In the 2nd and 3rd steps contrast thresholds were obtained at 3.2 and 10 times the corresponding minimum test pattern size.

We fitted a curve to the MAR threshold and contrast threshold data of each subject. Figure 6 shows the result of this exercise for the average data (geometric average over all subjects). From the fitted curve contrast thresholds at 2 and 4 times the MAR were deduced. This fitting procedure also has the advantage that noise in the data is reduced, leading to more reliable estimates.

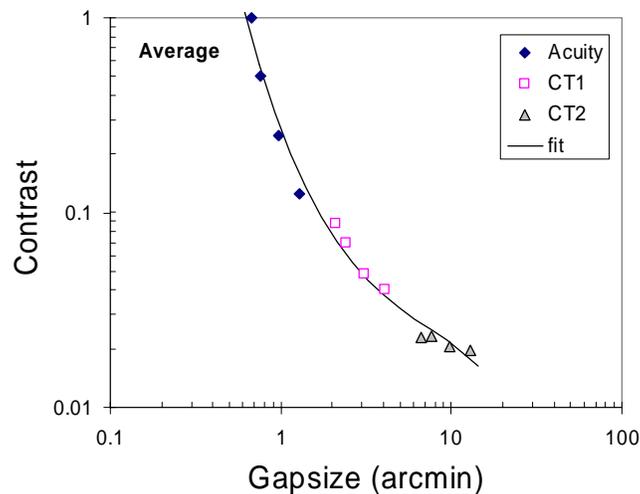


Figure 6. Contrast sensitivity curve connecting the MARs and contrast thresholds obtained by following the standard procedure for various contrast reduction factors (100%, 50%, 25% and 12.5%). The (least squares) fitted line is a 3rd order polynome function (in log-log coordinates).

2.4.3. RESULTS

Figure 7 shows the average blur threshold, CB1 and CB2 as a function of contrast reduction along with the average MARs and contrast thresholds at 2 and 4 times the resolution limit derived from the fitting procedure, in which the (geometric) average was taken over all 4 subjects. The blur threshold and MAR increase with decreasing contrast, while the contrast thresholds and contrast levels (CB1 and CB2) decrease with decreasing contrast. The contrast thresholds decrease with decreasing contrast because they correspond to larger test pattern sizes. Once more, the figure shows that the blur threshold increases in direct proportion with the MAR. Figure 8a shows this more clearly. Shown are the blur thresholds as a function of MAR for all subjects as well as for the average subject. The data of each subject for different contrast (reduction) values is connected. The lines are parallel to the line $y=0.5*x$, and the data fall approximately on this line (indicated by

² One of the subjects of Figure 4 is not included in this analysis, since only a single measurement was obtained per condition.

the dotted line in Figure 8a). This shows that for each subject the blur threshold is proportional to the resolution threshold for a given image, but that the proportionality factor may vary (to some extent) from subject to subject. The average ratios are 1.8 (NB), 1.9 (MAH), 1.9 (MvdH) and 1.5 (NvD). The fact that the same relationship holds between blur threshold and MAR as for the impaired subject (see Figure 4) can be taken as supportive evidence that this relationship is the same for various types of visual impairments. In particular some impairments, such as cataract, effectively lead to a retinal image with lower contrast.

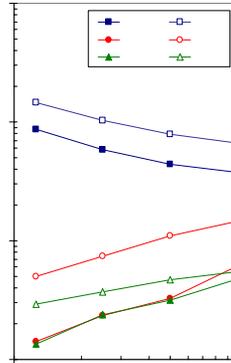


Figure 7. The three threshold values for sensitivity to changes in the natural image (sigma blur, CB1, CB2) along with the contrast sensitivity parameters (MAR, C2x, C4x) as a function of contrast (reduction) for the average subject (geometric averages over all subjects). C2x and C4x represent the contrast threshold values at 2 and 4 times the resolution limit.

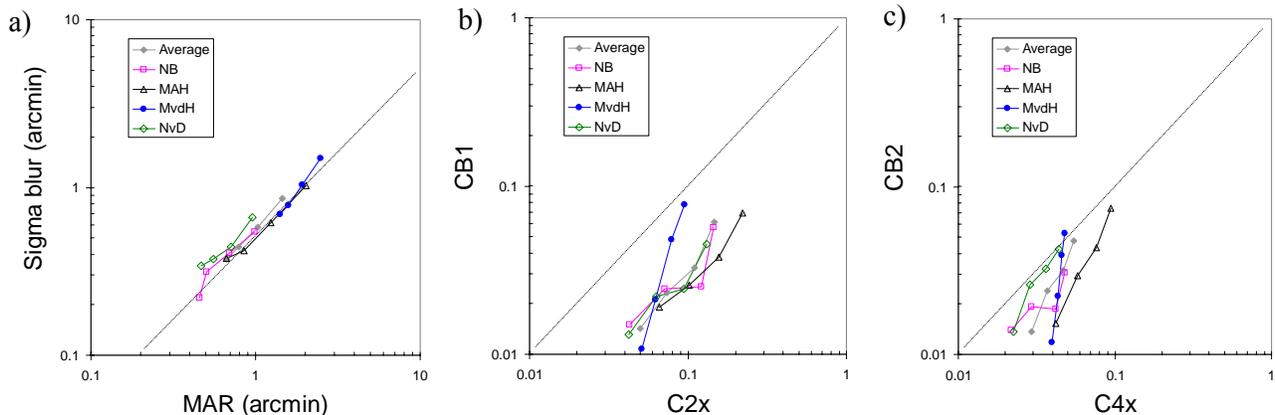


Figure 8. Thresholds for changes in the image as a function of contrast sensitivity thresholds: blur threshold as a function of MAR (a), CB1 vs. C2x (b) and CB2 vs. C4x (c). The dotted line in Figure 8a corresponds to a factor of 2 between x and y, and in the the other figures to a factor of one.

The contrast levels for SF-bands 1 and 2 (CB1 and CB2) decrease with decreasing contrast, as do the contrast threshold values. However, the C2x and C4x values cannot be taken as contrast levels CB1 and CB2 resp. On average CB1 is about 3.1 times smaller than C2x, and CB2 is about 1.5 times smaller than C4x. This means that, if the C2x and C4x thresholds were taken as values for CB1 and CB2, it will be easy to discriminate the transformed image from its original. Furthermore, whereas the contrast thresholds decrease with increasing pattern size for Landolt-C symbols, the contrast levels (CB1, CB2) in the simulation hardly decrease with increasing SF. This puts serious doubt to the general applicability of the model.

A potential solution is to use the model in combination with a steeper sensitivity curve. Peli (1996, 2001), for instance, uses a sensitivity curve that is based on detection of Gabor patterns. This curve is steeper than the Landolt-C curve (see Figure 9). Therefore CB1 and CB2 better match C2x and C4x at high contrast. (Note that a similar steepness with the Landolt-C curve can be obtained by taking the contrast thresholds to the power 1.5). Still, the fact remains that the blur threshold is proportional to the gap size of Landolt-C patterns (MAR). This fact does not hold for a steeper sensitivity curve (the slope in Figure 8a would be steeper than 1, e.g. 1.5). We therefore conclude that contrast sensitivity as measured with Landolt-C test patterns accurately predicts the blur thresholds over a wide range of image contrasts. Furthermore, it is not clear how the band limited contrast model should be used for modeling sensitivity to changes in natural images.

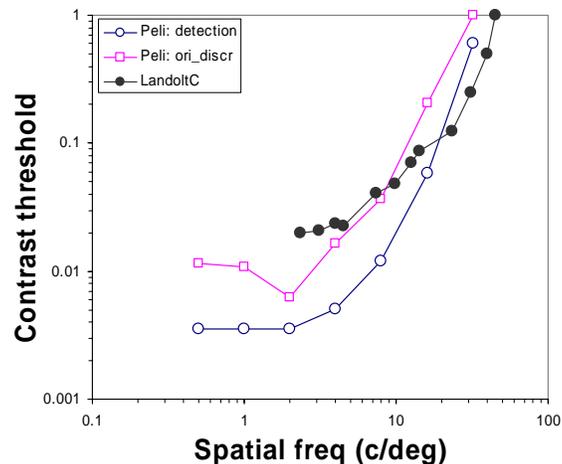


Figure 9. Comparison of various contrast sensitivity curves: detection and orientation discrimination of Gabor patches (data from Peli, 1996) and the average Landolt-C threshold data from this study (see Figure 6), in which the wavelength is taken as twice the gap size (the dominant SF-component).

2.5. Peripheral vision

The results of the main experiment (Figure 4) show that also for eccentric viewing the blur threshold is proportional to the MAR, and that the ratio between them is the same as for foveal viewing. Here we look in more detail into the effect of eccentricity with individual subjects.

2.5.1. SUBJECTS

Subjects MAH, NvD and NB (all unimpaired) participated in the experiment. In all cases vision was corrected to normal.

2.5.2. METHOD

We obtained thresholds for a smaller version of the original (face) image (scaled by 70%, i.e. 8.7 by 12.3 cm) to assure that the image spanned a limited range of eccentricities. A fixation point was added above the centre of the image. Viewing distance was 400 cm (500 cm for NB) for eccentric viewing conditions. The results were compared to the foveal thresholds obtained from a viewing distance of 1000 cm. We obtained only blur thresholds in the natural image, and these measurements were repeated 3 times for subjects MAH and NvD. For NB the measurements were not repeated. We also obtained measurements of acuity for the three subjects under the same conditions (eccentricity and viewing distance).

2.5.3. RESULTS

Figure 10 shows the blur threshold versus the MAR for the three subjects, in which the thresholds for the different eccentricities (0, 2, 4 and 8 deg) are connected (going from bottom-left to top-right with increasing eccentricity). The data are parallel to the diagonal indicating that the blur threshold is proportional to the MAR. The ratio between the two differs somewhat between subjects (on average: 2.6 for NB, 2.0 for MAH, 1.5 for NvD); on average the ratio is around 2. Interestingly, in the contrast reduction experiment average ratio's of 1.8 (NB), 1.9 (MAH) and 1.5 (NvD) were obtained, i.e. comparable for two of the three subjects (note that results of NB are less reliable due to the fact that only a single measurement was obtained). The results show that the blur threshold is proportional to the MAR for various eccentricities, and that the ratio varies (to some extent) from subject to subject. The results are very similar in this respect to the results from the contrast reduction experiment. These findings support the idea that the relationship between the blur threshold and MAR is independent of the type of visual impairment, given that with certain impairments (e.g. macula degeneration) an image will be inspected eccentrically.

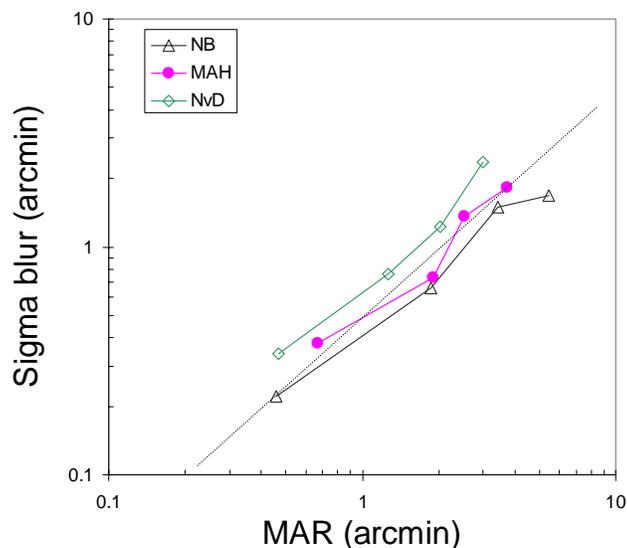


Figure 10. Blur threshold versus minimum angle of resolution (MAR) for all subjects for various eccentricities. Data corresponding to different eccentricities (0, 2, 4 or 8 deg) are connected, and increase from the bottom-left to the upper-right. The dotted line corresponds to a ratio of 2 between MAR and blur threshold.

2.6. Image size

One can assume that, for distances from which the pixels are smaller than the resolution, and when the angular image size remains constant, the model parameters (sigma, CB1, CB2) are independent of viewing distance. This assumption was confirmed in a limited pilot experiment with a single subject. More interestingly is the question whether the model parameters depend on the angular size of the image. This was determined in the following experiment.

2.6.1. SUBJECTS

Subjects MAH, JB and NB (all unimpaired) participated in the experiment. In all cases vision was corrected to normal.

2.6.2. METHOD

We obtained thresholds for the original (face) image and scaled versions of this image (by a factor of 0.7 and 1.47). Measurements were obtained at relative angular sizes of 0.5, 0.71, 1 and 1.41, in which a value of 1 corresponds to an image size of 1.0 by 1.4 deg. The following combinations of relative angular size, viewing distance and image size factor were used: (0.5, 520 cm, 1.47), (0.7, 735 cm, 1.47), (1, 707 cm, 1), (1.4, 700 cm, 0.7). The measurements were repeated 3 times for subjects MAH and NB. For JB the measurements were not repeated. Error bars indicate the SEM (MAH, NB) or the 95% confidence interval in the threshold estimate (JB).

2.6.3. RESULTS

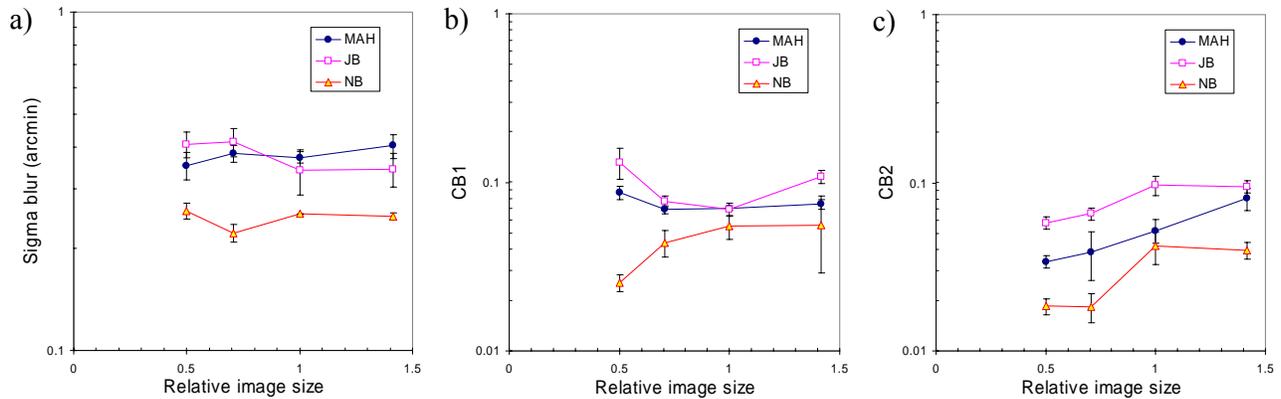


Figure 11. Threshold values versus relative image size: blur threshold (a), CB1 (b) and CB2 (c).

Figure 11 shows the blur threshold, CB1 and CB2 as a function of relative image size. Figure 11a shows no effect of angular image size on the blur threshold. This is yet another indication that the effect of image content on the threshold parameters is limited (see Figure 4b).

Contrast level CB1 does not show a consistent effect of angular size, while CB2 appears to increase with increasing angular size. The latter indicates that the contrast threshold value decreases when the angular size of the elements in the image decreases. This suggests that the extent to which low contrast medium SF-components can be removed from the image relies on image content.

For our purposes, the main result is that the blur threshold is largely independent of the angular size of the image.

2.7. Conclusions & discussion

The main result of this study is that the relationship between the blur threshold and acuity (as measured with Landolt-C symbols) is independent of the cause of reduced vision, whether this is due to visual impairment, contrast reduction, or eccentric viewing. Also, the ratio between the MAR and the blur threshold is largely independent of image content, and (angular) image size. On average the ratio between the MAR and the blur threshold is two, and varies to some (limited) extent from subject to subject. This simple relationship can be taken (as a rule of thumb) to predict how much blur can be introduced in an image before it starts to become noticeable. The fact that this relationship also holds for low contrast images indicates that the contrast sensitivity curve based on orientation discrimination of Landolt-C test patterns is suitable for modeling visibility of elements in natural images.

The fact that the relationship between blur threshold and acuity remains the same for low contrast images and for eccentric viewing gives additional support to the idea that this relationship is independent of the type of visual impairment. Certain impairments lead to a retinal image with lowered contrast (e.g. cataract) while others will be accompanied by eccentric inspection (e.g. macular degeneration).

We also measured contrast levels for removing low contrast SF components from the image (using the local band limited contrast model), and tried to establish the relationship between these contrast levels and contrast thresholds for orientation discrimination of Landolt-C test patterns. When the data of different subjects is considered (in the main experiment), no clear relationship was found between the contrast levels and contrast thresholds for Landolt-C patterns. A closer look at the effect of contrast reduction on the thresholds of individual subjects shows that the contrast levels (natural images) and contrast thresholds (Landolt-C) both decrease with decreasing contrast. However, the contrast levels were found to be lower than predicted by the contrast threshold data (the Landolt-C contrast thresholds at two and four times the resolution threshold size). This means that the contrast thresholds cannot be taken as contrast levels for the first two remaining SF-bands. At first sight this problem may be solved by using a different, steeper contrast sensitivity curve, such as the curve corresponding to detection of Gabors (Peli, 1996). However, such a curve does not predict the decrease in MAR with decreasing contrast accurately. This may be one of the reasons why the effect of contrast reduction in the study by Peli (2001) is less well predicted by the model. Summarizing, it is not clear how the local band limited contrast model can be used to accurately model visibility of elements in natural images.

Instead of using a fixed contrast sensitivity curve for modelling (an approach taken by Peli and colleagues) we tried to determine the appropriate model parameters in a series of (three) successive steps. The disadvantage of our approach is that the final result will be easier to discriminate from the original than the result of the first step, i.e. the discrimination score will be higher than 75%. To arrive at a 75% correct score somewhat lower values of the parameters should be taken. Discriminating images from which certain SFs were removed with low contrast (the task in steps 2 and 3) appears to be quite difficult. This is likely the reason why no correlation with contrast thresholds was found in the main experiment.

The fact that the blur threshold is inversely proportional to the acuity for various conditions indicates that the visibility of (high contrast) image elements is the same for different subjects and circumstances when the viewing distance is scaled inversely with acuity. A simple way to estimate which elements can be perceived by a person with lower acuity is by viewing the image from a larger distance, e.g. the same elements are visible for a person with acuity of 0.1 from a distance of 5 m as for a person with acuity of 1.0 from a distance of 50 m! In principle this rule only holds for high contrast elements, reflecting the fact that acuity is measured at a similar contrast. However, the data (Figure 5) shows that the contrast thresholds measured at a multiple times the resolution threshold are similar for impaired and unimpaired subjects. This indicates that the difference between contrast sensitivity curves from different subjects can be approximated by a shift along the (log) size axis. This indicates that the rule of thumb described above holds (reasonably well) for low contrast elements.

The finding that the blur threshold is proportional to the MAR for high and for low contrast values indicates that it is appropriate to blur to a degree that varies from location to location, depending on the local contrast. In this way, parts with low contrast are blurred more, and the simulation will better present the amount of information that can be removed before it becomes noticeable by the impaired subject. Such an image transformation seems similar to the use of the local band-limited contrast

model. However, in our case a single contrast measure is used for each location, instead for each location *and* SF-band. The image transformation method we propose can be regarded as a series of steps to remove more and more information from the image. The main impact results from the first approximation in which the image is blurred (homogenously) to an extent that is derived from the acuity for high contrast Landolt-C patterns. Next, more subtle changes can be made by introducing extra blurring in low contrast parts. Therefore, as a rough approximation, homogenous blur will suffice.

An unsolved issue is which (local) contrast measure is appropriate. One possibility is to use the maximum of all band-limited contrast measures at each location. By taking the maximum, one remains on the safe side, i.e. one can be sure the image is not blurred beyond the threshold for any of the SF bands. However, this measure varies largely from location to location (see Figure 12b). It therefore does not well implement the idea of successive degradation. It seems more appropriate to use a localized version of the root-mean-square contrast. This measure is equal to the standard deviation (sd) divided by the (absolute value of the) average (av): $C = |sd / av|$, in which the standard deviation and average are taken over a local (Gaussian) window (see Figure 12c). The effect of the image transformation is that sharp contrast edges remain in the image while regions of low contrast become homogeneous (see Figure 12d). By the way, the effect is similar to that of nonlinear diffusion (e.g. Perona & Malik, 1987), which may prove to be a good method for simulation.

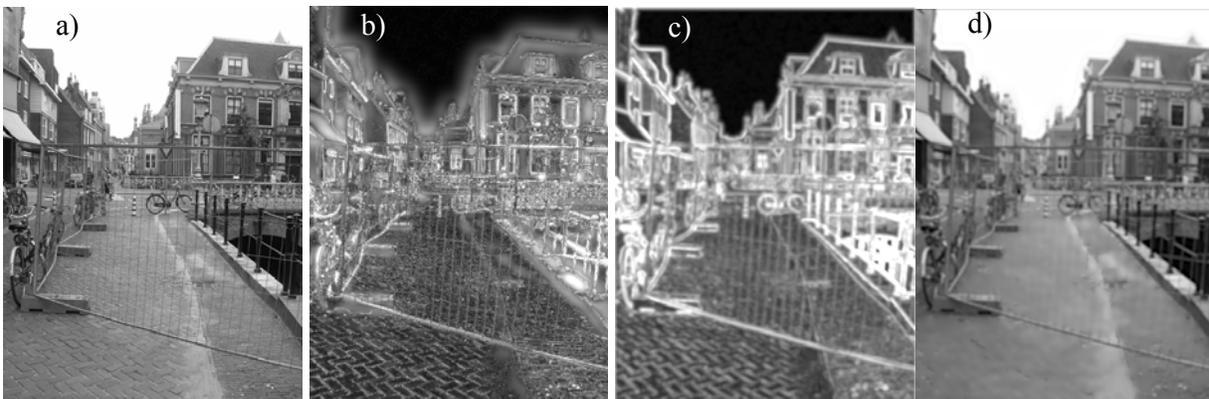


Figure 12. Various local contrast measures: a) original image, b) the maximum of the local band-limited contrast measures (maximum over all SF-bands), and c) the standard deviation divided by the average over a local area (using a Gaussian window). Figure 12d shows the result of local blurring using the local contrast measure based on the standard deviation (Figure 12c).

Summarizing, our results indicate a clear and simple relationship between acuity (measured with Landolt-C symbols) and the amount of just noticeable blur. The connection between contrast sensitivity and the contrast level below image elements can be removed (using the local band limited contrast model) is less clear. Based on these results we propose a simulation method in which the local blur adapts to the local contrast.

3. Fixation dependent simulation

3.1. Introduction

We investigate various ways to simulate reduced vision in the periphery by blurring the image by an amount that depends on the retinal location and a person's resolution at that location. A way to implement this is described by Perry & Geisler (2002). They developed an algorithm and software for creating and displaying, in real time, arbitrary variable resolution displays, contingent on the

direction of gaze. They demonstrated that the software can be used to simulate the visual fields of normal individuals as well as low vision patients. The resolution map of low vision patients is derived from Goldman perimeter data. The Goldman perimeter data is used to scale the resolution map of a normal visual field (for a value of -40 dB the resolution is multiplied by 0 and for a value of 0 dB the resolution is multiplied by 1; personal communication). The normal resolution map is based on a model fit to contrast sensitivity data from Robson & Graham (1981) for detecting sinusoidal gratings. The contrast threshold (CT) is described by:

$$CT(f, e) = CT_0 \exp\left(\alpha f \frac{e + e_2}{e_2}\right), \quad (1)$$

in which f represents spatial frequency (c/deg), e the retinal eccentricity (deg), CT_0 is the minimum contrast threshold, α is the spatial frequency decay constant, and e_2 is the half-resolution eccentricity. (The fit parameters are $\alpha = 0.106$, $e_2 = 2.3$ and $CT_0 = 1/64$). From the model the threshold eccentricity can be derived at which the contrast threshold for a particular SF equals 1. In the simulation the image is decomposed into various SF-bands using a Laplacian Pyramid model (Burt & Adelson, 1983). For each SF-band the threshold eccentricity is calculated belonging to the maximum SF (Nyquist frequency), elements at larger eccentricities are removed, and the resulting image is reconstructed.

This simulation is similar to that used by Peli and colleagues, in the sense that detection thresholds of sinusoids are used to predict the resolution threshold. We have found that blur thresholds are better predicted by orientation thresholds for Landolt-C patterns than by detection thresholds for sinusoidal patterns (see section 2). We therefore propose to use the relationship between acuity and blur threshold to arrive at a validated simulation of peripheral vision. Using this relationship, the (local) blur can be predicted from the (local) acuity (corresponding to orientation thresholds of Landolt-C patterns). An advantage of coupling the simulation to orientation thresholds of Landolt-C patterns is that this method is commonly used to measure acuity. As a first approximation one could use the resolution map of an unimpaired person and add scotomas, much like Perry & Geisler (2002) do, i.e. by transforming the resolution map on the basis of perimeter data. To be certain the minimum resolution fits that of the low vision patient, one could use the extra restriction that this is in correspondence with the patient's acuity (for unrestricted viewing). To arrive at an even more realistic simulation the link between perimeter data and acuity should be established and incorporated into the simulation. Below we describe two experiments. In the first we determine how acuity (measured with Landolt-Cs) varies with eccentricity for unimpaired persons. The second (pilot) experiment is focussed at establishing the link between perimeter data and Landolt-C acuity.

3.2. Acuity in the periphery

In the first experiment we investigate how acuity of unimpaired subjects varies across the visual field.

3.2.1. Method

The stimuli were displayed on a PC monitor. They consisted of a Landolt-C test pattern displayed in the centre at high contrast, dark on a white background. In the periphery a fixation cross was displayed. The position of the Landolt-C pattern relative to the fixation cross is given by the eccentricity ε and angle ϕ , representing the angle with the horizontal (e.g. for $\phi = 90$ deg the fixation cross is displayed beneath the Landolt-C pattern). In each session thresholds for a single eye were obtained used, and the other was occluded. The eye was put in the centre of the screen at a fixed distance. The subjects' head was fixed by a chinrest. At each trial a Landolt-C test pattern appeared

with the gap in one of eight orientations. The subject indicated the orientation of the gap. An adaptive staircase 8AFC method (QUEST, see Watson & Pelli, 1983) was used with 50 trials to determine the threshold gap size (Minimum Angle of Resolution) corresponding to a 75% correct score. Two (unimpaired) subjects participated in the experiment (MAH, NvD). MAR values were obtained for eccentricities of $\varepsilon = 2, 4, 8, 15$ and 30 deg, along directions $\phi = 180, 225, 270$ for the right eye (MAH and NvD), and $\phi = 0, 315, 270$ for the left eye (NvD). These values prevent that the test pattern falls into the blind spot. In each session thresholds were obtained for various combinations of eccentricities and angles ϕ . A viewing distance of 60 cm was chosen for eccentricities between 2 and 8 deg, viewing distance of 30 cm for $\varepsilon = 15$ deg, and 20 cm for $\varepsilon = 30$ deg. The PC-monitor was placed on a moving platform. This allowed us to change the viewing distance from condition to condition in random order.

3.2.2. Results

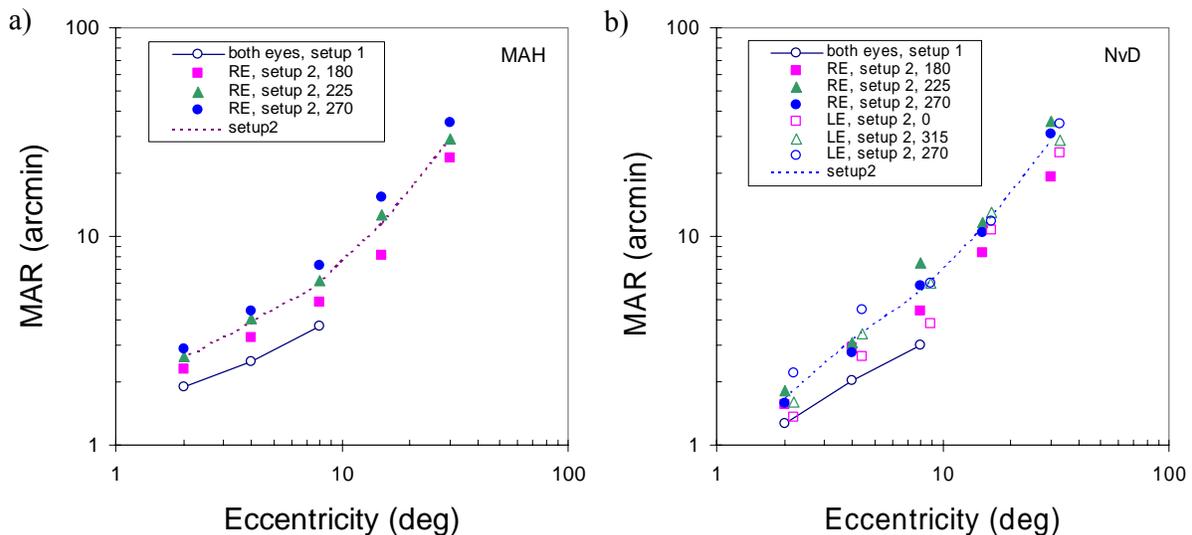


Figure 13. MAR as a function of eccentricity for the two subjects, for different angles (ϕ) and eyes (right = RE, left = LE) (the symbols for LE are shifted to the right to prevent overlap). The dotted line represents the average over all values. Also shown are the thresholds from the previous experiment taken with a different setup (setup 1). In the first setup thresholds were obtained from a large distance with two eyes with a 4AFC method, for an angle ϕ of 270 deg.

Figure 13 shows the thresholds as a function of eccentricity for different angles (ϕ) for each eye, along with the data of the previous experiment obtained with a different setup. The comparison between blur threshold and acuity was done using the first setup (see section 2). The current setup is more convenient for measuring acuity in the periphery. To estimate the blur threshold for a given retinal location it is necessary to determine how the thresholds measured with the second setup translate into thresholds measured with the first setup. The two setups differ in a number of ways. First, in setup 2 an 8AFC method was used (to arrive at stable threshold results with fewer measurements). Second, the data were obtained with a single eye instead of with both eyes. Also, a much shorter viewing distance was used. The thresholds obtained with setup 2 and setup 1 differ on average by a factor of 1.5 (MAH: 1.50; NvD:1.53).

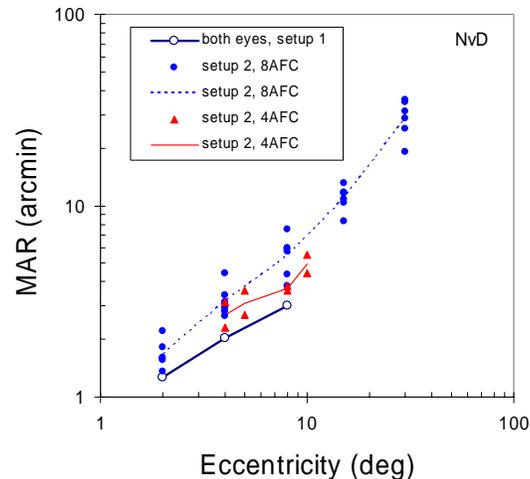


Figure 14. Minimum angle of resolution as a function of eccentricity for setup 2 obtained with an 8AFC and with a 4AFC procedure, along with the MARs obtained with setup 1 (with 4AFC-method). The different points correspond to different angles φ (see Figure 13). The lines represent average trends.

In a control experiment with a single subject (NvD) the effect of using an 8AFC-method instead of a 4AFC-method was investigated. Figure 14 shows the thresholds obtained with the 4AFC method and the 8AFC method with the second setup, along with the results obtained with the first setup (with 4AFC method). The average difference between setup 2 with 4AFC method and setup 1 is 1.3. The difference in factor (1.3 vs. 1.5) can thus be attributed to the use of the 8AFC-method. The remaining difference (by a factor of 1.3) is likely mainly due to the use of a single eye instead of two.

3.3. Perimeter data vs. acuity

Visual field analysers (such as the Goldman or the Humphries Field Analyser) obtain light intensity thresholds for detecting a disk of light. Here we try to establish how this relates to the acuity (measured with Landolt-C symbols). In principle one cannot expect these measures to be related. A perimeter threshold is akin to a contrast threshold whereas acuity is a measure of resolution. A system with low resolution can have good contrast sensitivity and vice versa. However, it is trivially the case that when no disk of light can be detected (in an (absolute) scotoma) the resolution threshold is indefinite. Furthermore, our previous results (see section 2, Figure 5) show that contrast thresholds are correlated to acuity, since they are about the same for pattern sizes that scale with the MAR. In the pilot experiment described below we will try to determine the statistical relationship between acuity and perimeter value.

3.3.1. Method

The Humphries Field Analyser was used with the SITA-standard test, with background luminance of 31.5 ASB (10 cd/m^2). Tested was the left eye of one observer, age 49, who suffered from a combination of visual impairments (surgically corrected peripheral retinal ablation, artificial implant lens to remove cataract, open angle glaucoma), and a central acuity of 0.33. Perimeter thresholds were obtained for eccentricities up to 30 deg. These thresholds were compared to eccentric acuity obtained with the setup and method described in section 3.1. Threshold measurements were obtained at 24 different grid points (without repeats). The viewing distance ranged from 20 cm for the highest eccentricity to 55 cm for the lowest eccentricity.

3.3.2. Results

Figure 15 shows the raw data from the perimeter experiment as well as the MAR obtained in a selection of (the same) locations. Figure 16a shows the relationship between the MAR and the perimeter threshold at corresponding locations separated according to visual field quadrant (1=top-right, 2=top-left, etc).

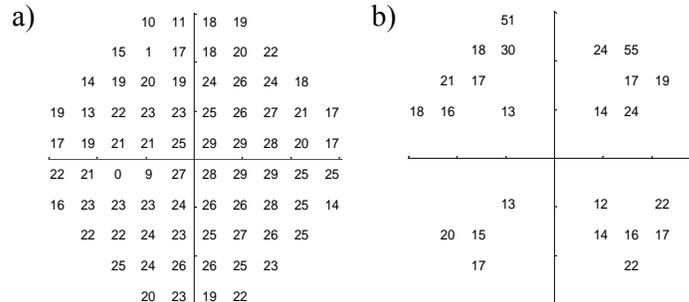


Figure 15. Raw results from the Humphries test (a): perimeter thresholds in dB, and acuity test (b): MAR values in arcmin.

The correlation between these parameters is not significant ($R = -0.26$). Figure 16b shows the relationship between MAR and eccentricity. As expected, the MAR increases with increasing eccentricity ($R = 0.55$). Whereas the (complete) perimeter results (not shown here) indicate significant pattern deviations in the upper left quadrant (quadrant 2), the MAR thresholds are not significantly worse in quadrant 2 than in quadrant 1. Figure 16b also shows the MAR data of (unimpaired) subject MAH. Surprisingly, the MAR thresholds of MAH are comparable to the thresholds of the impaired subject.

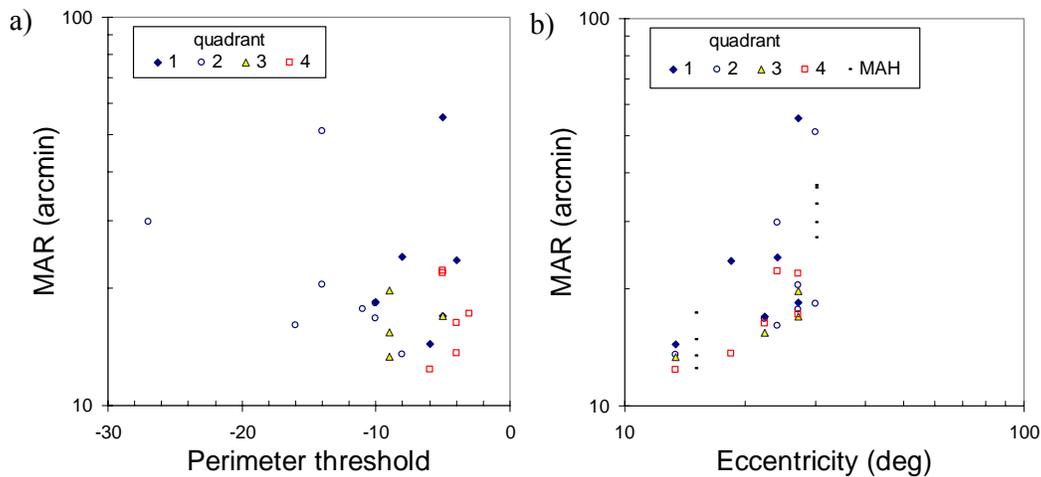


Figure 16. MAR as a function of perimeter threshold value (a) and as function of eccentricity (b) for the 4 quadrants (1 = top-right, 2 = top-left, etc), along with thresholds for unimpaired subject MAH (see Figure 13a).

We have to conclude that the results from this pilot experiment are insufficient to draw firm conclusions, and that a more extended experiment with more subjects is required.

3.4. Simulating eccentric viewing

We use the Visual Field Simulator (Geisler & Perry, 1998; Perry & Geisler, 2002) to model eccentric viewing. We make use of the “space variant imaging-toolbox” developed by Geisler & Perry (2005). The resolution map describes the way in which the resolution varies across the retina. Perry & Geisler use a normal resolution map based on a model fitted to the data from Robson & Graham (1981) corresponding to detection of sinusoids. Our results indicate that contrast sensitivity based on orientation discrimination of Landolt-Cs more accurately describes sensitivity to blur in natural images (at least for the conditions studied, i.e. foveal viewing and eccentricities up to 8 deg). Therefore, we will derive the resolution map of an unimpaired subject from the Landolt-C acuity data obtained in the periphery.

To calculate the resolution map for a given distance the following steps are taken. First, for each pixel the eccentricity is calculated. A smooth function is fitted to the MAR thresholds. Then, for each pixel the MAR is deduced using the fitted function and the eccentricity of the pixel. Our results (section 2) show that the ratio between the MAR and the sigma of the (Gaussian) blur threshold is 2. So, to arrive at the amount of (Gaussian) blur σ for each pixel the MAR is divided by 2. Finally we deduce the resolution from the amount of blur from: $\text{res} = 0.6418 / \sigma$ (personal communication with Perry).

An example of such an exercise is shown in Figure 17. Figure 17 shows contour graphs of the resolution maps covering an angle of 54 x 42 deg. Figure 17a shows the resolution map proposed by Perry & Geisler (2002). Figure 17b shows the resolution map derived from the acuity thresholds from subject MAH (right eye; see Figure 13a). The area around the fixation point for which there is no blur is much larger in the latter case. For larger eccentricities the resolution falls off more rapidly, and for large eccentricities the resolution is about the same in the two cases (the last contours, corresponding to a resolution of 0.2, approximately coincide).

Figures 17c and d show the result of applying these resolution maps to an image of 54 x 42 deg. Figure 17c shows the result for the resolution map used by Perry & Geisler, and Figure 17d shows the result when then a resolution map is used derived from the acuity data from subject MAH (right eye).

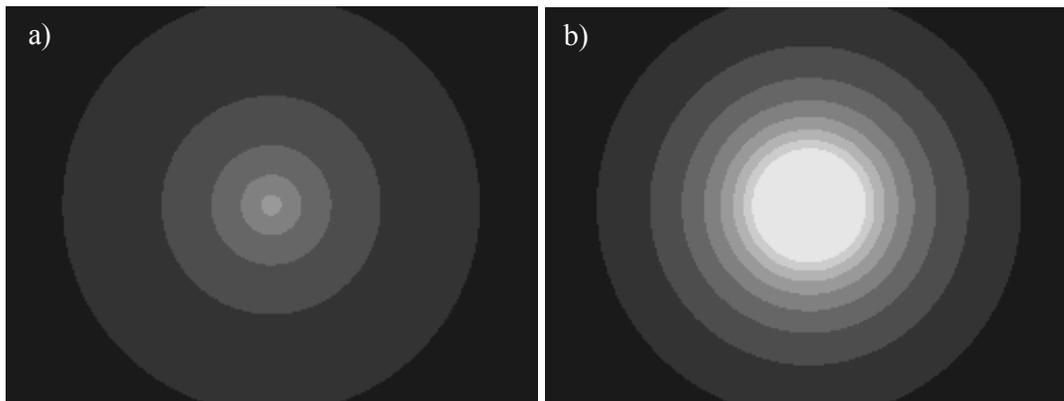




Figure 17. Resolution maps and the result of the transformation method using these resolution maps on an example image. The figure on the left shows the resolution map and result based on the model used by Perry & Geisler for an unimpaired subject, while the figures on the right represent the resolution map derived from the Landolt-C acuity and the transformed image. The image and resolution maps span a visual field of 54 x 42 deg.

The resolution map of low vision patients can be deduced from the normal resolution map using the same method as followed by Perry & Geisler (2002), or by simply adding the scotomas. To arrive at an improved simulation one could use the extra restriction that the minimum resolution of the resolution map cannot be lower than the one derived from acuity during free viewing.

4. Simulating Glare

Disability glare is glare that impairs vision. It is largely caused by scattering of light inside the eye because of the imperfect transparency of the optical components of the eye. It results in a kind of veiling luminance.

4.1. Model

The illuminance dE resulting from a light source with luminance L spanning an angle α is given by:

$$dE = L * d\alpha \quad (2)$$

The intensity of the veiling luminance L_{veil} depends on E_{glare} , the illuminance at the eye caused by the glare source, and the angle between the glare source and the line of sight θ (in first order approximation, see CIE, 2001):

$$L_{veil} = \frac{9.2 E_{glare}}{\theta^2} \quad (3)$$

We start with an image of luminance values. We calculate the contribution to the veiling luminance resulting from any other pixel, i.e. the other pixel is regarded as the light source causing the veiling luminance. We then calculate the total veiling luminance for a given pixel by summing the contribution of all other pixels:

$$L_{veil_total} = \iint 9.2 L(x, y) d\alpha(x, y) / \theta(x, y)^2 dx dy \quad (4)$$

Instead of using equation (2) an improved version formulated by the CIE can be used (CIE General Disability Glare equation):

$$\frac{L_{veil}}{E_{glare}} = \frac{10}{\theta^3} + \left[\frac{5}{\theta^2} + \frac{0.1p}{\theta} \right] \cdot \left[1 + \left(\frac{Age}{62.5} \right)^4 \right] + 0.0025p \quad (5)$$

in which p represents the *eye pigmentation factor* (0 for black, 0.5 for brown, 1 for light and 1.2 for very light eyes). This formula is valid for angles θ between 0.1 en 100 deg.

The formula shows that the disability glare rapidly increases with an increase in age. For an age of 62.5 the 2nd term is twice as large. A good approximation for angles θ between 1 and 30 deg is given by:

$$\frac{L_{veil}}{E_{glare}} = \frac{10}{\theta^2} \cdot \left[1 + \left(\frac{Age}{70} \right)^4 \right] \quad (6)$$

This equation is the so called CIE Age-adjusted Stiles-Holladay Disability Glare equation. This veiling luminance can be multiplied by a factor κ that depends on the amount of light scatter in someone's eye. For cataract patients a reasonable lies between 2 and 8 (see Yspeert, 1993):

$$\frac{L_{veil}}{E_{glare}} = \frac{10}{\theta^2} \cdot \kappa \quad (7)$$

In practice it may take a long time to perform the calculations for each pixel. To accelerate the calculating process one can choose to perform the calculation on a larger scale, effectively using an image reduced in size to calculate the veiling luminance.

4.2. Results

Figure 18b shows the result of the glare simulation applied to Figure 18a. Since the luminance values are unknown assumptions have to be made about the conversion of luminance values to pixel values. In certain circumstances the calibrated luminance values might be known and can be used. For instance, when a virtual design (resulting from e.g. ray tracing techniques) is evaluated, or when the luminance values are measured (e.g. with a luminance camera³). To make the effect clearly visible we assumed a gamma of 6 (i.e. the luminance equals the pixel value to the power 6) and a glare constant κ of 8 (representative of a cataract patient). Note that the veiling luminance varies over the image, and is larger close to the horizon than in the foreground.

³ This camera obtains images of calibrated luminance values.

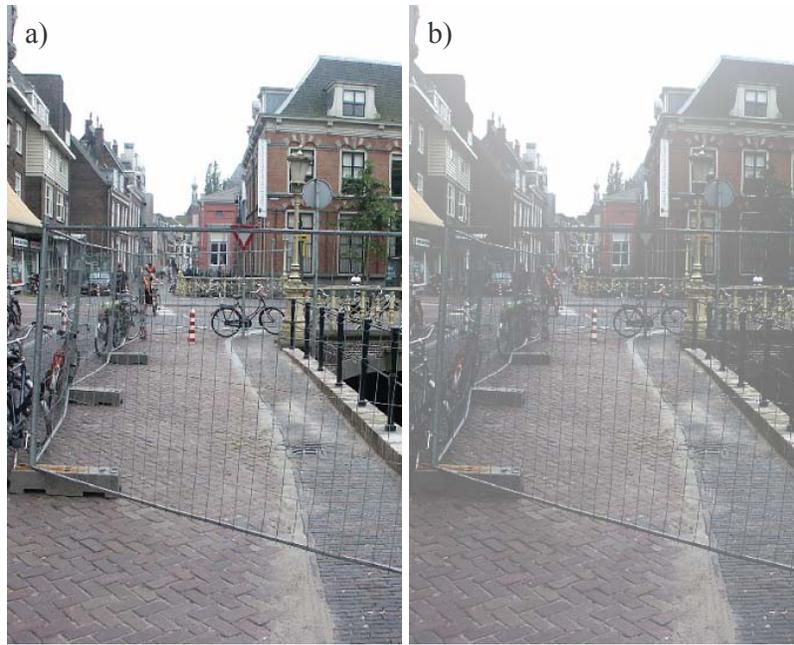


Figure 18 Example of disability glare. Figure 18a shows the original image, and Figure 18b shows the transformed image to which a veiling luminance is added. To make the results clear we assumed that the luminance is given by the pixel luminance value to the power γ , with $\gamma = 6$ (a more reasonable value lies around 2), and $\kappa = 8$.

5. Conclusion

The goal of this study was to find and implement ways to visualize the limitations of visually impaired persons. We investigated two image transformation methods: i) a fixation independent (global) and ii) a fixation dependent (local) one. The first type shows the information that is available when a person inspects the scene with his/her best part of the visual field. This simulation shows which elements are and which are not visible when attention is focussed on that part of the scene. As such, the simulation is related to human performance for recognising objects. The second type shows the information that is available everywhere in the visual field when a person fixates a particular part of the scene. This simulation gives insight into visual search for an item, i.e. the simulation is related to the conspicuity of elements in the scene. It also shows which information is available when the time to inspect a scene is short. Since both types of simulations give insight into other aspects of visual processing (object recognition vs. object detection) they both are valuable simulations (depending on the task one or the other will be more suitable).

We evaluated a simulation method developed by Peli and colleagues based on the local band limited contrast model. Instead of postulating a contrast sensitivity curve to be used for simulation the simulation parameters were measured in an experiment in which the image was degraded more and more in three successive steps. We tried to establish the relationship between the simulation parameters (blur, contrast level high SF-band, contrast level medium SF-band) and contrast thresholds for orientation discrimination of Landolt-C symbols. The main result of the experiments is that a simple relationship exists between acuity and amount of just noticeable blur, and that this relationship is largely independent of the cause of reduced vision (visual impairment, low contrast, eccentric viewing). The fact that this relationship holds for low contrast images indicates that the contrast sensitivity curve based on discrimination of Landolt-C symbols is appropriate for modelling visual limitations. The results however also show that the contrast thresholds can not be used directly in connection with the local band limited contrast model for simulation. We do not see how this model can be used to yield accurate predictions of one's visual limitations. Instead we propose a

(first type of) simulation in which the (local) blur adapts to the (local) contrast. This comes down to a direct (local) implementation of the relationship between amount of just detectable blur and the acuity (at that contrast level).

The visual field simulation developed by Perry & Geisler (2002) is used to implement the second type of (fixation dependent) simulation. This general method uses a resolution map that can be tailored at will and may include scotomas. The image transformation is linked to the eye fixation (which can be linked to eye movement equipment or to the PC-mouse) such that the local blur depends on the fixation location. The image transformation is implemented in real time software. Perry & Geisler use a resolution map based on detection thresholds for Gabor patches. Our results from the fixation independent study show that contrast sensitivity to Landolt-Cs may be more appropriate for modelling. At least, we established the relationship between blur and sensitivity to Landolt-C acuity. Therefore, we propose to use the relationship between acuity and blur threshold to yield an image deformation that corresponds to the discrimination threshold. We first measured Landolt-C acuity thresholds across the visual field to establish the resolution map of an unimpaired person. To arrive at the resolution map of impaired subjects Perry & Geisler scaled the resolution map of unimpaired subjects by using (Goldman) perimeter data. The result of this operation is a resolution map that includes scotomas. The scaling method has not been validated yet. So, we tried to establish the link between acuity and perimeter data. Since blur and acuity are related, this would validate the local blur method used by Perry & Geisler. We performed a pilot experiment with a single subject to establish this link, but could not find one. We have to conclude that more data is required to draw firm conclusions. For the moment we propose a simulation based on Landolt-C acuity of unimpaired subjects, and which is scaled by perimeter data (in a way similar to that used by Perry & Geisler). To arrive at an improved simulation the (fixation free) acuity could be taken into account, in order to be sure that the maximum amount of information available at each location does not exceed the information available from the fixation independent simulation.

We also investigated ways to simulate the effect of glare, since for certain types of impairments this has a big impact (e.g. cataract). For this simulation we resorted to well validated equations (CIE, 2001). A veiling luminance is calculated for each part of the image that results from the light originating from other parts of this image. The veiling luminance is therefore representative for the situation that an observer looks at that part of the image. This simulation can be used as a first step to creating both types of simulations (fixation independent or fixation dependent).

Acknowledgements

We wish to thank our collaborators dr. I.T.C. Hooge, and B.P.L.M. den Brinker for their contributions and helpful discussions and prof. dr. G.H.M.B. van Rens, ir. F.F. Jorritsma and T. Blom for their corporation and help in setting up the experiments with low vision subjects. We also very much appreciate the effort, work and discussions with trainees from University of Utrecht, J. van Elst, N. Boogaards and N. van Dijk, who conducted many of the experiments. We would also like to thank the students E.J.L. Brunenberg, H.J. Buisman, W.L. de Graaf, J.W. van Triest, J. Uittenbogaard from the Eindhoven University of Technology (Dept. of Biomedical Engineering) and their supervisor dr. L.M.J. Florack for their contribution to the development of a useful image transformation method.

References

Brainard, D. H. (1997) The Psychophysics Toolbox, *Spatial Vision*, **10**, 433-436.

Burt, P.J., Adelson, E.H. (1983) The Laplacian pyramid as a Compact Image Code. *IEEE Transactions on Communications*, **31** (4), 532-540.

Cannon Jr, M.W. (1985) Perceived contrast in the fovea and periphery. *J. Opt. Soc. Am. [A]*, **2**, 1760-1768

CIE (2001) Report of CIE Committee TC1-50 Disability Glare Formula

Geisler W.S. and Perry J.S. (1998) A real-time foveated multiresolution system for low bandwidth video communication. *SPIE Proceedings: Human Vision and Electronic Imaging*, **3299**, 294-305.

Geisler W.S. and Perry J.S. (2005) Space variant imaging software:
<http://svi.cps.utexas.edu/svtoolboxdoc/svtoolbox/Contents.html>

Ginsburg, A.P. (1975) Is the illusory triangle physical or imaginary? *Nature*, **257**, 219-220.

Peli, E. (1990) Contrast in complex images. *J. Opt. Soc. Am. A.*, **7**, 2032-2040

Peli, E., Yang, J., Goldstein, R. (1991). Image invariance with changes in size: the role of peripheral contrast thresholds. *J. Opt. Soc. Am. A* , **8**, 1762–1774.

Peli, E. (1995) Simulating normal and low vision. in *Vision Models for Target Detection and Recognition*, E. Peli, ed. (World Scientific, Singapore, 1995), pp. 63–87.

Peli, E. (1996) Test of a model of foveal vision by using simulations. *J. Opt. Soc. Am. A* , **13**, 1131–1138.

Peli, E., Geri, G.A. (2001) Discrimination of wide field images as a test of a peripheral-vision mode. *J. Opt. Soc. Am. A.*, **18** (2), 294-301.

Pelli, D. G. (1997) The VideoToolbox software for visual psychophysics: Transforming numbers into movies, *Spatial Vision*, **10**, 437-442.

Perona, P., Malik, J. (1987) Scale space and edge detection using anisotropic diffusion. In IEEE Computer Society Workshop on Computer Vision, 1987.

Perry, J.S. and Geisler, W.S. (2002) Gaze-contingent real-time simulation of arbitrary visual fields. In B. Rogowitz and T. Pappas (Eds.), *Human Vision and Electronic Imaging, SPIE Proceedings*, San Jose, CA.

Robson, J.G., Graham, N. (1981) Probability summation and regional variation in contrast sensitivity across the visual field. *Vision Research*, **21**, 409-418.

Simoncelli, E.P. (2004) *matlabPyrTools* - Matlab source code for multi-scale image processing.
<http://www.cns.nyu.edu/~lcv/software.html>

Walraven, J., Alferdinck, J.W.A.M. (1997) Color displays for the color blind. *Proceedings of the 5th Color Imaging Conference, Color Science, Systems, and Applications, November 17-20, 17-22*. Scottsdale, Arizona.

Walraven, J. (2000) *Kleurenblind – Zien en toch niet zien*. NIZW, Arnhem

Watson, A.B., Pelli, D.G. (1983) QUEST: a Bayesian adaptive psychometric method. *Perc. & Psych.*, **33** (2), 113-120.

Yspeert, J.K. (1993) *Light Scattering in the Human Eye*, PhD thesis.