Predicting Visibility in Designs of Public Spaces

Daniel Kersten, Robert Shakespeare, and
William Thompson

UUCS-13-001

School of Computing
University of Utah
Salt Lake City, UT 84112 USA

February 21, 2013

Abstract

We propose a method for anticipating potential visual hazards during the design phase of an architectural project. In particular, we focus on the needs of people with impaired vision. The method uses a model of human visual loss together with physically accurate rendering of scenes. It avoids the major challenge of automated computer recognition of obstacles, and instead uses easily computable surface geometry data to highlight regions that may require boosts in contrast to improve visibility.

Predicting visibility in designs of public spaces

D. Kersten[a,*], R. Shakespeare[b], W. Thompson[c]

[a]*Department. of Psychology, University of Minnesota, 75 East River Road, Minneapolis, MN 55455, USA*
[b]*Theater Department, Indiana University*
[c]*Computer Science Department, University of Utah*

Abstract

We propose a method for anticipating potential visual hazards during the design phase of an architectural project. In particular, we focus on the needs of people with impaired vision. The method uses a model of human visual loss together with physically accurate rendering of scenes. It avoids the major challenge of automated computer recognition of obstacles, and instead uses easily computable surface geometry data to highlight regions that may require boosts in contrast to improve visibility.

Keywords: low-vision, contrast sensitivity, hazard, design, obstacles

---

Introduction

Good vision depends on high spatial acuity, high contrast sensitivity, and an intact visual field [15]. Losses in any of these capacities creates difficulties in interpreting and navigating an environment. Public spaces, such as lobbies, offices, restaurants, and stores, pose unique challenges for a person with vision loss. Such spaces are often encountered by the person for the first time, and present unexpected hazards such as steps, ramps, doorways, glass walls, benches, and other obstacles. In addition, growing demands on energy conservation constrain illumination design[1], which in turn impacts visibility. Although there is no engineering solution to compensate for poor acuity per se, design choices can be made to boost brightness contrast in scene regions crucial to safe navigation, while at the same time avoiding sources of misleading contrast. We propose a method to assist designers in locating potential sources of visual ambiguity during the design phases of a project. The method builds on past work in characterizing human visual sensitivity together with physically realistic 3D graphics modeling and rendering.

Contrast, defined more precisely below, is a measure of the spatial change in intensity over some specified region in an image. The location and amount of contrast is a consequence of a complex interaction of illumination with surface geometry, material, and viewpoint. The type and arrangement of illumination sources is important for both esthetics and practical function. However, the effects of illumination can be hard to anticipate, requiring substantial experience augmented by computationally intensive, physically realistic simulations [20]. Added to this difficulty is the challenge for a designer with normal vision to anticipate the effects of illumination for persons with low-vision. Low light or poor fixture placement can result in the failure to see or accurately interpret potential hazards. Misinterpretations can also occur in well-lit scenes with high contrast. Cast shadows and reflections can create false interpretations of a scene, such as mistaking a shadow for a geometrical change such as a hole, step, or a material change (e.g. from dry to wet). Unfortunate material and shape choices can also create false perceptions of geometrical change, such as a step, when there is none. Such mistakes can occasionally be seen with normal high visual acuity (Figure 1). Such real-life "illusions" become increasingly problematic with clinical conditions that lead to a loss of spatial resolution and contrast, as well as normal vision under low light, and with age. We will call failure to detect a critical feature or obstacle a "miss", and incorrectly detecting one as a "false positive". Misses are often more costly than false positives; however, any kind of mistake can have unforeseen negative consequences. Figure 2 illustrates how light streaming through a window can be misinterpreted as a step, which contributes to missing the true step which is slightly higher in the visual field. We describe a tool to anticipate visibility problems, including both misses and false positives, during the design stages of a project.

For practical reasons, we focus on the effects of lighting under static, monochrome viewing conditions, although the principles and tools are extensible to motion and color. We propose a method by which visibility-related mistakes may be anticipated by lighting and material choices that maximize contrast in the right places, and minimize contrast in the wrong places. But how much contrast is enough or too much? And what is meant by the right or wrong places?

---

[*] Corresponding author.

Email addresses: `kersten@umn.edu` (D. Kersten), `shakespe@indiana.edu` (R. Shakespeare), `thompson@cs.utah.edu` (W. Thompson)

[1] See for example, http://energy.gov/public-services/building-design, "Energy Efficient Buildings Hub"; http://www.eebhub.org, and IESNA/ANSI RP-28-2007: Lighting and the Visual Environment for Senior Living. New York: Illuminating Engineering Society of North America (2007).

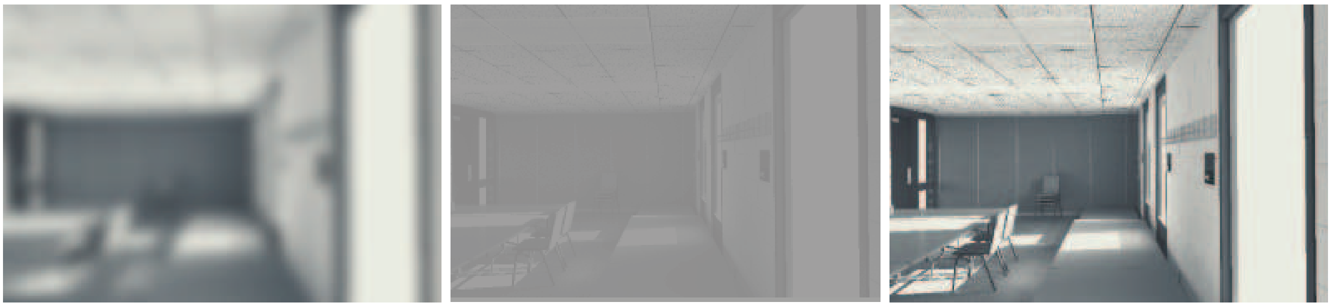

Figure 2: **Loss of spatial resolution and contrast**. **A.** Left panel illustrates a low spatial resolution image of a room with table and chairs on the left, and a raised ramp on the right. The ramp may or may not have a step at the end. The small horizontal bright patch in the lower right portion of the image, middle of the ramp, might be mistaken for the rise of step brightly illuminated from the side. **B.** Middle panel illustrates low contrast, which also contributes to errors in perceptual judgments. **C.** Right panel shows the scene with high contrast and high resolution. With normal vision, it is easy to correctly parse the bright region as due to bright illumination on a flat floor rather than due to a change in the surface orientation of a step.

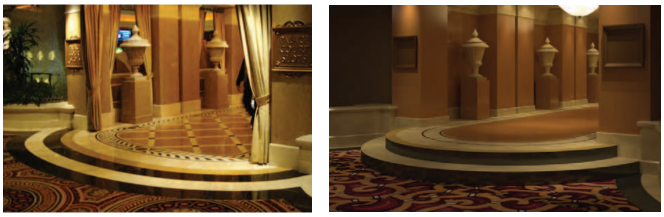

Figure 1: **Apparent and real steps.** The left panel shows a image of an actual scene with an apparent set of steps. These steps are illusory, due to material changes that mimic surface orientation and curvature changes associated with the step runs and rises. The right panel shows a simulation with model steps with the appropriate surface orientation changes. Additional occlusion cues, including corners provide strong cues for steps.

Answers to the first question come from several decades of research showing how visibility as a function of contrast and spatial frequency affects various tasks such as reading, object recognition, and navigation [15]. A key measure that takes both contrast and resolution loss into account in the *contrast sensitivity function* (CSF)[4] (Figure 6). The CSF measures the minimum amount of contrast required to detect a sine-wave grating[2]. Contrast sensitivity varies according to individual vision losses, and decreases as a function of light level in normal vision as well [4]. Our physically realistic renderings specify the predicted luminance values at each image location. Given a prescribed CSF and a physically realistic simulation of a scene, one can flag image regions whose contrast cannot be detected, and thus pose a risk for missing an obstacle.

However, as illustrated above in Figure 2, one can have high contrast regions that are easily detected, but that are nevertheless ambiguous as to cause due to lack of sufficient image resolution. While there may be sufficient information to recognize the scene given only low spatial frequency content, i.e. by its "gist", this knowledge can be insufficient for discerning the devil in the details, e.g. "I'm in a lobby, but is that a step? Is it up or down?". In a blurry image, one can lose contrast information important for precise scene parsing, including localization and identification of obstacles. A major challenge in applying measures of visibility in natural scenes is that one needs to specify the regions of an image that are important. Although there are existing results that could be applied to automating an analysis of overall contrast loss in rendered scenes (cf. [17], and see **Implementation** section below), one would like to have a method to automatically select regions corresponding to objects relevant to the person's task.

Answers to the second question–deciding the "right places"–requires knowledge of the 3D structure of a scene together with constraints based on the person's task. We first discuss 3D structure. Scene interpretation and safe passage requires detection and recognition of potential hazards. A basic premise is that accurate and reliable decisions critical for navigation rest on the visibility of changes in the depth and orientation of surfaces in the environment. Changes in intensity that coincide or correspond with surface transitions such as a step, a post boundary, or a ramp are informative for scene understanding and safe navigation. Changes primarily due to illumination are more difficult to interpret and can confound detection and recognition of objects because they are not stable with respect to surface features of objects [10, 6]. This is illustrated in Figure 3. Note that it is almost impossible to accurately interpret the scene from patches in the image whose contrast variation is primarily due to illumination (upper left panel of Figure 3). However, patches whose contrast source is surface orientation and/or depth change (3, lower right) provide better information for interpretability.

If one could automatically detect task-relevant regions in the image, those regions could be checked for adequate or misleading contrast. The problem is that, despite considerable advances in computer vision, an automatic solution to parsing an image into its causes, i.e. material, depth, orientation, and illumination in the scene remains elusive (cf. [3]). Further, parsing an image into regions that correspond to object or surface categories, such as "benches" or "stairs" also has no general, robust automatic solution. To appreciate the problem in our context, there is a huge space of possible obstacles. Consider just steps.

---

[2] A grating is an image pattern in which luminance is modulated sinusoidally about a fixed mean level. The contrast of a grating is defined to be: $(L_{max} - L_{mean})/L_{mean}$, where L represents luminance.

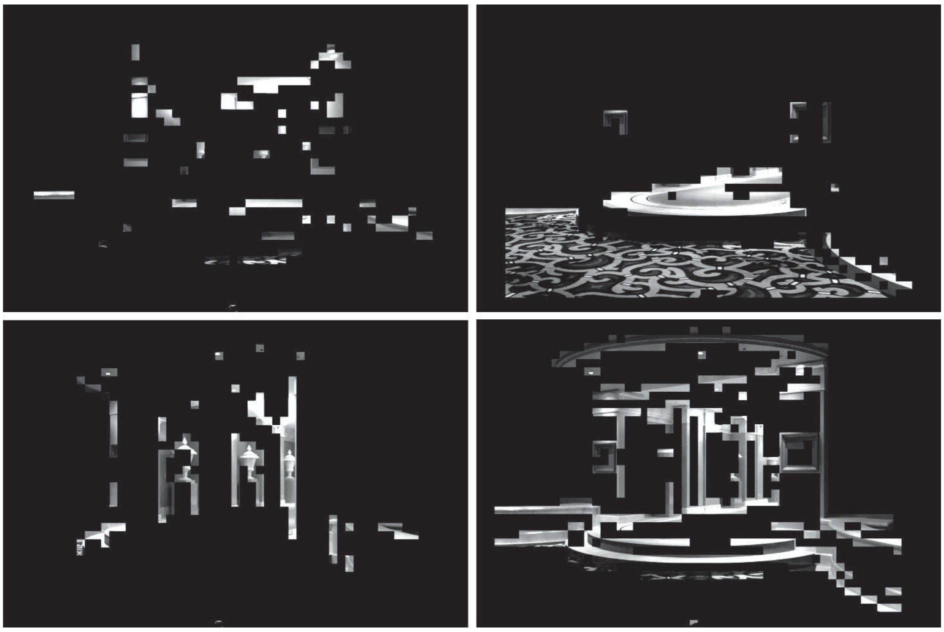

Figure 3: **Four causes of contrast.** The image from Figure 1 (right panel, with actual steps), was broken up into subpatches and classified into four groups based on the dominant cause of contrast in the patch: illumination, material, surface occlusion (i.e. abrupt depth change), and surface orientation change (i.e. concave or convex surface intersections). The classifications were based on subjective judgment based on viewing the full color original image. The upper left panel shows patches from a scene whose contrast is primarily due to spatial changes in illumination. The upper right panel shows patches of the image in Figure 1 whose contrast is primarily due to material change. Material changes often coincide with depth changes. The lower left panel shows patches whose contrast is primarily due to abrupt depth changes, i.e. to occlusion of one surface by another. The lower right panel shows patches whose contrast is primarily due to changes in surface orientation, such as going from the run to rise of a step.

Figures 2 and 3 show straight and convexly curved steps, but there are many other variations. Even with building code requirements, steps come a wide variety of forms. They can be concave, convex, and non-rectangular (winders), have riser and treads adjacent to walls or not, have varieties of balusters and newels, tread nosing or not, risers or not, treads can be mesh, and so forth. Compound the large variation in potential obstacle types with the infinite variation in the types of image features corresponding to a given type, due to viewpoint, and materials, and one begins to appreciate the problem. In our study, we circumvent this problem by exploiting the fact that designers work with scene descriptions, including object geometry and layout using computer aided design (CAD) from the start. Further, advances in physically realistic rendering of 3D spaces provide physically accurate images that can be used as input to models of human visual loss. CAD design files provide a potentially rich source of scene information, including object labels, shapes, positions, and materials. A major problem in using this information is that it comes in many formats that are software-dependent, and with labels and groupings that are designer-dependent. Extracting and representing such information in consistent standard format is a challenge by itself. Our solution is to focus on extracting information from the depth or "range" data. Range data is straightforward to obtain using rendering environments such as *Radiance*. Further, a representation of depth change correlates well with object boundaries (i.e. where there is often a depth discontinuity). Changes in surface orientation correspond with key shape features, such as the transition from run to rise in a step. Thus, a measure of depth and/or orientation change provides a generic intermediate proxy for objects, surfaces, and potential obstacles. We refer to such geometrical information as "ground-truth" and below propose a simple method for using such "ground-truth" to highlight problem regions in an image.

How about task? If one knew where a person is likely to go and where they will look, we could further narrow the selection of image regions. One could have a preselected list of labeled surfaces, such as steps and hallways that afford actions, and in principle even simulate potential actions with a space to check for hazards. Our solution is more modest–we use the available scene 3D depth data from a proposed design to automatically select regions of interest that are within a volume reachable by walking person, i.e. below a maximum height (e.g. 8'), and within a fixed number of steps (e.g. four steps, 3' each, or 12'). Alternatively, our approach also allows a designer to interactively select a specific region of interest (e.g. "these particular steps" or "this hallway"), and have the program automatically analyze this region over a range of lighting conditions for visibility problems.

In summary, our approach is to automatically identify surface geometry transitions important for successful scene interpretation and navigation and then check whether there is sufficient change in local contrast. But we also go the other way, to check whether regions with large contrast occur in the absence of geometric transitions, such as shadows across a floor. Figure 4 illustrates the main components of the procedure. An image rendering and a map of surface distances is obtained from a design's database. We refer to this latter map as "ground truth".

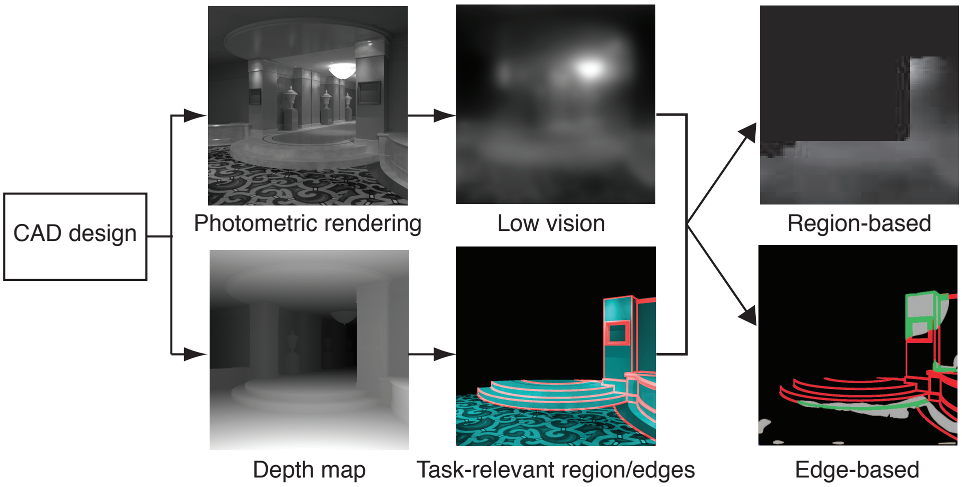

Figure 4: **Method.** The scene model contains information regarding objects, materials, layout, illumination and viewpoint. The scene model is used to produce an image rendering and a map of important locations extracted from the surface models in the scene. In the implementation tested, these regions are of two types: unobstructed navigable floor space within a criterion distance (in green), and regions with potential obstacles (in red). Surface transitions include discontinuities in depth, orientation, and curvature. Regions of interest can be constrained based on general a priori constraints such as "all regions below 8 feet and within 10 feet", or be targeted by a user such as "these particular steps" or "this walkway". Information loss is computed by filtering the rendered image with a filter representing the losses of a low-vision observer. The regions of interest are compared with the image representing image loss to check for ambiguities that might pose a risk. These ambiguities include predictions of both edge misses (red) and false positives (green).

In the implementation described below, the geometrical transitions are determined by detecting discontinuities in depth and orientation of surfaces from the range data. Task-relevant regions of interest are then selected based either on a priori navigable space constraints, or selection by the designer. Information loss in the image is computed by filtering the rendered image with a spatial filter that characterizes visual deficits in acuity and contrast. Information loss is represented by a low vision response map. The regions of interest are then compared with the low vision response map to assess potential risk. We consider two types of output and two types of application. The first output is simply a measure of overall contrast variation in the task-relevant region of interest. The second is local, and flags geometrical edges whose image contrast is below a criterion level (misses) or regions above criterion that have no edges (false positives). In an online interactive application, predictions of edge misses and false positives could be flagged for inspection by a designer or used in a summary score. In an automated application, a batch of files (e.g. ranging over a series of illumination conditions) is submitted, and the program returns visibility scores (e.g. contrast response variation or miss rate) for each image.

Implementation

Determining ground-truth and task-relevant regions

We use the physically accurate rendering model, Radiance, to produce 1) High Dynamic Range (HDR) image files that accurately represent the luminance intensity of the actual scene; 2) a registered range map from the z-buffer, which represents the distances from the viewer to surfaces in the scene; we also use maps representing x, y, and z components of surface normals at each point. In addition, we can also output the vertical and horizontal distances of surface points from the ground plane, useful for restricting relevant portions of the scene to a given task-specification, such as "all surface points reachable within 12 feet and below 8 feet in height".

Given a range data file $r(x, y)$ that specifies distance as a function of image coordinates x, y, we identify salient locations by finding where there are significant changes in depth and orientation. We have tried several edge detection methods, and found that the Canny edge detector [9], operating on the range map r, and on the surface normal maps provides reasonable edge maps of ground truth.[3] The result, combining edges over range and surface normal changes, is represented by O_E.

$$e(x, y) = O_E(r(x, y))$$

where $e(x, y) = 1$ and $e(x, y) = 0$ mark the presence and absence of an edge at point (x, y), respectively.

To pick out regions of the image that are important for navigation, we specify an operator, $ROI(\cdot)$, that selects an image domain reflecting a region of interest based on task constraints.

---

[3] There is a substantial body of computer vision research on edge detection in image intensity data, e.g. [13, 16]. No solution is perfect, and the problems are well-known. Two basic problems are selecting the appropriate spatial scale and spatial noise. If the scale is too coarse, important edges are missed and if the scale is too fine, noise introduces spurious edges. In the range data from architectural designs, noise could arise form fine-grain textural detail, such as foliage. So far, we have found that scale choice is less problematic because surfaces in designs tend to have smooth representations with sharp distinctions between objects. Further, range data does not have confounding factors, such as specularities and shadows, that complicate edge extraction from image intensity data. An additional advantage is that the range maps have an absolute spatial scale, which could be exploited to determine an appropriate degree of smoothing if needed.

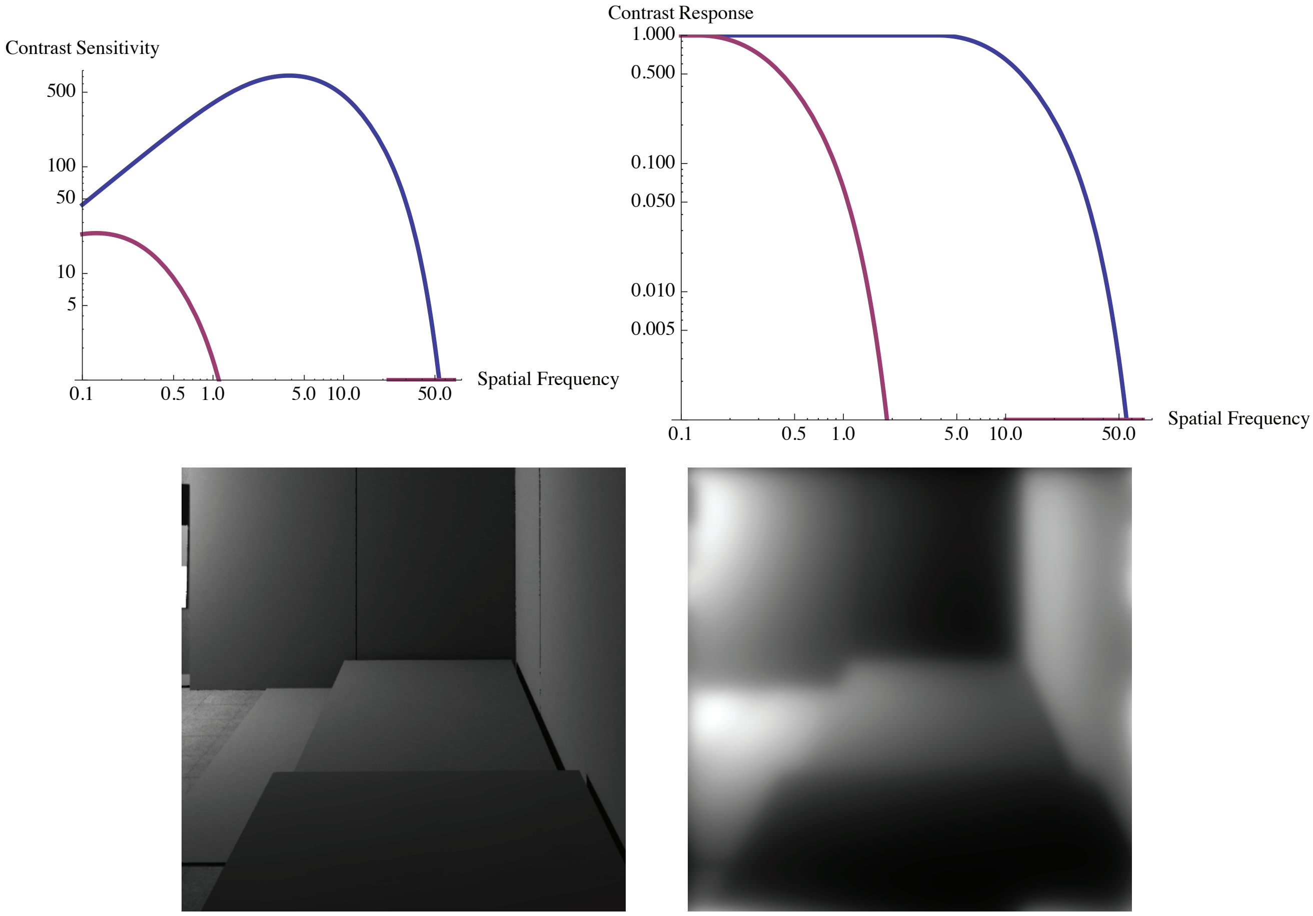

Figure 6: **Filter model of low-vision loss.** The **upper left** panel shows the contrast sensitivity for normal vision (upper curve) and a low-vision observer (lower curve). Contrast sensitivity is the reciprocal of contrast threshold. Contrast threshold is the minimum detectable contrast. Typically as spatial frequency increases (and the period of the grating gets smaller), the visual response becomes smaller until the white and dark bars of the grating can no longer be resolved even at 100% contrast, corresponding to the lower bound on contrast sensitivity of 1. Near normal visual response, corresponding to 20/20 vision, extends to about 50 $cycles/degree$. The lower curve shows the response of a low-vision model observer with 20/600 vision. The curves are based on the Barten model of the contrast sensitivity function (CSF) [4]. The **upper right** panel shows normalized human visual contrast response (CRF) to a sine-wave grating as a function of spatial frequency in cycles per degree of visual angle. In normal vision, the CSF is band-pass (see left panel), reflecting that response to low contrasts at low frequencies is suppressed. Although there are cases of low-vision with abnormal low-frequency response, we wanted a filter that represented the common, and critical loss at high frequencies. The upper right figure shows a modification in which the response left of the peak CSF are set to one. We assume radial symmetry (e.g. ignoring aberrations such as astigmatism) and use the 1D CRF to represent constant loss in each polar direction. The **lower left** and **lower right** panels show the filtered images of a room with a down-step. These images represent the information available for normal vision (left) and a low-vision observer (right). The image on the right is the convolution of the left image with a 2D version of the spatial filter shown in the upper right, representing the loss experienced by an observer with 20/600 vision.

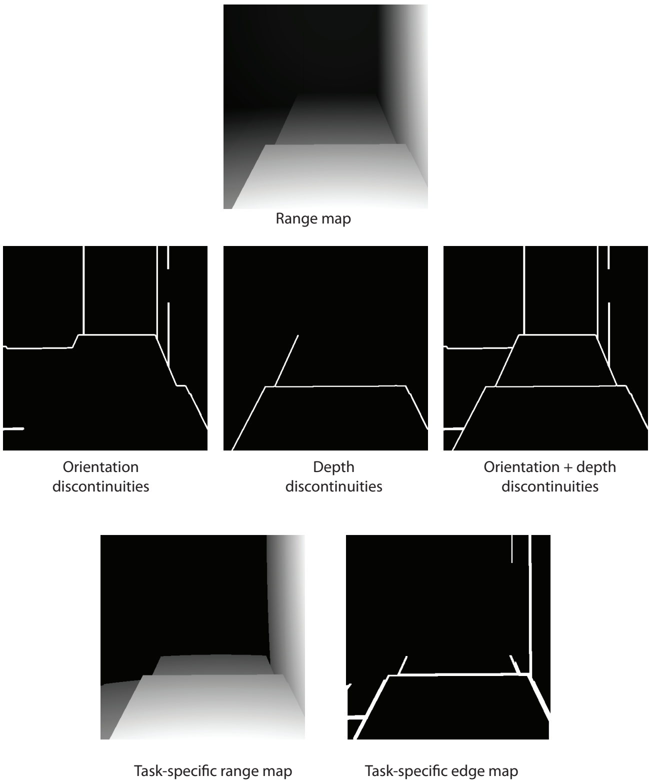

Figure 5: **Ground truth.** The **top** panel shows original range data, which is a map of the distance from the viewer indexed by position in the image. The **middle** panels show the extracted orientation, depth, and combined orientation and depth discontinuities, represented by $e(x, y)$. The **lower** panels show the range and edge maps restricted to the navigable region of interest. The metrics and annotations use the task-specific edge map, $e_t(x, y)$, together with a model of vision loss to produce annotations and feedback for the designer.

Here, the ROI is simply determined by those image locations that are reachable by a person, e.g. height $h(x, y)$ lower than h_c feet, and distances $d(x, y)$ within a predefined value, d_c from the viewer:

$$ROI(\cdot) = \begin{cases} \cdot, & \text{if } h(x, y) < h_c \text{ and } d(x, y) < d_c \\ Null, & \text{otherwise} \end{cases}$$

For interactive applications, a designer could select a region of concern for inspection, e.g. a pathway leading from the viewpoint to an intermediate destination in the domain of the image. Figure 5 shows examples of ground truth maps, and their restrictions to task-relevant regions.

The end-product for ground-truth calculation is given by:

$$e_t(x, y) = ROI(e(x, y)).$$

The bottom right panel of Figure 5 shows an example of e_t (right).

Modeling loss of spatial resolution

Although, the forms of visual loss are extremely diverse, we focused on the loss of spatial resolution, or more precisely the loss of contrast information at high spatial frequencies. Loss of spatial resolution or "acuity" is commonly measured as Snellen acuity, but a more general measure is the contrast sensitivity function or CSF (Figure 6). The diversity of vision loss together with the enormous range of scene variations in designs, underscores the need to initially keep the model of vision loss as simple as possible. Thus at this point, we ignore possible differential loss in multiple visual channels in human vision and the effects of large changes in illuminance, resulting for example in glare. We describe a simple "single-channel" model of visual loss. This provides a base model on which future elaborations can be developed and compared.[4]

The HDR image (from Radiance) is first converted to luminance contrast:

$$I_{input}(x, y) = \frac{HDR(x, y) - HDR_{avg}}{HDR_{avg}}$$

This is a "global" model of contrast. An image representing information loss, $I_{acuity}(x, y)$ is then computed by convolving the luminance contrast image with a "blur" kernel whose parameters are determined by the contrast response function:

$$I_{acuity}(x, y) = CRF_{acuity} * I_{input}$$

The shape of the *CRF* is determined by the CSF for a specified Snellen acuity (Figure 6).

As with the ground-truth maps, we use the ROI operator to restrict the region of image information loss to task-relevant areas:

---

[4] For example, a single-channel model doesn't take into account the dependency of spatial filtering on local luminance or multiple channel filtering by the human visual system, cf. Peli [17].

$$I_t(x, y) = ROI(I_{acuity}(x, y))$$

In the next section, we show how one can use I_t and e_t to provide designer feedback through graphical annotations and metrics.

Metrics and annotations

Region-based. Imagine now that a designer would like feedback demonstrating the information available to a low-vision observer with a pre-specified level of acuity. A simple form of feedback is to show the entire response image I_{acuity} (or I_t). However, a major problem in relying on designer inspection is that, in contrast to a low-vision observer, the designer has considerable familiarity with the scene and likely the rendered views as well. Prior knowledge is known to strongly bias the interpretation of a blurry and otherwise ambiguous image [8]. Further, visibility on the display itself depends on calibration. One solution is to provide the designer with objective feedback in the form of a visibility score which be used to evaluate the relative merits of various rendering choices. There is a very large and active literature on image quality metrics, motivated early on by the desire to reduce visible artifacts induced during image compression [1, 21, 22]. These metrics focus on measuring the difference between two images which can be used to predict discriminability for normal human observers. Almost nothing, however, is known about visibility metrics that reflect interpretability of a single natural image for low-vision observers. Further, the ultimate purpose of a metric is not to predict visibility per se for a particular image for a particular standard observer, but rather to provide a summary score that is correlated with functional visibility in the space once built. A simple way to score the overall visibility of the task-relevant region of I_{acuity} is to calculate the standard deviation of I_t (or root-mean-squared-contrast, "RMSC"):

$$RMSC = std(I_t).$$

Intuitively, this score captures the overall effective contrast variation within the task-relevant region of the blurred image. The number decreases with both loss of detail and contrast. High-values indicate large variations in response, and a uniformly constant response image would have a value of zero. Evaluation of the validity and merits of this or alternative global, image-based metrics will require systematic research.

Edge-based. Another way to reduce the effects of designer familiarity is to take advantage of our ability to calculate the task-relevant edge map, e_t. We can provide visual feedback regarding geometrical edges in the scene which might be missed or misinterpreted. Given a criterion threshold level, c, on the gradient of the response, we find those edges in the range map that have a contrast response change below criterion. This results in a map of misses:

$$misses(x, y) = \begin{cases} 1, & \text{if } |\nabla(I_t)| < c \text{ and } e_t(x, y) = 1 \\ 0, & \text{otherwise} \end{cases}$$

We can also highlight edge false positives:

$$falsepositives(x, y) = \begin{cases} 1, & \text{if } |\nabla(I_t)| > c \text{ and } e_t(x, y) = 0 \\ 0, & \text{otherwise} \end{cases}$$

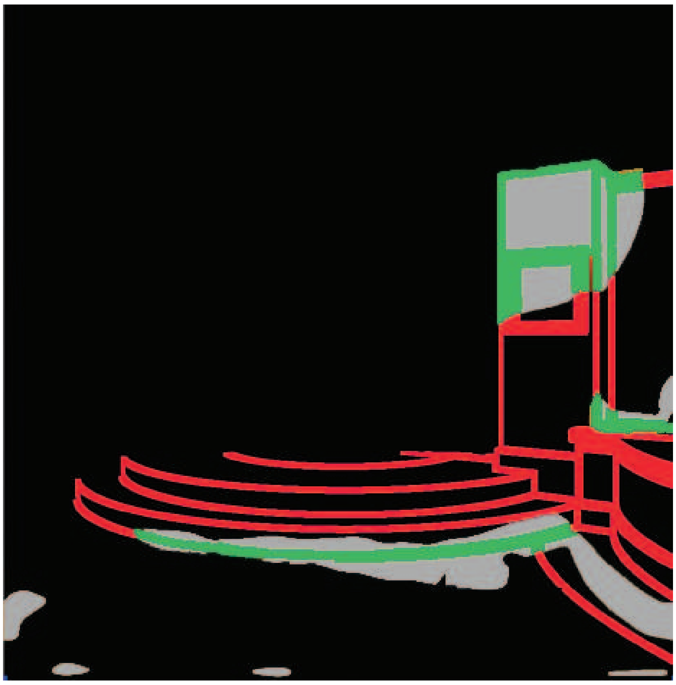

Figure 7: **Edge misses and false positives**. Misses, or edges where the response gradient falls below threshold are shown in red. False positives, i.e. edges whose contrast change is above threshold, are shown in green.

Given an assumption of the relative costs of the two types of errors, one could calculate a summary score based on both. For example, if the costs of a miss and false alarm were equal, a simple summary score is d', which is related to the area under the Receiver Operating Characteristic or ROC curve [11] . One advantage of such a measure is that it is independent of criterion, c. The disadvantage is that the costs of misses could be much higher than false alarms (e.g. missing a step vs. deciding there is a step when the floor is actually flat as in Figure 1).

Future Work

Software tool

We have provided what we think is a practical and compelling method that could be built into an interactive tool for designers to evaluate the relative impacts of design choices for people with impaired vision, or under low lighting.[5] It has the advantages of 1) reducing the effect of the unavoidable subjective biases a designer will have in trying to interpret a blurry image as a low-vision person might; 2) providing the basis for geometry-based metrics that could be used for interactive feedback, or the semi-automatic batch evaluation of a large space of

---

[5] The Barten model for the human CSF also takes into account loss of contrast sensitivity as a function of overall mean light level.

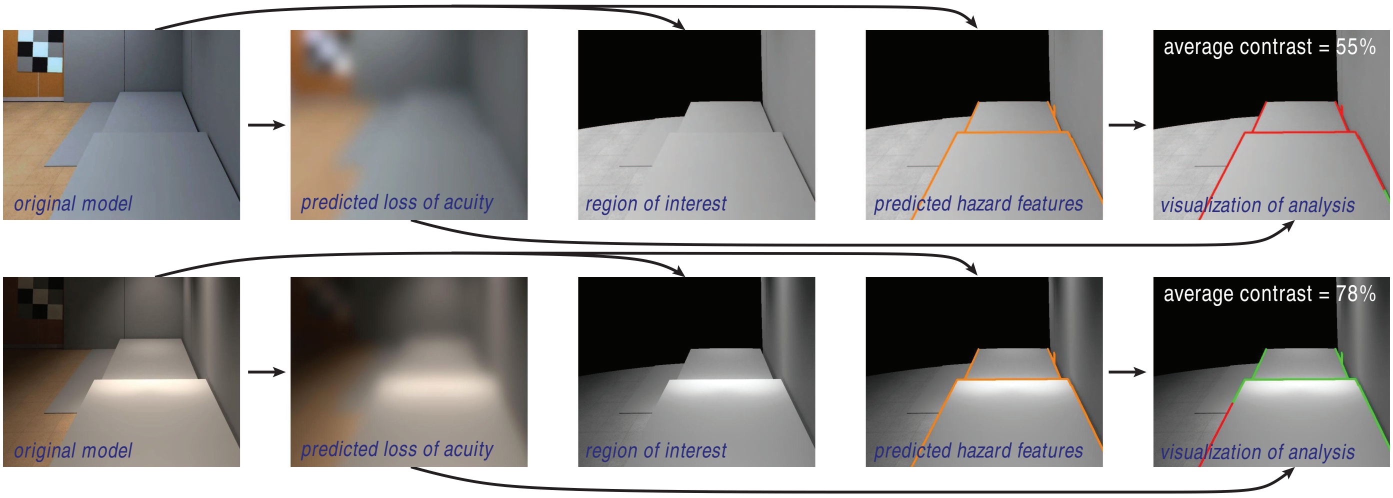

Figure 8: **Upper panel** shows a processing pipeline for step down with overhead lighting: Luminance image → region-of-interest →potential hazards within region-of-interest → visibility estimate of potential hazards. **Lower panel** shows the processing pipeline for step down with focused downlights to improve visibility. In the right panel, the green lines are hazard features predicted to be visible for someone with 20/600 acuity, while the red lines indicate misses, i.e. features predicted to be below the visibility threshold.

design choices (e.g. as in a daylight sequence in which several dozen images might be generated from a particular viewpoint, but over different times of day.).

Figure 8 illustrates a processing pipeline used for a proof-of-concept interactive tool. First, Radiance software is used to generate an accurate, high dynamic range luminance image for a viewpoint chosen by the designer, which is then processed to simulate a specified loss of acuity. The tool would provide a visualization of the surviving information, demonstrating the challenges to visibility. A second visualization shows the results of a quantitative analysis of the scene, indicating features that are likely to be invisible at the specified acuity level. The right panels of the upper and lower panels of Figure 8 illustrate two of the measures that can be computed, involving average contrast over a hazard and an estimation of visible (green) and invisible (red) geometric features.

Validations

Considerable work needs to be done, however, to refine and validate metrics and the utility of miss/false alarm annotations. One practical way of proceeding is to use low-vision human subjects (and/or normally sighted subjects with filtered images representing reduced acuity and contrast sensitivity) to label and trace object boundaries in images on a computer display calibrated for contrast and appropriate angular subtense to reflect a given viewing distance. Performance can be compared with the misses and false alarms produced by the method described above to test, and determine algorithm parameters. Recognition performance can also quantified as confusion matrices for targets, and by tracing accuracy scores for the boundaries, as in [5], to test for correlations between region-based and edge-based metrics and human performance. A key issue is whether the parameters for a visibility metric derived from one set of test images generalize to other contexts. Experimental data is essential to test and as needed improve the low vision filter models, and refine the algorithms for automatically computing visibility, in order to choose, and fine-tune the most useful visibility metrics.

Acknowledgements

We thank Shane Hoversten for help with code development and analysis, and Paul Beckman for help with models of visual loss. This research was supported by NIH R01 EY017835.

References

[1] Ahumada A. J. , Jr. (1993) Computational image quality metrics a review, SID International Symposium Digest of Technical Papers, vol. 24, pp. 305-308, Playa del Rey, CA.

[2] Alpert, S., Galun, M., Brandt, A., & Basri, R. (2012). Image Segmentation by Probabilistic Bottom-Up Aggregation and Cue Integration. IEEE Transactions on Pattern Analysis and Machine Intelligence, 34(2), 315327. doi:10.1109/TPAMI.2011.130

[3] Barron, J. T., & Malik, J. (2012). Shape, Albedo, and Illumination from a Single Image of an Unknown Object. CVPR, 18.

[4] Barten, P. G. J. (2004). Formula for the contrast sensitivity of the human eye. In Y. Miyake & D. R. Rasmussen (Eds.), (Vol. 5294, pp. 231238). Presented at the Image Quality and System Performance. Proc of SPIE-IS&T Electronic Imaging,Proceedings of SPIE.

[5] Brady, M. J., & Kersten, D. (2003). Bootstrapped learning of novel objects. Journal of Vision, 3(6), 413422.

[6] Braje, W. L., Legge, G. E. & Kersten, D. (2000). Invariant recognition of natural objects in the presence of shadows. Perception, 29, 383-98.

[7] den Brinker, B. P., & Daffertshofer, A. (2005, September). The IDED method to measure the visual accessibility of the built environment. In International Congress Series (Vol. 1282, pp. 992-996).

[8] Bruner, J. S., & Potter, M. C. (1964). Interference in Visual Recognition. Science, New Series, 144(3617), 424425.

[9] Canny, J. (1986). IEEE Xplore Download. Pattern Analysis and Machine Intelligence.8, 6, 679-698.

[10] Cavanagh, P. (1991). Whats up in top-down processing. In A. Gorea (Ed.), Representations of vision: Trends and tacit assumptions in vision research, pp. 295304.

[11] Green, D. M., & Swets, J. A. (1966). Signal detection theory and psychophysics. New York, Wiley.

[12] Kersten, D., Hess, R. F., & Plant, G. T. (1988). Assessing contrast sensitivity behind cloudy media. Clinical Vision Science, 2(3), 143158.

[13] Konishi, S., Yuille, A. L., Coughlan, J. M., & Zhu, S. C. (2003). Statistical edge detection: Learning and evaluating edge cues. Pattern Analysis and Machine Intelligence, IEEE Transactions on, 25(1), 5774.

[14] Legge, G. E., Yu, D., Kallie, C. S., Bochsler, T. M., & Gage, R. (2010). Visual accessibility of ramps and steps. Journal of Vision, 10(11):8, 119.

[15] Legge, Gordon E. (2007) Psychophysics of reading in normal and low vision. Mahwah, NJ, US: Lawrence Erlbaum Associates Publishers.

[16] Maini, R., & Aggarwal, H. (2009). Study and comparison of various image edge detection techniques. International Journal of Image Processing (IJIP), 3(1), 111.

[17] Peli, E. (1990). Contrast in complex images. J Opt Soc Am A, 7(10), 2032-2040.

[18] Shakespeare, R.A. (2011). DEVA-Automated Visual Hazard Detection. Presented at the 10th Annual International Radiance Workshop, Lawrence Berkeley National Labs, U.C. Berkeley

[19] Shakespeare, R.A. (2012). Towards Automated Visibility Metrics. Presented at the 11th Annual International Radiance Workshop. Royal Danish Academy of Fine Arts, School of Architecture, Copenhagen.

[20] Ward, G., & Shakespeare, R. (1998). Rendering with radiance: The art and science of lighting visualization. San Francisco, CA: Morgan Kaufmann Publishers, Inc.

[21] Watson A. B., (1993) Digital Images and Human Vision. MIT Press.

[22] Wang, Z., & Bovik, A. C. (2009). Mean squared error: love it or leave it? A new look at signal fidelity measures. IEEE Signal Processing Magazine, 26(1), 98117. doi:10.1109/MSP.2008.930649