

Simulating visibility under reduced acuity and contrast sensitivity

WILLIAM B. THOMPSON,^{1,*} GORDON E. LEGGE,² DANIEL J. KERSTEN,² ROBERT A. SHAKESPEARE,³ AND QUAN LEI²

¹School of Computing, University of Utah, 50 So. Central Campus Dr., Room MEB-3190, Salt Lake City, Utah 84112-9205, USA

²Department of Psychology, University of Minnesota, N218 Elliott Hall, 75 E. River Rd., Minneapolis, Minnesota 55455, USA

³Department of Theatre, Drama, and Contemporary Dance, Indiana University, 275 North Jordan, Bloomington, Indiana 47405-1101, USA

*Corresponding author: thompson@cs.utah.edu

Received 15 November 2016; revised 28 February 2017; accepted 28 February 2017; posted 1 March 2017 (Doc. ID 280894); published 22 March 2017

Architects and lighting designers have difficulty designing spaces that are accessible to those with low vision, since the complex nature of most architectural spaces requires a site-specific analysis of the visibility of mobility hazards and key landmarks needed for navigation. We describe a method that can be utilized in the architectural design process for simulating the effects of reduced acuity and contrast on visibility. The key contribution is the development of a way to parameterize the simulation using standard clinical measures of acuity and contrast sensitivity. While these measures are known to be imperfect predictors of visual function, they provide a way of characterizing general levels of visual performance that is familiar to both those working in low vision and our target end-users in the architectural and lighting-design communities. We validate the simulation using a letter-recognition task. © 2017 Optical Society of America

OCIS codes: (330.3790) Low vision; (330.1070) Vision - acuity; (330.1800) Vision - contrast sensitivity; (330.5020) Perception psychology.

<https://doi.org/10.1364/JOSAA.34.000583>

1. INTRODUCTION

Visual accessibility is a property of environmental spaces that allows the use of vision to travel efficiently and safely through such spaces, to perceive the spatial layout of key features in the environment, and to keep track of one's location in the layout. It plays a central role in independent mobility, which in turn is an important prerequisite for full participation in modern society. Reduced mobility and associated social isolation and economic disadvantage are among the most debilitating consequences of vision loss. In 2010 in the United States, approximately four million people had uncorrectable low vision, with projections up to seven million by 2030 and 13 million in 2050 [1]. Only a small percentage of those with low vision have total blindness, and most of those with low vision use residual visual capabilities for navigation and other functions [2]. This paper describes an approach to helping architects and lighting designers increase the utility of this residual visual capability for those with significant loss of acuity or contrast sensitivity. The approach simulates the low-vision visibility of features during the design process, allowing the identification of potential mobility hazards and landmarks that might go unrecognized by low-vision individuals. Identifying such hazards during

the design phase of a project make amelioration much easier than waiting until after actual construction.

Architects and lighting designers think in terms of manipulating the geometry, materials, and lighting of a space in order to achieve particular functional and aesthetic objectives. For those with visual impairment involving loss of acuity or contrast sensitivity, however, stimulus properties such as visual angle and contrast are critically important. Further complicating design for visual accessibility is the complex interaction among geometry, materials, and lighting arrangement that determines the light field surrounding the viewer. For normally sighted individuals, general guidelines relating to light levels, glare, and contrast are often sufficient to minimize visually indistinct hazards. For those with low vision, however, general guidelines are not sufficient. Because of the importance of angular feature size and limits on contrast sensitivity in low vision, the exact positioning and nature of light sources, surfaces, and the viewer can have a profound effect on visibility.

The value of lighting and the visual environment for the aging eye has seen a rise in importance in the architectural professions during the past decade, prompting focused symposiums, new guidelines, and recommended practices

(e.g., [3–5]). A remaining challenge is to evaluate the visibility outcomes when making design choices based on these recommendations. While many of the tools used to model architectural projects during the design phase can produce images of the project, a few can now also produce physically and photometrically accurate simulations of the space being designed. Developing new systems for simulating the effects of reduced acuity and contrast to build on these photometrically accurate renderings would provide a designer with an opportunity to evaluate design choices in the context of low vision and visibility. Working within the project's design palette, a designer could modify textures, colors, shapes, and lighting to optimize visibility, while retaining the character of the design. To evaluate the implementation of the final design specifications, acuity and contrast sensitivity filters can be applied to calibrated high dynamic range (HDR) photographs of the completed environment. Additionally, visibility studies using HDR images of existing environments would be of value when considering renovation or remodeling strategies. The integration of these tools into the architect's workflow will provide the missing link between general guidelines and the successful visual accessibility of a project.

Our approach to simulating the visibility impacts of loss of acuity and contrast sensitivity builds on the work of Peli [6,7], who described a method for transforming an image to simulate the visibility associated with a particular contrast sensitivity function (CSF). An original image is first transformed into a set of bandpass images, each representing an unnormalized contrast measure over a narrow range of spatial scales. Each pixel in each unnormalized contrast band is then divided by the local luminance of the original image surrounding the pixel location, providing a measure of local contrast closely related to Michelson contrast. Next, each pixel in each unnormalized contrast band is thresholded based on a criterion that compares the local contrast values to a CSF evaluated at peak band response frequency of the band filter. Finally, the thresholded unnormalized contrast bands are reassembled to produce an output image. The method uses variations in the embedded CSF to represent the reduction in pattern sensitivity of people with low vision.

This nonlinear method has three advantages over a linear filtering approach that uses the CSF as if it were a modulation transfer function (MTF) [7]: (1) Image contrast that is below the contrast specified by the CSF is removed, rather than just being attenuated. This reduces variability associated with viewing conditions and the viewer's own acuity and contrast sensitivity (called the *double filtering* effect [8]). (2) Image contrast that is above the CSF-specified threshold is left intact, thus better modeling the suprathreshold response of the visual system [9,10]. (3) CSF-based thresholding is done in a spatially localized manner that takes into account local luminance, which has a strong effect on contrast perception. Taken together, these properties serve to remove image features predicted to be not visible, while leaving features predicted to be visible clearly apparent in the output.

Our most important contribution extending the work of Peli [6,7] is to provide a way of parameterizing the method using standard clinical measures of acuity and contrast

sensitivity and to validate the parameterization using a letter-recognition task. While clinical measures of acuity and contrast sensitivity are known to be imperfect predictors of visual function for specific individuals, they provide a way of characterizing general levels of visual performance that is familiar to both those working in low vision and our target end-users in the architectural and lighting design communities. We also propose a way of reducing one type of artifact associated with the hard thresholding used in [7], we make suggestions for including color, which may aid acceptance by architects and lighting designers, and we provide an implementation that takes as input HDR images that are linearly encoded in luminance. Multichannel models such as [6] have advantages when dealing with broadband signals [11]. They also facilitate implementing adaptation to local intensity (see [12], which uses a spatial image pyramid approach).

2. PARAMETERIZING THE SIMULATION OF REDUCED ACUITY AND CONTRAST SENSITIVITY

The nature of the visibility filtering achieved by the method described in [6,7] is controlled by the CSF that it uses. There is substantial debate as to the appropriate functional form for CSFs modeling human vision (e.g., [11,13,14]). We chose to use the CSF described in Chung and Legge [15], since it is the only one that has been shown to fit empirical CSF data from a substantial group of low-vision subjects. While the Chung and Legge CSF was developed based on band-limited stimuli (sine wave gratings), when correctly calibrated (see Section 3), it proved sufficient to model recognition of local broadband stimuli such as letters.

Chung and Legge [15] propose a CSF of the following form:

$$S_l(f_l) = \begin{cases} SP_l - (f_l - FP_l)^2 w_L^2 & \text{if } f < FP \\ SP_l - (f_l - FP_l)^2 w_H^2 & \text{if } f \geq FP \end{cases} \quad (1)$$

where

S	contrast sensitivity,
S_l	$\log_{10}(S)$,
f	spatial frequency,
f_l	$\log_{10}(f)$,
SP	peak contrast sensitivity,
SP_l	$\log_{10}(SP)$,
FP	frequency of peak contrast sensitivity,
FP_l	$\log_{10}(FP)$,
w_L	constant for low frequency portion of CSF,
w_H	constant for high frequency portion of CSF.

As is common with CSF formulations, sensitivity is defined in terms of Michelson contrast. Based on a best-fit analysis to measured normal vision CSF data, Chung and Legge [15] use the following rate constants: $w_L = 0.68$ and $w_H = 1.28$.

Equation (1) has the shape of an asymmetric parabola when plotted in f_l - S_l space (see Fig. 1). Empirical evidence supports the claim that Eq. (1) can successfully approximate a wide range of normal and low vision by adjusting SP and FP [15]. In particular, low vision involving a reduction of acuity can be modeled by sliding the normal vision CSF function left in f_l - S_l

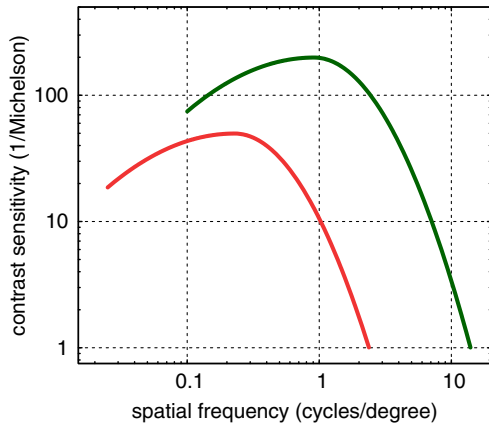


Fig. 1. Chung and Legge [15] CSF is an asymmetric parabola when plotted in f_l - S_l space. The plotted values show two instances of the CSF, one shifted left (lower acuity) and down (lower contrast sensitivity) compared to the other.

space, while low vision involving a reduction in contrast sensitivity can be modeled by sliding the normal vision CSF function down in f_l - S_l space. To emphasize this left-right/top-down sliding, we reparameterize Eq. (1) by replacing SP by $c \times SPN$ and FP by $a \times FPN$:

$$S_l(f_l) = \begin{cases} SPN_l + \log_{10} c - (f_l - FPN_l - \log_{10} a)^2 w_L^2 & \text{if } f < a \times FPN \\ SPN_l + \log_{10} c - (f_l - FPN_l - \log_{10} a)^2 w_H^2 & \text{if } f \geq a \times FPN \end{cases} \quad (2)$$

where

- SPN peak normal vision contrast sensitivity,
- SPN_l $\log_{10}(SPN)$,
- FPN frequency of peak normal vision contrast sensitivity,
- FPN_l $\log_{10}(FPN)$,
- c contrast sensitivity adjustment,
- a acuity adjustment.

A. Adjusting for Reduced Contrast Sensitivity

One common clinical measure of peak contrast sensitivity uses the Pelli–Robson Contrast Sensitivity Chart [16]. The chart consists of black or gray letter groups with decreasing contrast, all on a white background. Weber contrast is typically used to characterize the contrast of these darker optotypes viewed on a lighter background:

$$C_w = \frac{L_b - L_c}{L_b}, \quad (3)$$

where

- C_w Weber contrast,
- L_b luminance of background,
- L_c luminance of character.

The Pelli–Robson contrast sensitivity score is based on the threshold contrast for letter recognition, expressed as the log of the Weber contrast for threshold visibility (the negative of the log is commonly used to make the scores positive):

$$PR = -\log_{10} CT_w, \quad (4)$$

where

- PR Pelli–Robson score,
- CT_w Weber contrast for threshold visibility.

One important caveat is relevant to the use of the Pelli–Robson chart: “While it is supposed that the chart will normally be used at a distance of 3 m, it can be used at much nearer distances for assessment of low vision” [16]. For our purposes, we will assume that the Pelli–Robson score reflects a viewing distance as close as necessary to easily resolve the letters on the chart, though this is not always done in a clinical setting.

The contrast sensitivity adjustment parameter c in Eq. (2) represents the ratio of the peak contrast sensitivity being simulated to the normal vision peak contrast sensitivity:

$$c = \frac{SPL}{SPN}, \quad (5)$$

where

SPL simulated low vision peak contrast sensitivity.

Accounting for the difference between Michelson contrast, as used in the CSF, and Weber contrast, as used in the Pelli–Robson score, and assuming that a PR score of 2.0, which indicates a threshold Weber contrast of 1/100 and a threshold Michelson contrast of 1/199, corresponds to normal vision

$$C_m = \frac{L_b - L_c}{L_b + L_c} = \frac{C_w}{2 - C_w}, \quad (6)$$

$$SPL = \frac{1}{CT_m} = \frac{2 - \exp_{10}(-PR)}{\exp_{10}(-PR)}, \quad (7)$$

$$SPN = 199, \quad (8)$$

where

- C_m Michelson contrast,
- CT_m Michelson contrast for threshold visibility.

This yields

$$c = \frac{2 - \exp_{10}(-PR)}{199 \cdot \exp_{10}(-PR)}. \quad (9)$$

B. Adjusting for Reduced Acuity

The most common clinical measures of visual acuity utilize Snellen or logMAR letter charts and test for the smallest high-contrast letters that can be read accurately. Snellen scores are expressed as a ratio of the distance from which the chart is viewed to the distance from which the smallest readable characters subtend an angle of 5 arcmin. In the United States, this ratio is usually normalized to have a numerator of 20, corresponding to a distance of 20 ft, while elsewhere the common numerator is 6, corresponding to a distance of

6 m. A Snellen fraction evaluating to 1 indicates nominally normal acuity; smaller numeric values of the Snellen fraction correspond to lower acuity. LogMAR acuity scores are the negative of the base 10 logarithm of the Snellen fraction. A logMAR value of 0 indicates nominally normal acuity; larger logMAR values correspond to lower acuity. We will use the numeric value of the Snellen fraction in adjusting filtering for reduced acuity, noting that logMAR values are easily converted to this value if needed.

Shifting the CSF used for filtering in f_l - S_l space to simultaneously account for reductions in contrast sensitivity and loss of acuity is complicated by the fact that standard measures of contrast sensitivity such as Pelli–Robson scores and standard measures of acuity such as Snellen scores are associated with different parts of the CSF. The Pelli–Robson score provides information about the peak sensitivity of the CSF (lowest visible contrast), while the Snellen or logMAR score provides information about the high-frequency cutoff of the CSF (finest visible high-contrast pattern). This produces an interaction between acuity and peak contrast sensitivity as they affect the positioning of the CSF in f_l - S_l space. Figure 2 illustrates the problem. In Fig. 2(a), the CSF has been shifted directly downward so as to preserve the frequency associated with the peak contrast sensitivity. The decrease in the high-frequency cutoff is apparent. In Fig. 2(b), the CSF has been shifted downward and to the right so as to preserve the high-frequency cutoff frequency. In this case, the frequency associated with the peak contrast sensitivity increases. Figure 3 shows the CSF cutoff frequency as a function of peak contrast sensitivity for a peak contrast sensitivity frequency corresponding to normal vision.

The acuity adjustment parameter a in Eq. (2) specifies an acuity-related shift of the peak of the CSF. As indicated above, setting this value based on a measure of the high-frequency cutoff of the CSF is not straightforward. Given particular values for a and c in Eq. (2), the high-frequency cutoff of the CSF can be found by solving Eq. (2) for f_l , assuming $S_l(f_l) = 0$. This yields

$$FC_l = FPN_l + \log_{10}(a) + \frac{(\log_{10}(c) + SPN_l)^{\frac{1}{2}}}{w_H}, \quad (10)$$

where

$$FC \quad \text{high-frequency cutoff of CSF,} \\ FC_l \quad \log_{10}(FC).$$

Given a high-frequency cutoff frequency, FC , the corresponding peak sensitivity frequency, FP , can be found by solving the equation

$$SPN_l + \log_{10} c - (FP_l - FC_l)^2 w_H^2 = 0 \quad (11)$$

for FP_l . This yields

$$FP_l = FC_l - \frac{(\log_{10}(c) + SPN_l)^{\frac{1}{2}}}{w_H}. \quad (12)$$

When there is no decrease in contrast sensitivity (i.e., $c = 1$), a is equal to the numeric Snellen value, since the ratio of low-vision peak sensitivity frequency to normal vision peak sensitivity frequency is the same as for the corresponding cutoff

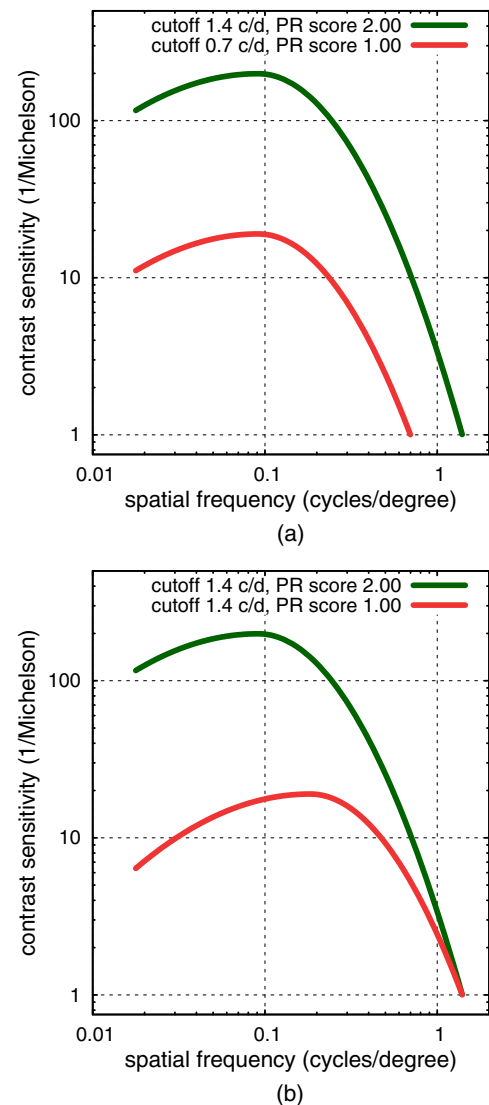


Fig. 2. (a) Contrast sensitivity plots for different peak contrast sensitivities, but the same peak contrast sensitivity frequencies; (b) contrast sensitivity plots for different peak contrast sensitivities, but the same acuity as measured by cutoff frequency.

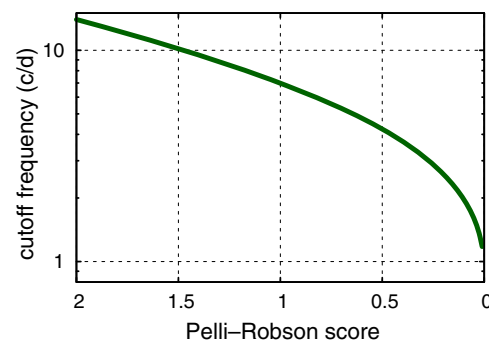


Fig. 3. CSF cutoff frequency as a function of peak contrast sensitivity for a peak contrast sensitivity frequency corresponding to normal vision.

frequencies. When $c < 1$, the low-vision contrast sensitivity cutoff frequency value is

$$FCR = \text{Snellen_value} \times FCN, \quad (13)$$

where

FCR low-vision contrast sensitivity cutoff frequency,
 FCN normal vision contrast sensitivity cutoff frequency.

The corresponding low-vision contrast peak sensitivity frequency, FPR , can be found using Eq. (12). Finally, the value of the parameter a can be computed using

$$a = \frac{FPR}{FPN}. \quad (14)$$

3. CALIBRATING THE SIMULATION OF REDUCED ACUITY AND CONTRAST SENSITIVITY

Section 2 described how to adjust the filter for different acuities and contrast sensitivities by shifting the CSF used for filtering in f_l - S_l space relative to nominal normal vision values for SPN (peak contrast sensitivity) and FPN (the frequency at which SPN occurs). Calibrating the filter thus requires choosing appropriate values for SPN and FPN . There is an extensive literature on the quantitative nature of normal vision CSF as it relates to the detectability of sinusoidal contrast gratings (e.g., [17]). Much less is known, however, about the relationship between contrast sensitivity and acuity for letter charts, as used in the measures for specifying the degree of visual degradation in our filter (see [18–20]).

We used the following approach to setting SPN and FPN . First, SPN was set to 199, which is the 1/Michelson contrast sensitivity corresponding to a Pelli–Robson score of 2.0. We then varied FPN , testing the readability of variously sized high-contrast letters filtered with various simulated acuity reductions. Trials consisted of presenting subjects with an image of the 10 Sloan characters in random order on a computer screen, with all 10 characters visible at once. The characters were shown in two rows. Subjects were asked to read the characters in order, indicating when a particular character was clearly illegible. They were asked to identify each of the displayed letters one at a time, without comparing it to the

other letters being displayed. The images were filtered to simulate a particular level of low-vision acuity (see Fig. 4). Individual trials used characters of a single logMAR size, with the logMAR size varying between trials.

Filtering was done using HDR linearly encoded floating point luminance values, with the output converted to a low dynamic range (LDR) 8 bit/pixel format using the sRGB non-linear luminance encoding. The maximum possible value of the HDR images was set to 95% of the corresponding maximum LDR displayable values so as to avoid problems with saturation that sometimes occur in LCD monitors at the high end of the display range. Display was done on an Asus PA246Q 24" LCD monitor, set to sRGB mode. Screen size and viewing distance were such that all but the largest character size subtended approximately the correct angle. (Because of the inability to fit all 10 of the filtered logMAR 1.6 characters on the screen at this viewing distance, these characters were displayed to subjects at half size to avoid disruptions in the stimuli presentation due to large changes in viewing distance that would otherwise be needed. The filtering was done with the correct visual angle.) Results were insensitive to moderate changes in viewing distance.

For each filter setting, results were based on the smallest characters for which a subject could correctly read seven or more of the 10 letters. If the filter is correctly adjusting for the effects of acuity loss, we would expect that the smallest readable filtered characters would correspond to the acuity specified for the filtering. Figure 5 shows the results for $FPN = 0.915$ cycles/deg, corresponding to a high-frequency cutoff of 14.0 cycles/deg. The data are averaged from the results obtained from six normal vision subjects (average age 24.8 years). All participants gave written informed consent with procedures approved by the University of Utah's Institutional Review Board. As can be seen, the $FPN = 0.915$ cycles/deg setting produced results that were quite close to this prediction. Subjects were very consistent in their response, with standard error at each acuity value ranging from 0.0 to 0.022.

The high-frequency normal vision CSF cutoff is commonly assumed to be about 30 cycles/deg or greater. The best-fit normal vision cutoff of 14.0 cycles/deg that we found for our task involving the legibility of filtered characters is substantially




Fig. 4. Screenshots of two stimuli used to evaluate the setting of FPN . (a) shows logMAR 1.3-sized characters, filtered to simulate an acuity of logMAR 1.2; (b) shows logMAR 1.1-sized characters, filtered to simulate an acuity of logMAR 1.2. Letters in (a) are readily recognizable, while those in (b) are not. ($FPN = 0.915$ cycles/deg for both examples.)




Fig. 5. Empirically determined acuity for simulated low vision with $FPN = 0.915$ cycles/deg.

lower than this. This may be due to low-frequency information below 2.5 cycles/character being sufficient for character recognition (14.0 cycles/deg corresponds to 1.2 cycles/character for logMAR 0.0 characters), or perhaps a more general disassociation between grating-based cutoffs and letter-based cutoffs (see [18,21,22]).

This first test of filter calibration involved only high-contrast targets. A second test, involving the same six subjects, simulated a logMAR 1.1 acuity with two different levels of peak contrast sensitivity, one corresponding to a Pelli–Robson score of 2.0 (normal vision), and the other corresponding to a Pelli–Robson score of 1.0 (moderate loss). The readability of letters with three different contrasts was evaluated: one corresponding to a Pelli–Robson score of 0.75 (0.178 Weber contrast), one corresponding to a Pelli–Robson score of 0.50 (0.316 Weber contrast), and one corresponding to a Pelli–Robson score of 0.00 (1.000 Weber contrast). Table 1 shows the predicted smallest legible character size along with the average actual smallest size readable by the six subjects. Figure 6 provides the same information in graphical form, plus a plot of the two CSFs.

As an additional check on calibration, the same six subjects judged the lowest contrast at which characters of different sizes

Table 1. Predicted and Actual Smallest Legible Characters for Reduced Acuity and Contrast Sensitivity

Filter logMAR	Filter Peak		Predicted Target logMAR	Actual Average Target logMAR
	Contrast Sensitivity (Pelli–Robson Score)	Target Contrast (Pelli–Robson Score)		
1.1	2.0	0.75	1.40	1.42
1.1	2.0	0.50	1.30	1.23
1.1	2.0	0.00	1.10	1.13
1.1	1.0	0.75	1.58	1.70
1.1	1.0	0.50	1.40	1.43
1.1	1.0	0.00	1.10	1.12

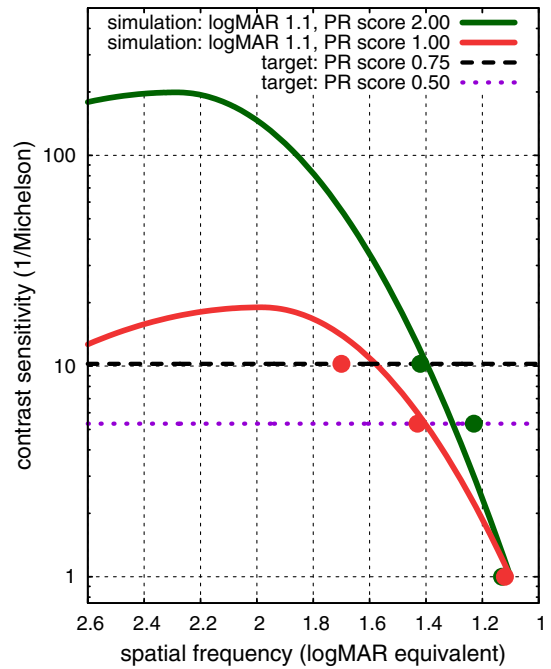


Fig. 6. Empirically determined smallest legible characters for reduced acuity and contrast sensitivity.

were legible when filtered to simulate normal vision by using the unshifted CSF. The contrast of the displayed images was increased after filtering, making this a test primarily of the information retained in the filtering and thus minimizing confounds with the subjects’ own contrast sensitivity. The results are shown in Fig. 7, compared with tests of human performance in letter-recognition tasks as reported by [19] and the author-collected data reported in [20].

Finally, Fig. 8 shows the result of applying the filter to the image of a logMAR chart. In Fig. 8(a), the character size of the

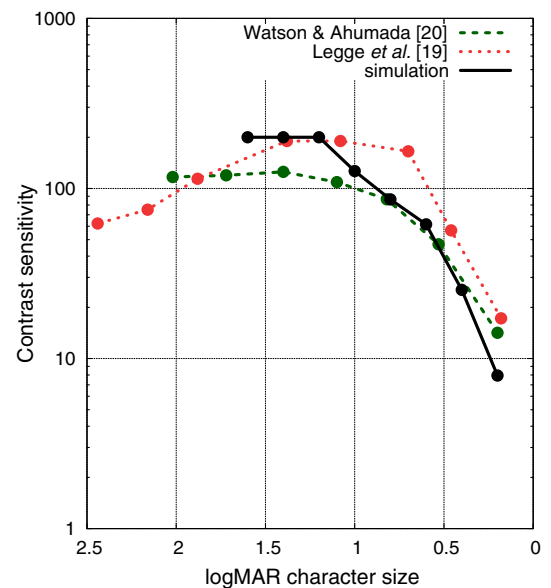


Fig. 7. Empirically determined contrast sensitivity for simulated normal vision.

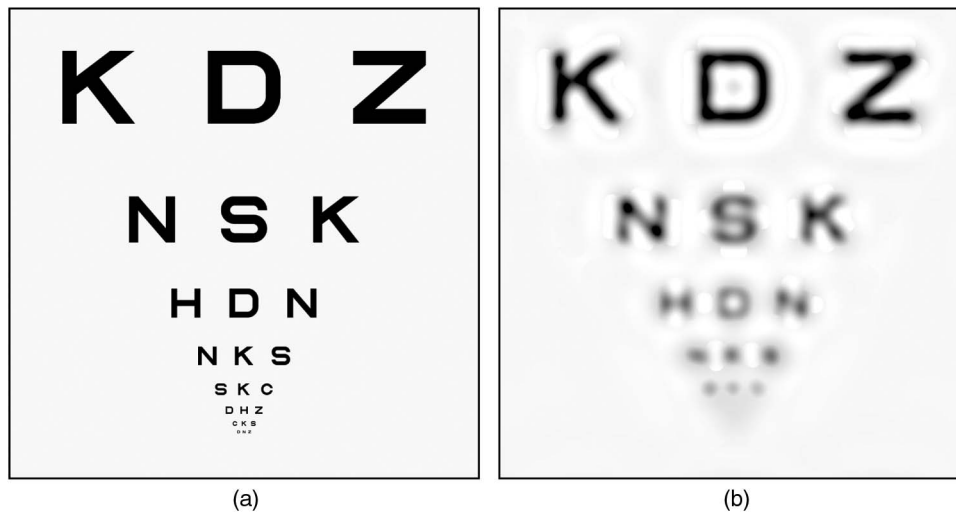


Fig. 8. (a) Original logMAR chart, with third line from top corresponding to logMAR 1.1 and the fourth line from the top corresponding to logMAR 0.9. For correct character size, view the chart from a distance equivalent to 3.33 times the width of the chart image. (b) Original logMAR chart, filtered to simulate an acuity of logMAR 1.0. The third line is readable; the fourth line is not.

top row is logMAR 1.5. In subsequent rows, the character size drops by 0.2 logMAR units per row. Figure 8(b) shows the image in Fig. 8(a) filtered with a simulated logMAR acuity of 1.0 but no reduction in peak contrast sensitivity. The third line from the top corresponds to an original character size of logMAR 1.1 and is clearly readable in the filtered image. The fourth line from the top corresponds to an original character size of logMAR 0.9 and is clearly illegible in the filtered image.

4. OTHER ISSUES

We have extended the filtering approach of [6,7] in several other ways of significance to our target user audience of architects and lighting designers. These include the reduction of one type of artifact that appears when simulating low vision, the addition of color, and the ability to effectively process high dynamic range input.

A. Thresholding Artifacts

The filtering method described in [6] can produce artifacts when simulating significant low vision that are not easily noticed in simulations of normal or less degraded vision. Figure 9 shows an example of one type of these artifacts. Figure 9(a) shows an input image with two equal-sized bars of differing contrast with respect to the background. Figure 9(b) shows the results of filtering Fig. 9(a) using the contrast thresholding method described in [6], using settings that would be predicted to leave the high-contrast bar visible but suppress the lower-contrast bar. While the visibility of the two bars in Fig. 9(b) is as intended, there are also multiple banding artifacts surrounding the visible bar. In Fig. 9(c), the banding artifacts are much less noticeable.

Figure 10 illustrates the source of these banding artifacts. Figure 10(a) is a plot of the luminance across one row in the original image. Figure 10(b) shows a plot of one row of the unnormalized contrast band associated with a peak band




Fig. 9. (a) Vertical bars of same width with two different contrasts with respect to the background; (b) low-vision simulation using thresholding in [6]; (c) low-vision simulation using improved thresholding.

response frequency of 4.0 cycles/image (0.13 cycles/deg based on the field of view of the input image). Figure 10(c) shows the result of thresholding the unnormalized contrast band based on a per-pixel comparison of the local-luminance-normalized contrast band with the CSF. Note that the signal associated with the low-contrast bar is gone, but that there are also notches near the zero crossings of the above-threshold contrast associated with the high-contrast bar. This is the primary cause of the banding seen in Fig. 9(b). These banding artifacts can be eliminated by preserving below-threshold contrast values when they are spatially near above-threshold values of the same sign [see Fig. 10(d)]. In our case, we define near to be within 25% of the wavelength at the peak response point of the band. This can

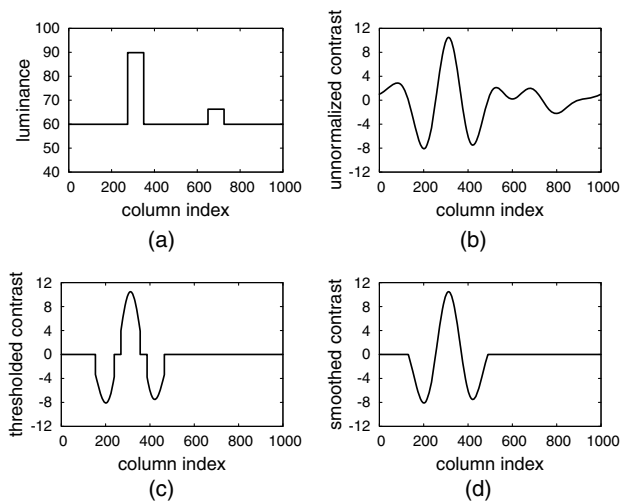


Fig. 10. (a) Luminance profile of Fig. 9(a); (b) plot of one of the bands produced by the low-vision simulation filter; (c) [6] style thresholding of band; (d) improved thresholding of band.

be done relatively efficiently by using $O(n)$ distance transforms such as in [23].

B. Color

The filtering method described in [6] deals only with luminance. Architects and lighting designers are used to working exclusively with color images and find the grayscale display of simulated low vision to be distracting. There is an extensive literature on designing visual displays for those with color-deficient vision, even when the designer has normal color vision (e.g., [24]). Rather than duplicating that functionality in our low-vision simulation filter, we have chosen to implement a more generic approach allowing easy creation of color output by a variety of methods. We start by transforming the input image into the CIE xyY (Commission Internationale de l'Éclairage) color space, which uses one luminance channel and two normalized chromaticity channels. Simulation of loss of acuity and contrast sensitivity is done by filtering the luminance channel as above. We do not apply the same filtering to the chromaticity channels, since the CSF does not predict the threshold chromatic sensitivity as a function of spatial frequency. Instead, we use a linear filtering approach applied to each of the chromaticity channels in isolation, in which the CSF is used as a MTF. This can result in some pixels ending up outside the xy gamut, in which case they are clipped to the nearest in-gamut value. Finally, a color image is reassembled in which the new Y (luminance) channel is the nonlinearly filtered original Y channel and the new x and y channels are the linearly filtered and clipped versions of the original x and y channels. For demonstration purposes, we have shown how this approach can be used to provide any desired level of color saturation.

C. High Dynamic Range

Real architectural spaces have a much higher dynamic range of light levels than can be either displayed on conventional devices or represented in standard 8-bit/color image file formats. It is important that this HDR be accounted for in simulations of

low vision. While there are a variety of ways of creating and representing HDR images of existing spaces (see [25]), our focus on design applications has led us to use the RADIANCE modeling and rendering system [26]. Unlike almost all other modeling systems, RADIANCE allows precise photometric specifications of lighting and material properties and supports photometrically accurate simulations of light transport. In addition, RADIANCE (as with other HDR software) uses a linear representation of luminance, so filtering is not distorted by the nonlinear luminance encoding of LDR image representations. This is particularly important when quantifying actual contrast values.

One problem with filtering HDR images is that extremely bright areas of the input image can result in excessive ringing in the output. We deal with this by clipping these extremely bright areas to an input-dependent maximum luminance level. (Glare, which is not simulated in the current model, represents a substantial problem for low vision individuals. We elaborate on this in the Discussion.) This level is determined using a variant of the RADIANCE glare identification heuristic. First, average luminance over the input image is computed and used to set a preliminary glare threshold value. This average is not a robust estimator, since it is strongly affected by very bright glare pixels or glare pixels covering a large portion of the image. To compensate for this, a second pass is done in which a revised average luminance is computed based only on pixels less than or equal to the preliminary glare threshold. This revised average luminance is then used to compute a revised glare threshold, which is used as the clipping value for preprocessing the image.

Another issue with using HDR images in simulations of low vision is the need to display the results on LDR displays. This requires some form of *tone mapping* [25]. Tone mapping is still an area of active research. Fortunately, the fact that the filtering approach we are using removes rather than just attenuates contrast predicted to be invisible at a particular level of acuity and contrast sensitivity makes the approach largely insensitive to the specific tone-mapping method used.

5. EXAMPLES


The filtering approach described in this paper was implemented in an open-source C program (see [27]). Figure 11 shows the results of applying this program to two different RADIANCE models of a Washington, D.C., subway station. The top images show renderings of the original models. The middle images show the original renderings filtered to simulate moderate low vision (Snellen acuity of 20/250, Pelli–Robson contrast sensitivity score of 1.0, and color saturation of 40%). The bottom images show the original renderings filtered to simulate severe low vision (Snellen acuity of 20/800, Pelli–Robson contrast sensitivity score of 0.5, and completely unsaturated color). In Fig. 11, the left column is based on a photometrically correct model of the lighting in an actual station. The right column is based on a modification of the original model that adds direct lighting to highlight features such as benches, while reducing indirect lighting to match the original electrical load. The result is an increase of illumination and contrast in the pedestrian area. In the left column, the bench at the lower left of the image is hard to see in the original, and even harder to see as the



Original RADIANCE renderings.



Original rendering, filtered to simulate moderate low vision.



Original rendering, filtered to simulate severe low vision.

Fig. 11. Examples of simulated loss of acuity and contrast sensitivity for a RADIANCE model of a Washington, D.C., Metro station (left column) and the model modified to provide improved lighting (right column).

simulated level of visual impairment increases. In contrast, in the right column the bench continues to be visible even at high levels of visual impairment, providing evidence that residual visual function would be an aid to mobility under the modified lighting condition. This tool provides the ability to explore

options involving the location of luminaires and fenestration, through changes in the reflective properties of surfaces, and through changes in the shape and orientation of a potential hazard or way-finding element. Architects typically make these choices, hopefully with low-vision guidelines in mind, but they

currently have no tool to test the visibility consequences of their choices.

6. DISCUSSION

Our goal is to simulate the loss of visual information associated with reduced acuity and contrast sensitivity as an aid in creating visually accessible architectural spaces. To the extent that our findings generalize to actual low vision and to real-world environments, we can predict the features of such spaces that would not be seen by people with specified levels of reduced acuity and contrast sensitivity. Doing so can provide architects and lighting designers with key information not currently available. This is a challenging task. Predicting the limits of low-vision visibility requires a quantitative analysis of the geometry and photometric information available to the viewer, along with quantitative modeling of the effects of low vision.

There are a number of limitations to the work described here that still need to be addressed. Most importantly, our model does not yet account for glare, absolute luminance and adaptation, or field loss. Glare is a multifaceted and complex phenomenon. While a few automated tools have been developed to assist in quantitatively analyzing the relationship between glare and visibility in architectural spaces (e.g., [28,29]), little is currently known about how to extend these tools to account for significant levels of low vision. Contrast sensitivity decreases at low light levels [30] and is affected by both spatial and temporal adaptation [12,31,32]. While none of these effects have been incorporated into our simulation, it would be relatively easy to incorporate existing techniques if needed.

Our simulation does not explicitly address the impact of field loss. People with central-field loss often adopt a retinal region outside of a central scotoma for fixation, termed the preferred retinal locus (PRL). Typically, acuity and contrast sensitivity are reduced at the PRL [33]. Our model may be useful in simulating the visibility of features viewed at the PRL, but our model does not address the loss of information within a central scotoma or the eye-movement recalibration required to bring features of interest to the PRL. Similarly, people with peripheral field loss may entirely miss seeing targets outside of their restricted field of view. Our model only simulates what is seen when gaze direction brings these targets into view.

The simulation addresses the visibility of targets, but does not simulate the subjective experience of low vision. Images filtered to represent low acuity may appear blurry to a normally sighted viewer, but people with low vision do not necessarily describe their perception as blurry. We also use color as a default in our simulations, but not with the intent of simulating color appearance in low vision. Retinal and other forms of eye disease often distort normal color vision; our simulation is not intended to capture these effects.

Finally, it is important to note that our modeling needs to be validated by testing subjects with actual low vision and measured values of acuity and contrast sensitivity. To date, there is little empirical data on predicting the effects of low vision on the visibility of hazards in real-world situations. Obtaining more such data will be critical to our ability to improve the visual accessibility of architectural spaces.

Funding. National Eye Institute (NEI) (BRP 2 R01 EY017835-06A1).

Acknowledgment. We gratefully acknowledge Eli Peli for providing us with access to the original implementation of the nonlinear filtering described in [6]. Test data for use in evaluating the legibility of filtered Sloan characters were created with the assistance of software provided by Precision Vision (<http://precision-vision.com/>). Erica Barhorst-Cates collected the human subject data reported on in the section on calibration. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

REFERENCES

1. National Eye Institute, "Statistics and data," 2016, <https://nei.nih.gov/eyedata/>.
2. R. Pararajasegaram, "Low vision care: the need to maximise visual potential," *Community Eye Health* **17**, 1–2 (2004).
3. American Optometric Association and Illuminating Engineering Society, "Lighting your way to better vision," paper IES CG-1-09.
4. Illuminating Engineering Society, "Light + seniors: a vision for the future," in *IES Research Symposium I*, March 6–7, 2012.
5. National Institute of Building Sciences, "Design guidelines for the visual environment," 2015, https://c.ymcdn.com/sites/www.nibs.org/resource/resmgr/LVDC/LVDP_Guidelines_052815.pdf.
6. E. Peli, "Contrast in complex images," *J. Opt. Soc. Am. A* **7**, 2032–2040 (1990).
7. E. Peli, "Simulating normal and low vision," in *Vision Models for Target Detection and Recognition*, E. Peli, ed. (World Scientific, 1995), Vol. **2**, pp. 63–87.
8. C. W. Tyler, "Is the illusory triangle physical or imaginary?" *Perception* **6**, 603–604 (1977).
9. M. W. Cannon, "Perceived contrast in the fovea and periphery," *J. Opt. Soc. Am. A* **2**, 1760–1768 (1985).
10. M. A. Georgeson and G. D. Sullivan, "Contrast constancy: deblurring in human vision by spatial frequency channels," *J. Physiol.* **252**, 627–656 (1975).
11. A. B. Watson and A. J. Ahumada, Jr., "A standard model for foveal detection of spatial contrast," *J. Vis.* **5**(6), 740 (2005).
12. G. W. Larson, H. Rushmeier, and C. Piatko, "A visibility matching tone reproduction operator for high dynamic range scenes," *IEEE Trans. Vis. Comput. Graphics* **3**, 291–306 (1997).
13. D. G. Pelli and P. Bex, "Measuring contrast sensitivity," *Vis. Res.* **90**, 10–14 (2013).
14. A. M. Rohaly and C. Owsley, "Modeling the contrast-sensitivity functions of older adults," *J. Opt. Soc. Am. A* **10**, 1591–1599 (1993).
15. S. T. L. Chung and G. E. Legge, "Comparing the shape of contrast sensitivity functions for normal and low vision," *Invest. Ophthalmol. Visual Sci.* **57**, 198–207 (2016).
16. D. G. Pelli, J. G. Robson, and A. J. Wilkins, "The design of a new letter chart for measuring contrast sensitivity," *Clin. Vis. Sci.* **2**, 187–199 (1988).
17. P. G. J. Barten, "Formula for the contrast sensitivity of the human eye," *Proc. SPIE* **5294**, 231–238 (2004).
18. M. Kwon and G. E. Legge, "Spatial-frequency cutoff requirements for pattern recognition in central and peripheral vision," *Vis. Res.* **51**, 1995–2007 (2011).
19. G. E. Legge, G. S. Rubin, and A. Luebker, "Psychophysics of reading—V: the role of contrast in normal vision," *Vis. Res.* **27**, 1165–1177 (1987).
20. A. B. Watson and A. J. Ahumada, "Letter identification and the neural image classifier," *J. Vis.* **15**(2), 15 (2015).
21. J. J. McAnany, K. R. Alexander, J. I. Lim, and M. Shahidi, "Object frequency characteristics of visual acuity," *Invest. Ophthalmol. Visual Sci.* **52**, 9534–9538 (2011).

22. F. Thorn and F. Schwartz, "Effects of dioptric blur on Snellen and grating acuity," *Optom. Vis. Sci.* **67**, 3–7 (1990).
23. P. F. Felzenszwalb and D. P. Huttenlocher, "Distance transforms of sampled functions," *Theory Comput.* **8**, 415–428 (2012).
24. WebAIM, "Visual disabilities: color-blindness," 2013, <http://webaim.org/articles/visual/colorblind>.
25. E. Reinhard, W. Heidrich, P. Debevec, S. Pattanaik, G. Ward, and K. Myszkowski, *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting* (Morgan Kaufmann, 2010).
26. G. W. Larson and R. Shakespeare, *Rendering with Radiance: The Art and Science of Lighting Visualization* (Booksurge LLC, 2007).
27. W. B. Thompson, "Deva-filter source code," <https://github.com/visual-accessibility/deva-filter>.
28. J. Suk and M. Schiler, "Investigation of Evalglare software, daylight glare probability and high dynamic range imaging for daylight glare analysis," *Light. Res. Technol.* **45**, 450–463 (2013).
29. J. Wienold, "Evalglare—a new RADIANCE-based tool to evaluate daylight glare in office spaces," in *3rd International RADIANCE Workshop*, October 11, 2004.
30. J. A. Ferwerda, S. N. Pattanaik, P. Shirley, and D. P. Greenberg, "A model of visual adaptation for realistic image synthesis," in *Proceedings of ACM SIGGRAPH* (1996), pp. 249–258.
31. P. Irawan, J. A. Ferwerda, and S. R. Marschner, "Perceptually based tone mapping of high dynamic range image streams," in *Proceedings of the 16th Eurographics Conference on Rendering Techniques (EGST)* (2005), pp. 231–242.
32. S. N. Pattanaik, J. Tumblin, H. Yee, and D. P. Greenberg, "Time-dependent visual adaptation for fast realistic image display," in *Proceedings ACM SIGGRAPH* (2000), pp. 47–54.
33. S.-H. Cheung and G. E. Legge, "Functional and cortical adaptations to central vision loss," *Visual Neurosci.* **22**, 187–201 (2005).