On the Limits of Resolution and Visual Angle in Visualization

CHRISTOPHER G. HEALEY and AMIT P. SAWANT, North Carolina State University

This article describes a perceptual level-of-detail approach for visualizing data. Properties of a dataset that cannot be resolved in the current display environment need not be shown, for example, when too few pixels are used to render a data element, or when the element's subtended visual angle falls below the acuity limits of our visual system. To identify these situations, we asked: (1) What type of information can a human user perceive in a particular display environment? (2) Can we design visualizations that control what they represent relative to these limits? and (3) Is it possible to dynamically update a visualization as the display environment changes, to continue to effectively utilize our perceptual abilities? To answer these questions, we conducted controlled experiments that identified the pixel resolution and subtended visual angle needed to distinguish different values of luminance, hue, size, and orientation. This information is summarized in a perceptual display hierarchy, a formalization describing how many pixels—*resolution*—and how much physical area on a viewer's retina—*visual angle*—is required for an element's visual properties to be readily seen. We demonstrate our theoretical results by visualizing historical climatology data from the International Panel for Climate Change.

Categories and Subject Descriptors: H.1.2 [**Models and Principles**]: User/Machine Systems—*Human information processing*; I.3.3 [**Computer Graphics**]: Picture/Image Generation—*Viewing algorithms*; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Psychology*

General Terms: Experimentation, Human Factors

Additional Key Words and Phrases: Hue, orientation, luminance, resolution, size, visual acuity, visual angle, visual perception, visualization

ACM Reference Format:

Healey, C. G. and Sawant, A. P. 2012. On the limits of resolution and visual angle in visualization. ACM Trans. Appl. Percept. 9, 4, Article 20 (October 2012), 21 pages.

DOI = 10.1145/2355598.2355603 http://doi.acm.org/10.1145/2355598.2355603

1. INTRODUCTION

Scientific and information visualization convert large collections of strings and numbers into visual representations that allow viewers to discover patterns within their data. The focus of this article is the visualization of a *multidimensional* dataset containing m data elements and n data attributes, $n > 1$. As the *size* m and the *dimensionality* n of the dataset increase, so too does the challenge of finding techniques to display even some of the data in a way that is easy to comprehend [Johnson et al. 2006]. One promising approach to this problem is to apply rules of perception to generate visualizations that

Authors' addresses: C. G. Healey, Department of Computer Science, 890 Oval Drive #8206, North Carolina State University, Raleigh, NC 27695-8206; email: healey@csc.ncsu.ed; A. P. Sawant, NetApp RTP, Research Triangle Park, NC, 27709.

DOI 10.1145/2355598.2355603 http://doi.acm.org/10.1145/2355598.2355603

build on the strengths of human vision. Defining perceptual guidelines to "take full advantage of the bandwidth of the human visual system" has been cited as an important area of current and future research in visualization [Johnson et al. 2006; McCormick et al. 1987; Smith and Van Rosendale 1998; Thomas and Cook 2005].

Unfortunately, defining visual bandwidth is almost certainly not as simple as computing a fixed bits-per-second result. Based on previous work in human perception, we hypothesize that the abilities of our visual system, properties of the display environment, and the particular visualization strategy being employed combine to define limits on a visualization's information carrying capacity. A minimum number of pixels—*resolution*—and subtended physical area on the retina—*visual angle*—are required to perceive different values of a visual feature. More importantly, these limits can easily be exceeded in a typical visualization. Because the specific resolution and visual angle limits we seek have not yet been reported, we conducted a series of psychophysical experiments to identify the boundaries for properties of color and texture. Our initial studies focus on foveal vision, the area of acute vision centered about the current gaze direction. By formalizing and integrating our results with existing perceptual knowledge, we can design multidimensional visualizations that harness the strengths and respect the limitations of human visual processing within a given display environment.

Issues critical to designing effective visualizations include the following.

(1) *Resolution,* the total number of pixels allocated to an on-screen object or element needed to produce just noticeable differences in its visual appearance.
(2) *Visual angle,* the physical size, measured in degrees of subtended visual angle, that an object or element forms on the viewer's retina to produce just noticeable differences in its visual appearance.
(3) *Viewing parameters,* the amount of data being displayed, its on-screen format, and the speed at which it is shown to the viewer.
(4) *Visualization technique,* the type and number of visual features being displayed, and the methods used to map data values to their corresponding visual representations.
(5) *Data properties* of the dataset being visualized and *analysis tasks* being performed by the viewer.

To date, significant effort has been dedicated to items four and five, that is, how properties of the data and the viewer's analysis needs affect different visualization designs. Less effort has been devoted to studying the first three items. Many visualization techniques simply assume that sufficient resolution, visual angle, and time will be available for viewers to comprehend the images being produced. We study the first three items in the context of multidimensional visualization. Understanding how each property affects image comprehension would allow us to perform a number of important tasks during visualization design:

(1) provide the fundamental knowledge needed to construct perceptually effective multidimensional visualizations;
(2) verify whether the display device and the human visual system meet the requirements of a given visualization algorithm;
(3) characterize to what extent a visualization technique saturates "visual bandwidth;" this can be done in ways that are more sophisticated than simply calculating how many data elements or data attributes the technique tries to display.

This leads to a second major goal of our investigations: understanding how to visualize information across a range of display devices, for example, large, collaborative displays like a multi-projector powerwall, or small, portable displays like a smartphone. These nontraditional display environments are becoming more and more popular. Differences in a display's properties can have a significant effect

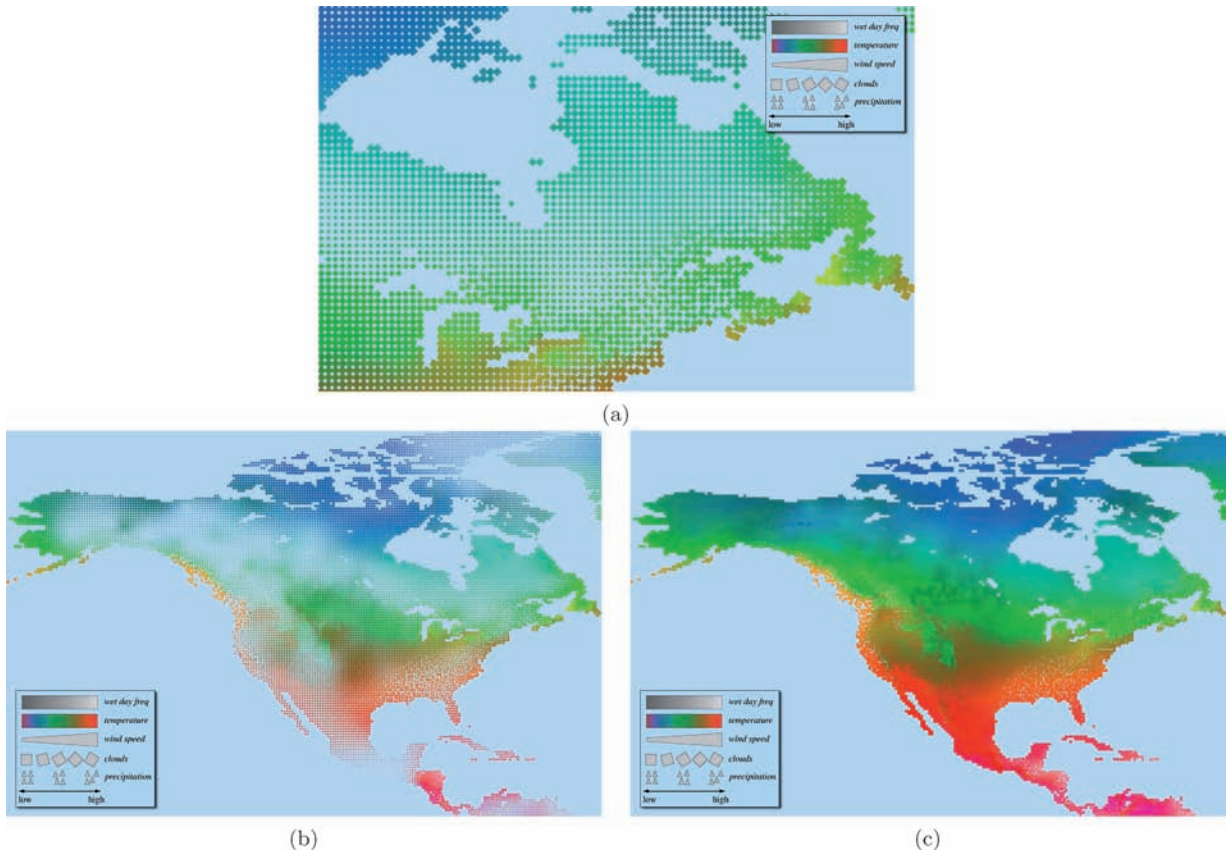

Fig. 1. Weather visualizations, *wet day frequency* → luminance, *temperature* → hue, *precipitation* → regularity, *wind speed* → size, *cloud coverage* → orientation: (a) near Hudson's Bay, all elements distinguishable; (b) North America with low resolution and visual angle, size and orientation interfere with luminance, hue, and regularity; (c) size and orientation removed.

on what it can visualize, however. Results from our studies can be used to determine how a display's physical size, pixel count, and standard viewing distance affect its visualization capabilities. This in turn will define which fraction of a dataset a display can visualize effectively.

Figure 1 shows a simple example of the effects of limited resolution and visual angle. Figure 1(a) visualizes historical weather conditions for January around Hudson Bay and the Great Lakes. Here, sufficient resolution and visual angle are available to distinguish the luminance, hue, size, orientation, and regularity of each element. Bright regions (higher *wet day frequency*) are visible on the west coast of Newfoundland, tilted elements (moderate *cloud coverage*) occur throughout the map, and irregularly placed elements (higher *precipitation*) are seen in Pennsylvania (in the lower-center of the map). Figure 1(b) zooms out to display all of North America. Now, limited resolution and visual angle hide certain visual properties. Differences in orientation are difficult to identify. More importantly, small elements hide other features that might be visible in isolation. Figure 1(c) removes orientation and size, since neither are considered usable. The remaining features—luminance, hue, and regularity—are easier to interpret.

Results from our investigations represent perceptual rules that describe how best to apply fundamental properties of color and texture during visualization. The intent is not to design a single new

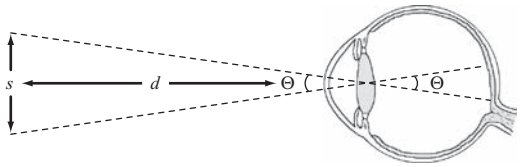

Fig. 2. Subtended visual angle θ, based on an element's physical size s and viewing distance d.

Table I. Visual Acuity Limits for Basic Visual Properties

Property	Description	Acuity
Point acuity	Resolve two point targets	$1'$
Grating acuity	Distinguish bright and dark bars in a uniform gray patch	$1–2'$
Letter acuity	Resolve letters, 20/20 vision means a 5 arc-minute letter can be seen with 90% accuracy	$5'$
Stereo acuity	Resolve a just-noticeable depth difference through binocular disparity	$10''$
Vernier acuity	Resolve if lines are collinear	$10''$

visualization technique, nor to focus on a specific domain area. Instead, our findings are meant to compliment any visualization technique, current or future, that uses these visual features to display information. We can also suggest ways to optimize the amount of information being visualized as the display environment changes, for example, as the data is moved between different display devices, or as a viewer zooms in and out on their data. Perceptual level-of-detail hierarchies, together with a viewer's interests and analysis needs, can be combined to ensure that the most important data is given priority during visualization.

2. BACKGROUND

Issues related to perception, visual acuity, and the display of information have been studied by different communities. We briefly review related research from human vision and visualization.

2.1 Human Vision

Light entering our eyes falls on the retina, a thin layer of photosensitive cells made up of approximately 120 million rods and 6 million cones [Glassner 1995; Ware 2012]. Cones are responsible for vision in bright light, and are divided into three types, each of which is most sensitive to three different wavelengths of light centered in the short, medium, and long regions of the visible spectrum. In bright light, the cones are used to determine both color and luminance information. Rods are used in low light situations, and are sensitive to variations in luminance. Since the cones do not function in low light, we see luminance differences alone. Information leaving the eye is reorganized into three channels: a luminance channel, and two opponent color channels that correspond to red–green and blue–yellow. Processing of visual properties that depend on local spatial arrangements (e.g., size, orientation, and motion) begins in the primary visual cortex (V1) [Marr 1982]. Output from this stage of vision is fed into the higher-level cognitive system.

Visual acuity is a measurement of our ability to see detail. It defines limits on the information densities that can be perceived. Acuity is measured with visual angle, the angle subtended by an object on a viewer's retina specified in arc-degrees, arc-minutes $'$ ($1^\circ = 60'$), or arc-seconds $''$ ($1' = 60''$). Given an object size s and a viewing distance d (Figure 2), visual angle θ can be calculated as $\theta = 2 \arctan(^{s/2}/_{d})$.

Some basic acuities are summarized in Table I [Ware 2012]. Most acuities fall in the $1–2'$ range, which corresponds roughly to the spacing of receptors in the center of the fovea. A number of

superacuities exist, however, including stereo and vernier acuity. Here, post-receptor mechanisms integrate input from multiple receptors to resolve at higher resolutions.

Neural postprocessing combines input from both eyes. The area of overlap is approximately 120° with 30–35° of monocular vision on either side. This allows for combined horizontal and vertical fields-of-view of approximately 180–190° and 120–135°, respectively. Acuity within the overlap region has been shown to be up to 7% more accurate, compared to the monocular boundaries [Campbell and Green 1965].

2.2 Perception in Visualization

A long-standing goal in visualization is the construction of guidelines from human visual perception, to be used, for example, to choose visual features that are best suited to representing different types of data, or to supporting different analysis tasks.

A well-known example of perception in visualization is the use of perceptually balanced colormaps to visualize continuous scalar values. Color models like CIE LUV, CIE Lab, and Munsell were built to provide a rough measure of perceptual balance, where a unit step anywhere along the color scale produces a perceptually uniform difference in color [CIE 1978]. More recent models like s-CIE Lab perform an initial filtering step to simulate the blurring that occurs when viewing high spatial frequency color patterns [Zhang and Wandell 1997]. Rheingans and Tebbs refined this basic idea by plotting a path through a perceptually balanced color model, then asking viewers to define how attribute values map to positions on the path [Rheingans and Tebbs 1990]. Rogowitz and Treinish proposed rules to automatically select a colormap based on an attribute's spatial frequency, its continuous or discrete nature, and the analysis tasks viewers need to perform [Rogowitz and Treinish 1993]. Ware ran experiments that asked viewers to distinguish between individual colors and shapes formed by colors. He used the results to build a colormap that spirals up the luminance axis, providing perceptual balance and controlling simultaneous contrast error [Ware 1988]. Healey conducted a visual search experiment to determine the number of colors a viewer can distinguish simultaneously. His results showed that viewers can rapidly choose between up to seven isoluminant colors [Healey 1996].

Similar studies have explored the perceptual properties of texture. These were derived by measuring statistical properties to perform texture segregation (e.g., Julész et al. [1973]), through experiments that asked viewers to discriminate texture images (e.g., Rao and Lohse [1993]), or by defining filters that mimic the low-level neural mechanisms of human vision (e.g., radially symmetric and directionally tuned difference of offset Gaussians used in Malik and Perona [1990]). Results identified size, directionality, contrast, and regularity, among others. These individual properties have been studied extensively. For example, Watson and Ahumada constructed a standardized contrast image set, then optimized a model that included effects of contrast filtering, orientation, and distance from fixation [Watson and Ahumada 2005]. Liu et al. conducted a study of the effects of aspect ratio and "roundness" of an element on orientation detection [Liu et al. 2002]. Even at aspect ratios as small as 1.5:1, viewers could align a probe to a target element with an average orientation error of $\pm 3^\circ$.

Interestingly, although various texture features are strong discriminators in image segmentation algorithms, some of them are difficult to perceive. For example, experiments by Healey and Enns showed that size, directionality, and contrast are perceptually salient, but regularity is not [Healey and Enns 1999]. Healey and Enns used their results to construct "pexels," perceptual texture elements that vary their luminance, hue, size, orientation, and spatial density to represent multiple data attributes. Other researchers have used a similar strategy to visualize data with texture. Grinstein et al. designed stick men whose joint angles encode attribute values. When the stick men are arrayed across a display, they form texture patterns whose spatial groupings and boundaries highlight attribute correlations [Grinstein et al. 1989]. Ware and Knight used Gabor filters that varied their orientation, size, and

contrast to visualize data elements with three separate attribute values [Ware and Knight 1995]. An intriguing approach proposed by Interrante uses "natural" textures like weaves or brick patterns to visualize data [Interrante 2000].

2.3 Level-of-Detail Visualization

The need to visualize large datasets is common. Many techniques attempt to make better use of available display resources. We briefly discuss approaches that address this problem by allowing a viewer to selectively increase the resolution and visual angle of some of the elements in a visualization.

Focus + context techniques construct visualizations that combine a high-level overview and lower-level details in a single image. For example, magnifying lenses increase resolution and visual angle in user-selected areas of detail (e.g., semantic fisheye lenses [Furnas 1986]). A related technique maps from Cartesian coordinates to hyperbolic space to magnify elements at the center of the display, pushing more distant elements towards the edge of the screen [Lamping and Rao 1996].

A separate study by Yost et al. investigated visualizing data on large displays [Yost et al. 2007]. Pixel densities were chosen to exceed a viewer's point acuity at the initial viewing position. Yost did not try to identify limits on acuity, however. Instead, viewers were allowed to physically navigate—to walk around the display—to resolve different parts of the visualization. Results showed that as the display size increased, the amount of additional time needed to complete tasks rose more slowly than the amount of additional data being presented. This suggests that large, high-resolution displays can make visual exploration more efficient.

Our study of perceptual limits is not meant to replace these approaches. Instead, our results could be integrated into these techniques, for example, by defining the level of magnification needed to properly perceive the visual properties of objects within a fisheye lens or hyperbolic region. Our investigations are also focused on a slightly different question. Rather than asking "What *can* our display environment show?" we want to answer "What can our viewers *perceive* within the display environment?" This will allow us to filter our visualizations to make best use of the display resources we have available.

3. EXPERIMENT DESIGN

Our goal is to determine, for a given approach, the type and amount of visual information a human viewer can comprehend. We believe this depends critically on two properties: resolution and visual angle.

Resolution involves individual pixels on a display device. A minimum number of pixels are needed to distinguish differences in a visual feature. For example, a single pixel cannot show differences in size or orientation, since by definition an individual pixel cannot be resized or rotated. A single pixel might be sufficient to present distinguishable differences in luminance or hue, however. We will identify the minimum number of pixels needed for a collection of common visual features.

Visual angle defines the minimum physical size an object must subtend on the viewer's retina to produce distinguishable differences in a visual feature. Consider objects displayed on a multi-projector powerwall. If a viewer stands back to observe the entire wall, small elements may be too small to perceive. The subtended visual angle of an object is calculated from the object's on-screen size in pixels, the physical width and height of the display device, and the distance from the display to the viewer. Since objects can be rotated, we define visual angle to be the subtended angle of the bounding box of an object in all its possible orientations. We want to identify the minimum visual angle needed for the same collection of common visual features.

We decided to investigate four visual features during our initial experiments: luminance, hue, size, and orientation. These features include aspects of color and texture, and each feature is common to many visualization techniques and systems. We believe understanding the resolution and visual

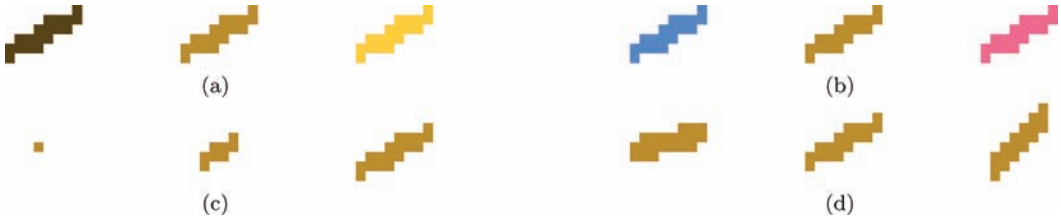

Fig. 3. Visual features used during the resolution experiment; all elements shown at 8×2 resolution except for size feature: (a) dark, medium, and bright luminance, CIE $xyY = (0.44, 0.48, 6.7)$, $(0.43, 0.48, 13.6)$, $(0.38, 0.43, 19.5)$, respectively; (b) blue, brown, and pink hue, CIE $xyY = (0.22, 0.23, 13.7)$, $(0.43, 0.48, 13.6)$, $(0.34, 0.28, 13.6)$, respectively; (c) 1, 4, and 16 pixel size; (d) 15°, 35°, and 55° orientations counterclockwise from horizontal.

angle limits for these features will help to guide future studies of related color, texture, and motion properties.

3.1 Selecting Visual Features

We began our studies by selecting a small set of values for each feature to be tested. We had to ensure these values were distinguishable when applied to an element with a large resolution and visual angle. This is necessary to draw meaningful conclusions from the experiments that use these values. Suppose, for example, we were searching for the hue resolution limit. Because our hues are known to be distinguishable for large objects, if a viewer cannot see a color difference, we can assume this is due to insufficient resolution.

To identify distinguishable feature values, we ran a preliminary pairwise-comparison experiment. First, a saturated yellow "anchor" value f_1 was selected with a luminance roughly 65% of the displayable maximum (Figure 3(a), center element). f_1 had a resolution of 32×8 pixels subtending a visual angle of 0.97° at a 22-inch viewing distance. We then used a modified Cornsweet staircase method to obtain distinguishability with 75% accuracy between f_1 and a second feature f_2 [Pollack 1968; Wales and Blake 1970].

—A value of f_2 was chosen with f_2 slightly different than f_1, that is, $\Delta F = |f_1 - f_2| < \epsilon$, where ϵ is a distance that is too small to produce distinguishable differences.

—Two large rectangular elements were displayed for 200 msec on a white background, then replaced with a grey mask. One element was randomly selected to show f_1. The other element showed f_2.

—Viewers were asked whether the elements were identical or different.

—ΔF was changed by moving f_2 relative to f_1 until viewers could reliably identify the two rectangles as being different.

—An identical procedure was applied, but in the opposite direction of f_1–f_2 to choose another value f_3.

This produced three values f_1, f_2, and f_3 that are distinguishable from one another when sufficient resolution and visual angle are provided.

Applying this approach to each of the four visual features we tested produced eight pairings—two for each feature—with equal perceived distinguishability, both within each feature and across different features. We used the Munsell color space to choose luminances and hues, since it provides perceptual balance and subdivides hue into discrete, named regions. Munsell values specified in CIE LUV [Wyszecki and Stiles 1982] were converted to CIE xyY, then to monitor RGB values calibrated to generate the expected colors. Figure 3 shows approximations of these values.

3.2 Novel Target Detection

Trials in our experiments contain a random collection of distractor elements, all with identical luminance, hue, size, and orientation. Half the trials are randomly selected to contain a target region, a group of target elements with a visual appearance different from the distractors. For example, hue targets would display a different hue than the distractors.

In traditional visual search experiments, a viewer is told what type of target to look for—search for blue targets, for example—with response time and accuracy used to measure performance. In a visualization environment, however, the visual properties of a region of interest normally cannot be anticipated, particularly during exploratory analysis. To simulate this incomplete knowledge, we do not prime our viewers by identifying the target. Instead, viewers are told "You will be asked to report whether you saw a small group of elements that differ from all the other elements. If you see this group, you will also be asked whether it had a horizontal or vertical orientation."

By varying the specific feature that differs between the target and the distractors, we can measure how rapidly viewers *detect* arbitrary targets within some criteria for success. It is also important to differentiate between a viewer's ability to detect a target—I see the target—versus the ability to identify properties of the target—I see this kind of target. Asking about the target group's orientation allows us to determine whether viewers can also determine the shape of the target group and its basic orientation. We hope that results from this type of limited knowledge search experiment will be more directly applicable to real-world visualization domains.

4. RESOLUTION EXPERIMENTS

The resolution experiments tested the effect of resolution on viewer response time and accuracy. Three different resolutions were tested: a single pixel (a 1×1 region), four pixels (a 4×1 region), and sixteen pixels (an 8×2 region).

4.1 Design

Trials from the resolution experiment were shown on a 19-inch LCD monitor (14.8×11.9 inches subtending $37.38 \times 30.28^{\circ}$) with dimensions of 1280×1024 pixels, or 86 pixels-per-inch. Data elements in each trial were displayed as 2D rectangular glyphs with an aspect ratio of 4:1. Since we are only interested in how resolution affects viewer performance, we chose an element size that guaranteed a sufficient visual angle. Given our fixed viewing distance of 22-inches from the monitor, we needed to use "virtual pixels" formed by a 2×2 block of physical pixels. Any further reference to "pixels" in this section implies virtual pixels.

Each trial contained a 20×20 array of elements, placed within a grid of 48×48 pixel cells subtending $28.47 \times 28.47^{\circ}$. Each element was randomly jittered to introduce irregularity in the layout. The background in each trial was white, with CIE $xyY = (0.29, 0.30, 24.9)$. A viewer was presented with 384 trials in random order. Half the trials were chosen to contain a target patch—target present trials—while the other half—target absent trials—did not. For target present trials, the target patch was randomly located and was positioned at least one row and one column away from the edge of the grid.

Viewers were told to determine whether a region of elements different from the distractors was present or absent in each trial. Viewers pressed one of two keys to record their answer. Viewers were told to respond as quickly as possible, while still maintaining a high level of accuracy. After an answer was entered, the display was replaced with a grey mask. If the viewer correctly reported the target patch, he was then asked whether it had a horizontal or vertical orientation. Response time and accuracy were recorded on each trial for later analysis.

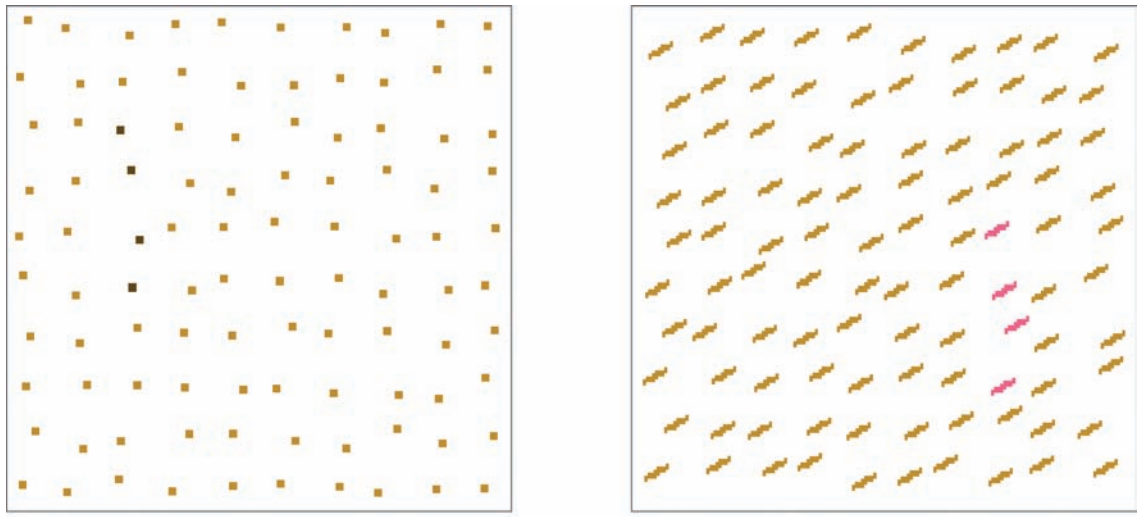

Fig. 4. Example resolution trials, actual trials were 20×20 arrays of elements: (a) luminance trial with 1×1 resolution, dark vertical target present; (b) hue trial with 8×2 resolution, pink vertical target present.

We constructed eight different feature sets for the three different feature values f_1, f_2, and f_3 as follows.

(1) Distractors use feature f_1, targets use feature f_2.
(2) Distractors use feature f_2, targets use feature f_1.
(3) Distractors use feature f_1, targets use feature f_3.
(4) Distractors use feature f_3, targets use feature f_1.
(5) Distractors use feature f_1.
(6) Distractors use feature f_2.
(7) Distractors use feature f_1.
(8) Distractors use feature f_3.

The first four feature sets represent target present trials; the last four represent target absent trials.

Resolution varied during the experiment. Each feature set used three different resolutions: elements formed by 1×1 pixels (Figure 4(a)), 4×1 pixels, and 8×2 pixels (Figure 4(b)). This produced eight feature sets by three resolutions by two target patch orientations for a total of 48 different trial types. Two trials for each trial type were shown, producing 96 trials. Given the four visual features we tested—luminance, hue, size, and orientation—the total number of trials in the experiment was 384.

For certain resolutions, target values of size and orientation are not possible. For example, for size trials with a 1×1 resolution, a viewer cannot see larger or smaller targets because both target and distractor are constrained to be one pixel in size. These trials are included to properly balance the experiment conditions during statistical analysis. The question arises, then, should answering "target absent" for these trials be considered correct or incorrect? We decided to mark these trials as incorrect. Although it is impossible for the viewer to distinguish the target, this situation mirrors a real visualization environment where elements with a 1×1 resolution are shown, and size is used to represent attribute values. Since "error" in our experiment means "the visual feature is not distinguishable at

Table II. Raw $\overline{rt}$ and $\overline{a}$ by Resolution for the Four Feature Types: (a) by Target Present; (b) by Target Absent

Feature	1×1	4×1	8×2
luminance	2.81 / 51%	1.88 / 85%	1.47 / 95%
hue	3.29 / 26%	2.38 / 73%	1.30 / 100%
size	3.43 / 2%	2.69 / 50%	1.21 / 99%
orientation	3.59 / 0%	3.24 / 79%	1.52 / 99%

(a)

Feature	1×1	4×1	8×2
luminance	3.82 / 99%	3.44 / 98%	2.98 / 99%
hue	3.63 / 99%	3.60 / 99%	2.85 / 99%
size	3.40 / 97%	3.68 / 99%	3.40 / 99%
orientation	3.47 / 100%	3.83 / 97%	3.07 / 99%

(b)

the given resolution," this choice seems reasonable. For elements with a 1×1 resolution, 16 size trials contain a target that cannot be seen. Similarly, for elements with a 4×1 resolution, there are 8 trials where a larger 8×2 target cannot be shown, and therefore cannot be seen. The same situation occurs for orientation trials at a 1×1 resolution: 16 target present trials contain targets that are the same orientation as the distractors. This effect can be seen in Table II, where the accuracies for size targets at 1×1 and 4×1 resolutions are 2% and 50%, respectively, and the accuracy for orientation targets at a 1×1 resolution is 0%.

Eleven graduate students (eight males and three females) participated during the experiment. Although some of the viewers were familiar with traditional visual search experiments, none had previous experience with the novel detection task we employed. All viewers had normal or corrected-to-normal vision, and none of the viewers were color blind. Each viewer completed trials for all four feature types. Trials for the different features were intermixed, to support our novel target detection approach. Since we are running a within-subjects experiment, our main concern with respect to sample size is the total number of trials completed for each experiment condition. Our eleven participants produced 22 repetitions of each trial type, for a total of 4,224 trials.

The experiment was preceded by a small set of practice trials, to allow viewers to gain familiarity with the different features being tested, and with the format and speed of the experiment. Viewers were shown 16 practice trials—four features by target present and target absent by two target patch orientations. Viewers continued the practice trial session until they reported they were comfortable with the procedure and achieved an acceptable accuracy level. Following the practice session, the main experiment was conducted. The total running time of the practice session and the main experiment was approximately 60 minutes.

4.2 Results

Mean response time $\overline{rt}$ and mean viewer accuracy $\overline{a}$ were calculated for each unique combination of target feature type, target present or absent, resolution, and target patch orientation. Multifactor analysis of variance (ANOVA) was used to identify statistically significant differences in performance. When significant differences were found, post-hoc Tukey honestly significant difference (HSD) tests were used to identify pairwise significance. In summary, we identified the following statistically significant results:

(1) $\overline{rt}$ decreased monotonically as resolution increased;
(2) $\overline{a}$ increased monotonically for target present trials as resolution increased;
(3) $\overline{rt}$ varied by feature type: it was lowest for luminance trials and highest for orientation trials;
(4) $\overline{a}$ varied by feature type: it was highest for luminance trials and lowest for size trials; and
(5) $\overline{rt}$ and $\overline{a}$ were higher for target absent trials, compared to target present trials.

4.3 Detailed Statistics

We started by testing to see if different types of target present trials (e.g., f_1 distractors and f_2 targets versus f_1 distractors and f_3 targets) produced significantly different $\overline{rt}$ or $\overline{a}$. If so, it would suggest that

Table III. Combined $\overline{rt}$ and $\bar{a}$ for the Four Feature Types: (a) by Target Present or Absent; (b) by Resolution

Feature	P	A
luminance	2.05 / 77%	3.41 / 99%
hue	2.32 / 66%	3.36 / 99%
size	2.44 / 50%	3.49 / 98%
orientation	2.78 / 59%	3.46 / 98%

(a)

Feature	1×1	4×1	8×2
luminance	3.31 / 75%	2.66 / 92%	2.22 / 97%
hue	3.46 / 63%	2.99 / 86%	2.08 / 100%
size	3.42 / 49%	3.18 / 75%	2.31 / 99%
orientation	3.53 / 50%	3.53 / 88%	2.29 / 99%

(b)

some feature pairings were easier to distinguish than others. No significant differences were found. Based on this, we combined and averaged results over the four target present trial types.

ANOVAs showed a significant effect of resolution on both $\overline{rt}$ and $\bar{a}$, $F(2, 20) = 13.4$, $p < 0.005$ and $F(2, 20) = 508.79$, $p < 0.005$. As resolution increased, viewers were faster and more accurate. Similarly, ANOVAs identified significant differences by feature type for both $\overline{rt}$ and $\bar{a}$, $F(3, 30) = 4.18$, $p = 0.014$ and $F(3, 30) = 62.55$, $p < 0.005$. Tukey HSD comparisons showed a significant difference between luminance and orientation for $\overline{rt}$ ($p = 0.034$) and between luminance and size for $\bar{a}$ ($p = 0.005$).

For target present trials, $\overline{rt}$ was 2.4 seconds (s) and $\bar{a}$ was 64%. For target absent trials, they were 3.43 s and 99%. Target presence or absence was significant for both measures, $F(1, 10) = 36.21$, $p < 0.005$ for $\overline{rt}$ and $F(1, 10) = 470.09$, $p < 0.005$ for $\bar{a}$. This type of finding is not uncommon for target present–absent experiments. When viewers do not initially see a target, they spend some additional time confirming their result. Also, when $\bar{a}$ is low, $\overline{rt}$ differences cannot always be trusted. Since we are focusing on situations with higher $\bar{a}$, this does not affect our guidelines on feature use.

Significant interaction effects for both $\overline{rt}$ and $\bar{a}$ were also identified: a feature type $\times$ target present or absent effect, and a feature type $\times$ resolution effect. Table III(a) shows $\overline{rt}$ and $\bar{a}$ for each feature type, divided by target present or absent. Differences in $\overline{rt}$ and $\bar{a}$ occurred across feature type for target present trials, but were nearly uniform for target absent trials, with $F(3, 30) = 5.11$, $p = 0.006$ for $\overline{rt}$ and $F(3, 30) = 46.06$, $p < 0.005$ for $\bar{a}$. Table III(b) shows $\overline{rt}$ and $\bar{a}$ for each feature type, divided by resolution. Differences in $\overline{rt}$ and $\bar{a}$ were larger for smaller resolutions, $F(6, 60) = 3.24$, $p = 0.008$ and $F(6, 60) = 26.31$, $p < 0.005$, respectively. An accuracy ceiling is being approached for resolutions of 8×2. This is not surprising, since the values for each feature were chosen to produce roughly uniform distinguishability, that is, to be equally detectable.

Finally, there was no significant difference between target detection and identifying a target group's orientation as horizontal or vertical ($t(20) = 0.584$, $p = 0.57$), with $\bar{a} = 98.8\%$ and $\bar{a} = 97.2\%$, respectively, for trials where viewers correctly detected a target patch.

5. VISUAL ANGLE EXPERIMENTS

The visual angle experiments tested the effect of visual angle on viewer response time and accuracy. Three different visual angles were tested: 0.06125°, 0.1225°, and 0.245°.

5.1 Design

As before, trials were presented on a 19-inch LCD monitor with 1280×1024 pixels. Elements in each trial were displayed as 2D rectangular glyphs with an aspect ratio of 4:1. Since we are only interested in how the visual angle affects viewer performance, we used the monitor's native resolution and sized each element to render without a pixelated or jagged appearance at each visual angle. Our resulting elements were too large to produce the required visual acuities at the original viewing distance, so viewers were moved back to a distance of approximately 86.5-inches, producing a 0.06125° visual angle. Elements displayed at this visual angle had a resolution of 12×3 pixels. Since this is above the required resolution of 8×2 pixels identified during the previous experiment, an inability to identify targets would be due to an insufficient visual angle. The element's resolution was doubled and

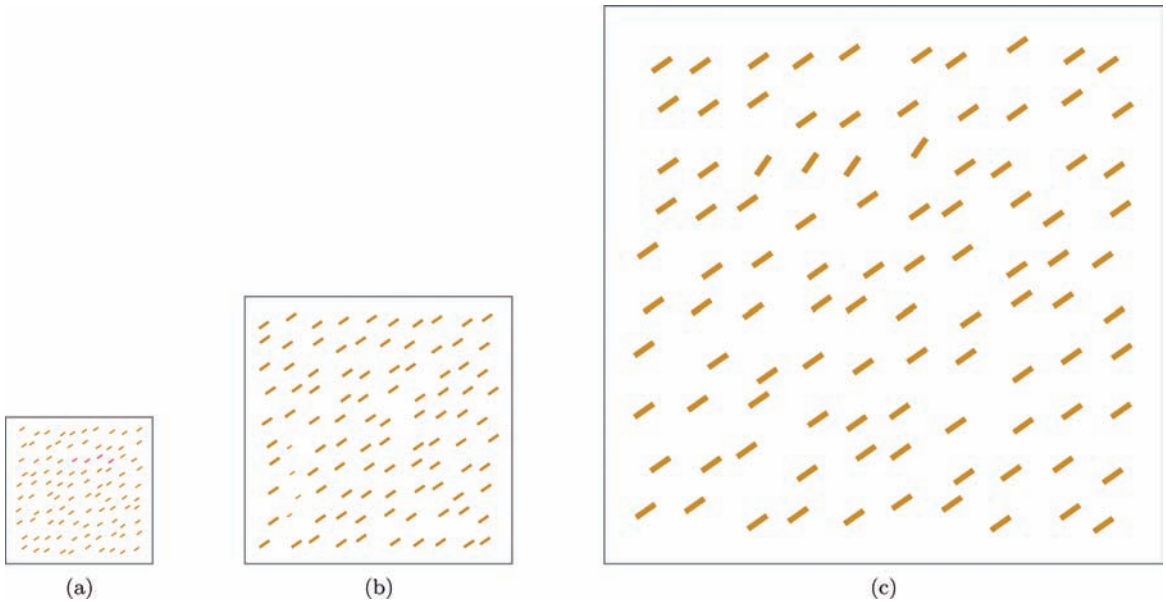

Fig. 5. Example visual angle trials, actual trials were 20 × 20 arrays of elements: (a) hue trial with a 0.06125° visual angle, pink horizontal target present; (b) size trial with a 0.1225° visual angle, smaller vertical target present; (c) luminance trial with a 0.245° visual angle, brighter horizontal target present.

quadrupled—24 × 6 pixels and 48 × 12 pixels—to produce visual angles of 0.1225° and 0.245°. Figure 5 shows examples of hue, size, and luminance trials using these visual angles.

Apart from these differences, the conditions in the visual angle experiment were identical to the resolution experiment: four visual features—luminance, hue, size, and orientation—two target conditions—present and absent—four target–distractor pairings, and two target patch orientations—horizontal and vertical. The number of repetitions of each trial type and the practice trials presented to a viewer prior to running the experiment were also the same.

Eleven graduate students—eight males and three females—participated during the experiment. Seven of the viewers also participated in the resolution experiment. For these viewers, the order of the experiments was randomized. Four viewers completed the resolution experiment first. The remaining three completed the visual angle experiments first. All viewers had normal or corrected-to-normal vision, and none of the viewers were color blind.

5.2 Results

Mean response time $\overline{rt}$ and mean viewer accuracy $\overline{a}$ were calculated for each trial type. Multifactor ANOVA and Tukey HSD were used to identify statistically significant differences in performance. In summary, we identified the following statistically significant results:

(1) $\overline{rt}$ decreased monotonically as visual angle increased;

(2) $\overline{a}$ increased monotonically as visual angle increased;

(3) $\overline{rt}$ varied by feature type: it was lowest for size trials and highest for orientation trials;

(4) $\overline{a}$ varied by feature type: it was highest for size trials and lowest for hue trials; and

(5) $\overline{rt}$ and $\overline{a}$ were higher for target absent trials, compared to target present trials.

Table IV. Raw $\overline{rt}$ and $\overline{a}$ by Visual Angle for the Four Feature Types: (a) by Target Present; (b) by Target Absent

Feature	0.06125°	0.1225°	0.245°
luminance	2.68 / 34%	2.09 / 73%	1.45 / 93%
hue	3.25 / 11%	2.60 / 52%	1.51 / 87%
size	2.56 / 40%	1.46 / 85%	1.24 / 99%
orientation	3.84 / 3%	2.71 / 66%	1.66 / 95%

(a)

Feature	0.06125°	0.1225°	0.245°
luminance	3.03 / 100%	3.30 / 99%	2.47 / 99%
hue	3.02 / 98%	3.03 / 98%	2.69 / 98%
size	2.83 / 99%	2.80 / 99%	3.00 / 100%
orientation	2.91 / 99%	3.06 / 99%	2.55 / 99%

(b)

Table V. Combined $\overline{rt}$ and $\overline{a}$: (a) Visual Angle by Target Present or Absent; (b) Feature Type by Target Present or Absent

Angle	P	A
0.06125°	3.08 / 22%	2.95 / 99%
0.1225°	2.21 / 69%	3.05 / 99%
0.245°	1.47 / 94%	2.68 / 99%

(a)

Feature	P	A
luminance	2.08 / 66%	2.93 / 99%
hue	2.46 / 50%	2.91 / 98%
size	1.75 / 75%	2.87 / 99%
orientation	2.74 / 55%	2.84 / 99%

(b)

5.3 Detailed Statistics

As before, we confirmed that different target pairings for target present trials did not produce significantly different $\overline{rt}$ or $\overline{a}$, allowing us to combine and average results over the four target present trial types.

ANOVAs showed a significant effect of visual angle on both $\overline{rt}$ and $\overline{a}$, with $F(2, 20) = 7.78$, $p < 0.005$ and $F(2, 20) = 359.29$, $p < 0.005$, respectively (Table IV). As the visual angle increased, viewers were faster and more accurate. Different feature types also produced significantly different $\overline{rt}$ and $\overline{a}$, with $F(3, 30) = 7.72$, $p < 0.005$ and $F(3, 30) = 44.08$, $p < 0.005$, respectively. Tukey HSD comparisons found a significant size–orientation difference in $\overline{rt}$ ($p < 0.005$), and significant size–orientation and size–hue differences in $\overline{a}$ ($p = 0.039$ and $p < 0.005$, respectively). Identification of targets based on size was significantly faster and more accurate than some of the other features.

Target present trials were faster but less accurate than target absent trials (Table IV), with $\overline{rt}$ of 2.25 s versus 2.89 s for target present or absent ($F(1, 10) = 14.9$, $p < 0.005$) and $\overline{a}$ of 61% versus 99% for target present or absent ($F(1, 10) = 95.28$, $p < 0.005$). As with the resolution experiment, viewers tended to confirm their intuition that no target is present, producing slower but more accurate results.

Significant interaction effects were identified in both $\overline{rt}$ and $\overline{a}$ for visual angle × target present or absent, and feature type × target present or absent.

Table V(a) shows $\overline{rt}$ and $\overline{a}$ for visual angles 0.06125°, 0.1225°, and 0.245°, divided by target present or absent. Differences in $\overline{rt}$ and $\overline{a}$ changed more rapidly for target present versus target absent trials—$F(2, 20) = 5.67$, $p = 0.011$ and $F(2, 20) = 146.75$, $p < 0.005$, respectively. Table V(b) shows $\overline{rt}$ and $\overline{a}$ for each feature type, divided by target present or absent. Differences in $\overline{rt}$ and $\overline{a}$ were large for target present trials versus nearly uniform for target absent trials, with $F(3, 30) = 5.4$, $p < 0.005$ and $F(3, 30) = 17.25$, $p < 0.005$, respectively.

Finally, there was no significant difference between target detection and identifying a target group's orientation as horizontal or vertical ($t(20) = 0.335$, $p = 0.74$), with $\overline{a} = 98.9\%$ and $\overline{a} = 97.3\%$, respectively, for trials where viewers correctly detected a target patch.

5.4 Different Feature Values and Tasks

Our experiments studied three values for each feature type: an anchor f_1, and two values f_2 and f_3 that are equally distinguishable from f_1. We might ask "What if you used a different set of colors, or a different set of sizes? Would the resolution and visual angle limits be the same for these new feature values?"

We believe the answer is Yes; that is, resolution and visual angle limits depend on feature values being distinguishable from one another, and not on the specific values being displayed. Recall that we chose feature values using just-noticeable differences (JNDs). Picking, for example, a different anchor hue f_1 and new hue values f_2 and f_3 will produce the same amount of JND between f_1–f_2 and f_1–f_3. If resolution and visual angle limits are dependent on distinguishability, then using these new feature values in our experiments should produce identical resolution and visual angle limits for hue. We have applied our results to visualization domains that do use a wide range of feature values. Performance has closely tracked our theoretical findings. To fully test our hypothesis, however, we would need to choose multiple sets of feature values, then conduct experiments to measure and compare the resolution and visual angle limits of each set.

A related finding in our experiments was that, at the smallest resolution and visual angle, accuracy for luminance trials varied significantly based on the specific luminance value used for the target. This did not occur for the other three features. Since all of the targets had poor accuracy at this resolution and visual angle, it did not affect our guidelines on the use of luminance. If follow-on experiments with new feature values were conducted, however, it might provide an opportunity to better study this anomaly.

Another question to consider is "Are these results limited to the specific task viewers performed during the experiments?" We asked viewers to search for target regions, and report on their orientations. This was done in part because it approximates searching for differences in a visualization, and in part because visual search is a common experimental method for measuring feature salience. As with different feature values, we are optimistic that these results will extend to additional tasks like spatial boundary detection, object tracking, and estimation. Although we have anecdotal evidence from applying our results to visualizations where different tasks are performed, additional experiments are needed to properly test this capability.

6. DISPLAY HIERARCHIES

Based on our experimental results, we are now able to define minimum requirements in resolution and visual angle to produce a required level of accuracy. These limits are encoded into a distinguishability graph that subdivides 2D space into regions of "distinguishable" and "not distinguishable" for each feature. A display environment's properties—the display's physical size and pixel dimensions, the viewing distance, and the number of pixels allotted to each element—define a "display point" in the distinguishability graph. For each visual feature, this display point lies within either the distinguishable or not distinguishable region. This defines whether the feature should be used to visualize data.

As the display environment changes—as a viewer zooms to change the resolution and visual angle of the elements, or when a viewer moves to a different display device—the display point moves. As the display point enters or leaves different distinguishable regions, the corresponding visual features can be enabled or disabled. In this way the visualization dynamically updates to ensure that only those features that produce distinguishable differences are displayed.

6.1 Distinguishability Graph

Although any threshold accuracy can be used to define a feature's minimum required resolution and visual angle, we selected $\bar{a} \geq 66\%$ on target present trials as our cutoff (Figure 6(b)).

One interesting finding, at least in our experiments, was that neither luminance nor hue allowed for accurate detection for elements with a single-pixel resolution, even when a very small number of highly distinguishable values were shown. Size required the largest minimum resolution—8×2, but this is not surprising, since elements with 1×1 and 4×1 resolutions were unable to show some or all of the smallest and largest size values we studied. Hue required the largest minimum visual

Feature	Resolution	Angle
luminance	4×1	0.1225°
hue	4×1	0.245°
size	8×2	0.1225°
orientation	4×1	0.1225°

(a)

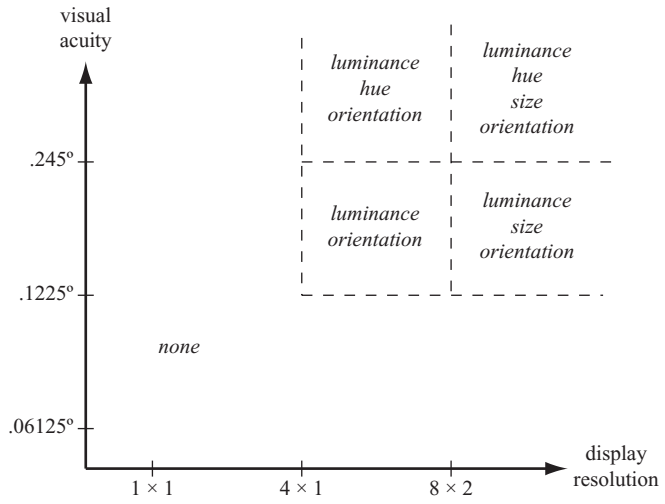

(b)

Fig. 6. Distinguishability limits: (a) minimum resolution and visual angle for the visual features luminance, hue, size, and orientation; (b) distinguishability graph showing the visual features that are recognizable within each region of the graph.

angle—0.245°. Identification was good when a sufficient visual angle was provided—87% for elements with a 4×1 resolution, but dropped off once the visual angle fell to 0.1225°—52%, even though distinguishable hues and large resolutions were being displayed.

6.2 Psychometric Curve Fitting

An alternative suggestion to determine the resolution and visual angle limits needed to identify target patches is to fit a psychometric curve to percentage correct results to locate the 66% accuracy rate [Ulrich and Miller 2004]. Following this suggestion, we fit a Weibull curve to our data. The Weibull cumulative distribution function $y = 1 - e^{-(x/\lambda)^k}$ can be rewritten to apply to a two-alternative forced choice experiment:

$$y = 1 - (1 - g)e^{-(bx/t)^k}$$
$$b = -\log\left(\frac{1-a}{1-g}\right)^{1/k} \quad (1)$$

Here, $g = 0.5$ is chance performance, t is the resolution or visual angle cutoff for our desired target accuracy $a = 0.66$, and k is the slope of the Weibull function. Our goal is to choose t and k to maximize the function's fit to the likelihood of our viewer's actual responses x.

Figure 7(a) shows an example of fitting a Weibull curve to accuracy data for all size trials—both target present and target absent—across the three display resolutions we tested. The $\bar{a} = 66\%$ cutoff is shown just to the left of 4×1 targets. Figures 7(b) and 7(c) show results for target present trials and all trials, respectively, rounded to the next highest resolution or visual angle we tested. Resolution limits for target present trials are identical to our distinguishability graph (Figure 6(a)), but visual angle limits are higher for luminance and size—0.245° for curve fitting versus 0.1225° for direct comparison to $\bar{a}$. When all trials are analyzed, better performance on target absent trials reduces the resolution limits for luminance and size, and the visual angle limits for all four features. Since our visualization require the ability to identify targets with a given visual feature, we are most interested in target

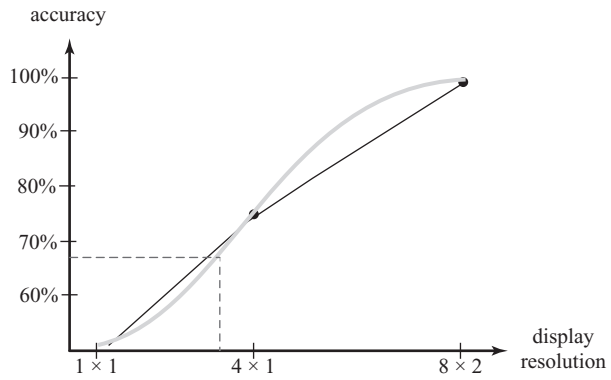

(a)

Feature	Resolution	Angle
luminance	4 × 1	0.245°
hue	4 × 1	0.245°
size	8 × 2	0.1225°
orientation	4 × 1	0.245°

(b)

Feature	Resolution	Angle
luminance	1 × 1	0.06125°
hue	4 × 1	0.1225°
size	4 × 1	0.06125°
orientation	4 × 1	0.1225°

(c)

Fig. 7. Distinguishability limits from psychometric curve fitting: (a) example Weibull curve fit to accuracy data for size targets across all trials and three display resolutions; (b) minimum resolution and visual angle for target present trials; (c) minimum resolution and visual angle for all trials

present performance (Figures 6(a) and 7(b)). We chose to apply the distinguishability graph limits in the examples below, since they hint at the possibility of reasonable accuracy for slightly lower visual angle limits.

6.3 Visualizing Climatology Data

To investigate the practical aspects of our experimental results, we visualized climatology datasets provided by the International Panel on Climate Control. The datasets contain eleven historical monthly weather conditions sampled at $1/2^\circ$ latitude and longitude steps for positive elevations throughout the world.

Figure 1 shows an example of how visual features can be added (Figure 1(a)) or removed (Figure 1(c)) as a viewer zooms in or out on their data. Rather than continuing to investigate zooming in a single display environment, we explore visualizing data on a small screen smartphone, and on a large screen powerwall.

6.3.1 *Smartphone.* Figure 8 shows simulated examples of February's historical weather conditions over North America visualized on an iPhone with a 4-inch × 3-inch screen containing 480 × 320 pixels.

Figure 8(a) visualizes four data attributes: *wet day frequency* → luminance, *temperature* → hue, *wind speed* → size, and *precipitation* → orientation. We built a continuous color scale around the luminance axis in Munsell space near the boundary of our monitor's gamut, subdivided into perceptually uniform luminance and hue steps. This produces colors that (1) are perceptually balanced, and (2) maintain a roughly constant simultaneous contrast error [Healey and Enns 1999]. Elements are assigned a 2×2 pixel region with a visual angle of 0.084° for a 5.8-inch viewing distance. Not surprisingly, variations in size are difficult to distinguish, since both the resolution and the visual angle are below the minimums required for distinguishable sizes—8×2 and 0.1225°, respectively (Figure 6(a)). Even worse, elements with small sizes obscure their luminance, hue, and orientation.

In Figure 8(b) size is removed. Although the resulting display is easier to interpret, the elements still do not have sufficient resolution or visual angle to be fully distinguishable. In Figure 8(c) the visualization has been zoomed 200%, increasing the elements to a resolution of 4×4 pixels and a visual angle of 0.1685°. This produces some noticeable improvements, for example, in orientation's

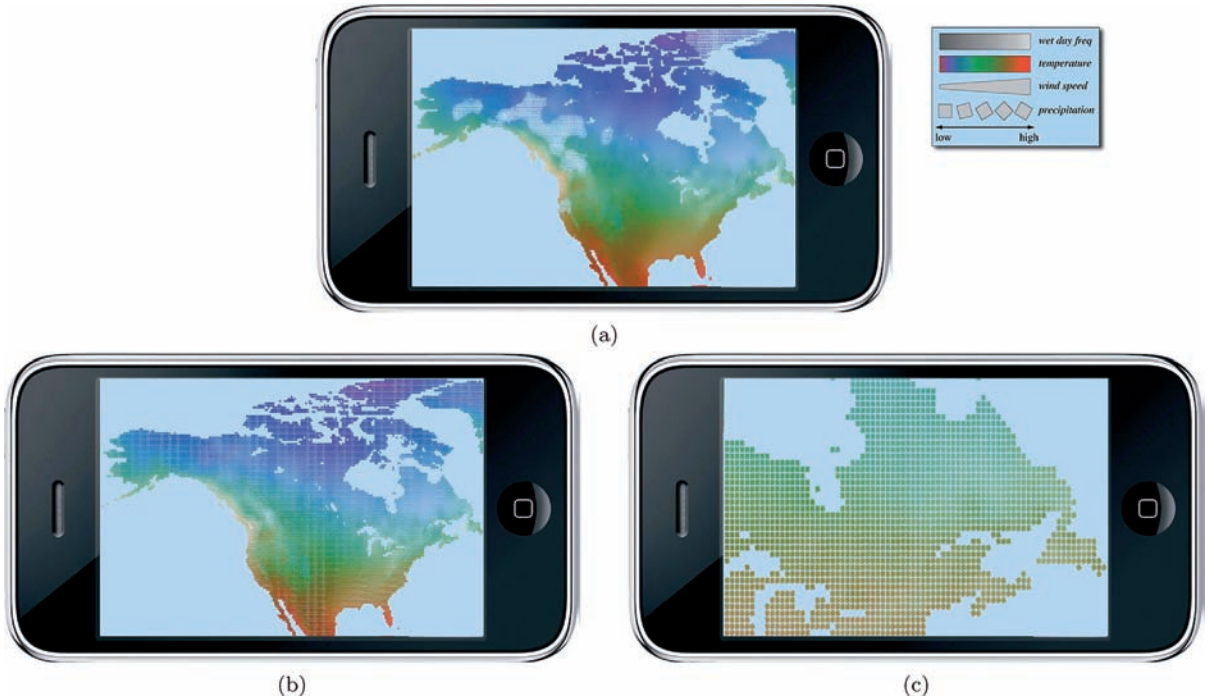

Fig. 8. Mockups of visualizations on an iPhone, view from 5.8-inches: (a) all visual features enabled; (b) luminance, hue, and orientation enabled, but with resolution and visual angle below required minimums; (c) 200% zoom producing elements above the required minimums.

distinguishability along the coastlines of Newfoundland and Nova Scotia. Although variations in hue are clearly visible, the visual angle is below the hue's minimum visual angle of 0.245°, suggesting that a higher zoom level could improve color performance. The tighter packing of elements can also cause the the perception of a target's color to shift, based on the colors of the local neighbors that surround it.

6.3.2 *Powerwall.* We also studied a 4-projector powerwall with projector resolutions of 1440×1050 pixels. Assuming a 100 pixel overlap between neighboring projectors, the resolution of the powerwall is 2780×1910 pixels, and its physical size is 96×72-inches. At a normal viewing distance of 200-inches, the visual angle of each pixel is 0.0096°. Figure 9 simulates visualizing January's historical weather conditions. The same data-feature mapping is used: *wet day frequency* → luminance, *temperature* → hue, *wind speed* → size, and *precipitation* → orientation. Each element is assigned a 4×4 resolution, with a visual angle of 0.0382°.

The left side of Figure 9 visualizes all four attributes. As before, variations in size are difficult to identify and obscure other features for small elements. On the right side of Figure 9 size has been removed. Elements now have sufficient resolution to visualize differences in luminance, hue, and orientation. Individual elements are still well below the minimum visual angle of 0.1225°, however. This is demonstrated for hue using the four test patches in the upper-left corner of the image. The leftmost and rightmost patches are solid pink and solid brown, respectively. The inner patches are made up of 1×1 pixel and 4×4 pixel checkerboard patterns of pink and brown. The checkerboard pattern is not obvious in either patch because the visual system is integrating the patches' colors into an additive result.

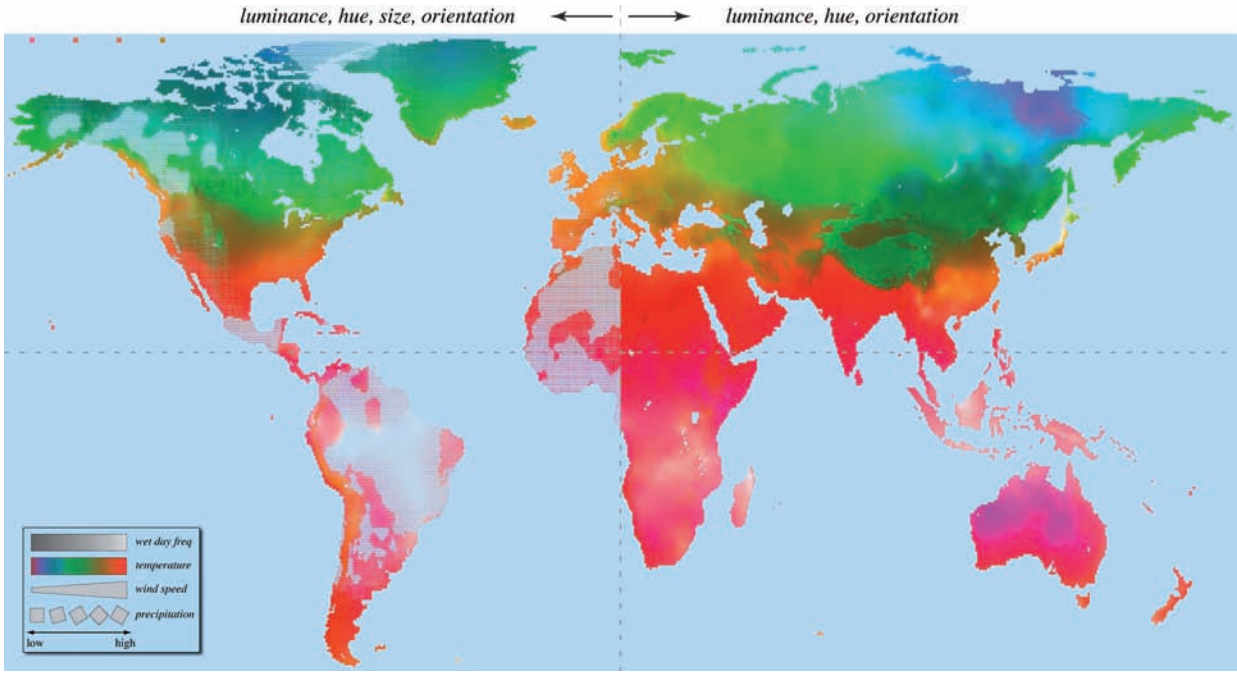

Fig. 9. Visualizations on a powerwall, view from 4-inches: (left) all visual features enabled; (right) luminance, hue, and orientation enabled, but with resolution and visual angle below required minimums.

The right side of Figure 9 show that increasing the pixel count may not lead to a corresponding increase in information content. Although the powerwall contains ∼5.8 million pixels, we cannot visualize 5.8 million data values. Our experiments suggest that we need a 12×12 resolution to produce visual angles that show differences in luminance, hue, and orientation. At this resolution, only 40,200 elements can be displayed.

One advantage of a powerwall is that it allows viewers to physically "zoom" or "pan" where they are looking. Ball et al. describe a study of physical versus virtual navigation. When available, physical navigation was preferred, and it significantly improved performance [Ball et al. 2007]. Figure 10 visualizes a subset of the powerwall at a simulated 48-inch viewing distance and a per-element visual angle of 0.1592°, slightly below (for hue) or above (for luminance and orientation) the recommended limits. Notice how the 4×4 checkerboard pattern is now distinguishable as pink and brown, while the 1×1 checkerboard is still hard to recognize.

7. CONCLUSIONS

This article describes an investigation of the limits of resolution and visual angle in visualization. We conducted psychophysical experiments to choose a small set of luminances, hues, sizes, and orientations that are easily distinguished in isolation. We then ran novel target-detection experiments to determine the minimum size of a 2D geometric glyph—both in pixel resolution and in subtended visual angle—needed to support rapid and accurate identification of the presence or absence of an element with a unique visual appearance.

Figure 6(a) shows the minimum resolutions and visual angles our experiments identified. These results were used to build a perceptual display hierarchy, a subdivision of view space to automatically

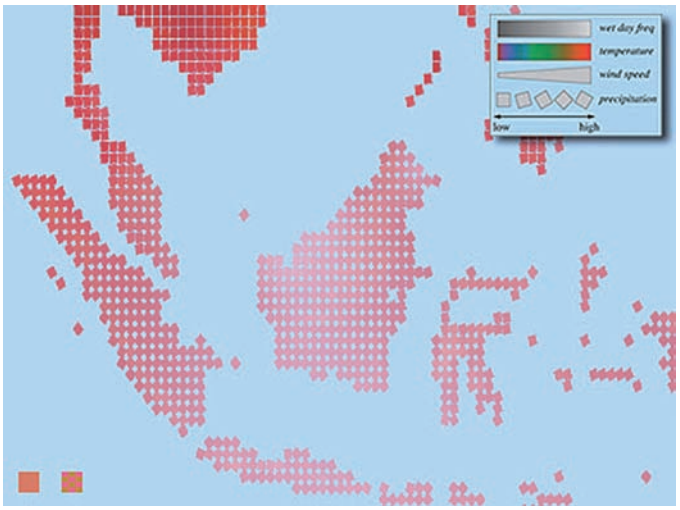

Fig. 10. A cropped view from the powerwall visualization with a simulated decrease in viewing distance, producing a resolution of 4×4 pixels and a visual angle of 0.1592°, view from 4-inches.

add or remove features from a visualization, for example, as elements grow or shrink during zooming, or as resolutions and visual angles change when the visualizations are moved between different display environments.

7.1 Future Work

Our experimental results provide a solid foundation from which to conduct additional studies. One obvious extension is to study additional visual features like saturation, contrast, flicker, direction of motion, and velocity of motion. These have been shown to be salient [Healey and Enns 1999; Huber and Healey 2005], highlighting their potential for visualizing data. In addition to this, we are currently focusing on three possibilities: (1) What additional factors should we consider to better generalize our results? (2) How does the spatial packing density of elements affect viewer performance? and (3) Can results from experiments on text legibility be applied to our goal of data visualization?

7.1.1 *Generalizability.* Our initial findings may depend on other factors that we have not yet tested. For example, field of view affects performance, with detail falling off rapidly at the periphery due to a reduction in the total number of receptor cells outside the fovea, and in their packing density—the so-called cortical magnification factor [Anstis 1998]. Researchers have studied how eccentricity affects the visual angle needed to detect different visual features. In fact, it may be possible to use existing discrimination models to estimate limits on resolution and visual angle for both foveal and peripheral vision, then compare these to our experimental findings. An element's appearance can also impact its distinguishability, for example, different shapes or heterogeneous sets of elements may be interpreted differently than collections of rectangular glyphs. Finally, JND does not necessarily predict performance, so theoretical results need to be validated with users visualizing real data for real analysis tasks. We are considering all of these issues as areas for future study.

7.1.2 *Spatial Packing Density.* A separate issue that arose during our experiments is the effect of spatial packing density. For example, during the resolution experiments the packing density decreased when the element size fell (Figure 4). This may affect viewer performance. Sagi and Julész suggested that orientation detection is a short-range phenomena, occurring within a local neighborhood defined

by average element size [Sagi and Julész 1987]. Joseph et al. also hint at a target detection advantage for dense versus sparse element packing [Joseph et al. 1997]. Simple modifications should allow us to measure the effect of spatial packing density. Advantages for higher density displays would suggest, for example, that element sizes should grow or shrink to maintain an appropriately packing density during viewpoint changes. Color surround—a shift in a target's color appearance based on its neighbors' colors—would also need to be considered.

7.1.3 *Legibility.* Research on text legibility may offer insights into visual acuity that are directly relevant to our goal of data visualization. Legge et al. have studied how visual features like contrast, size, color, and luminance affect text legibility [Legge et al. 1987, 1990]. Legge documents thresholds on a viewer's reading ability, for example, the effects of character contrast and size on reading speed, and the interaction that occurs between these two features. Although Legge's display environment differs from ours—following text moving across a display versus scanning to explore visual properties of glyphs—we believe it would be interesting, and potentially fruitful, to see whether results from legibility research extend to the visualization domain.

REFERENCES

ANSTIS, S. 1998. Picturing peripheral acuity. *Perception 27*, 7, 817–825.

BALL, R., NORTH, C., AND BOWMAN, D. 2007. Move to improve: Promoting physical navigation to increase user performance with large displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (*CHI'07*). ACM, New York, 191–200.

CAMPBELL, F. W. AND GREEN, D. G. 1965. Monocular versus binocular visual acuity. *Nature 208*, 9, 191–192.

CIE. 1978. CIE Publication No. 15, Supplement Number 2 (E-1.3.1, 1971). Official recommendations on uniform color spaces, color-difference equations, and metric color terms. Commission Internationale de L'Èclairge.

FURNAS, G. W. 1986. Generalized fisheye views. In *Proceedings of the SIGCHI* (*CHI'86*). ACM, New York, 16–34.

GLASSNER, A. S. 1995. *Principles of Digital Image Synthesis*. Morgan Kaufmann, San Francisco, CA.

GRINSTEIN, G., PICKETT, R., AND WILLIAMS, M. 1989. EXVIS: An exploratory data visualization environment. In *Proceedings of the Graphics Interface'89*. 254–261.

HEALEY, C. G. 1996. Choosing effective colours for data visualization. In *Proceedings of the 7th IEEE Visualization Conference* (*Vis'96*). IEEE, Los Alamitos, CA, 263–270.

HEALEY, C. G. AND ENNS, J. T. 1999. Large datasets at a glance: Combining textures and colors in scientific visualization. *IEEE Trans.Visualization Comput. Graph. 5*, 2, 145–167.

HUBER, D. E. AND HEALEY, C. G. 2005. Visualizing data with motion. In *Proceedings of the 16th IEEE Visualization Conference* (*Vis'05*). IEEE, Los Alamitos, CA, 527–534.

INTERRANTE, V. 2000. Harnessing natural textures for multivariate visualization. *IEEE Comput. Graph. Appl. 20*, 6, 6–11.

JOHNSON, C. R., MOORHEAD, R., MUNZNER, T., PFSITER, H., RHEINGANS, P., AND YOO, T. S. 2006. NIH/NSF visualization research challenges report. IEEE, Los Alamitos, CA.

JOSEPH, J. S., CHUN, M. M., AND NAKAYAMA, K. 1997. Attentional requirements in a 'preattentive' feature search task. *Nature 387*, 805–807.

JULÉSZ, B., GILBERT, E. N., AND SHEPP, L. A. 1973. Inability of humans to discriminate between visual textures that agree in second-order statistics—revisited. *Perception 2*, 391–405.

LAMPING, J. AND RAO, R. 1996. The hyperbolic browser: A focus+context technique for visualizing large hierarchies. *J. Visual Lang. Comput. 7*, 1, 33–55.

LEGGE, G. E., PARISH, D. H., LUEBKER, A., AND WURM, L. H. 1990. Psychophysics of reading. XI. Comparing color contrast and luminance contrast. *J. Optical Soc. Amer. A 7*, 10, 2002–2010.

LEGGE, G. E., RUBIN, G. S., AND LUEBKER, A. 1987. Psychophysics of reading. V: The role of contrast in normal vision. *Vision Res. 27*, 7, 1165–1177.

LIU, B., DIJKSTRA, T. M. H., AND OOMES, A. H. J. 2002. The beholder's share in the perception of orientation of 2D shapes. *Perception Psychophys. 64*, 8, 1227–1247.

MALIK, J. AND PERONA, P. 1990. Preattentive texture discrimination with early vision mechanisms. *J. Optical Soc. Amer. A 7*, 5, 923–932.

MARR, D. 1982. *Vision*. W. H. Freeman, San Francisco, CA.

MCCORMICK, B. H., DEFANTI, T. A., AND BROWN, M. D. 1987. Visualization in scientific computing. *Comput. Graph. 21*, 6, 1–14.

POLLACK, I. 1968. Methodological examination of the PEST (parametric estimation by sequential tracking) procedure. *Perception Psychophys*. *3*, 285–289.

RAO, A. R. AND LOHSE, G. L. 1993. Towards a texture naming system: Identifying relevant dimensions of texture. In *Proceedings of the 4th IEEE Visualization Conference* (*Vis'93*), IEEE, Los Alamitos, CA, 220–227.

RHEINGANS, P. AND TEBBS, B. 1990. A tool for dynamic explorations of color mappings. *Comput. Graph. 24*, 2, 145–146.

ROGOWITZ, B. E. AND TREINISH, L. A. 1993. An architecture for perceptual rule-based visualization. In *Proceedings of the 4th IEEE Visualization Conference* (*Vis'93*), IEEE, Los Alamitos, CA, 236–243.

SAGI, D. AND JULÉSZ, B. 1987. Short-range limitation on detection of feature differences. *Spatial Vision 2*, 1, 39–49

SMITH, P. H. AND VAN ROSENDALE, J. 1998. Data and visualization corridors report on the 1998 CVD workshop series (sponsored by DOE and NSF). Tech. rep. CACR-164, Center for Advanced Computing Research, California Institute of Technology.

THOMAS, J. J. AND COOK, K. A. 2005. *Illuminating the Path: Research and Development Agenda for Visual Analytics*. IEEE, Los Alamitos, CA.

ULRICH, R. AND MILLER, J. 2004. Threshold estimation in two-alternative forced choice (2AFC) tasks: The Spearman-Kárber method. *Perception Psychophys. 66*, 3, 517–533.

WALES, R. AND BLAKE, R. R. 1970. Rule for obtaining 75% threshold with the staircase method. *J. Optical Soc. America 60*, 2, 284–285.

WARE, C. 1988. Color sequences for univariate maps: Theory, experiments, and principles. *IEEE Comput. Graph. Appl. 8*, 5, 41–49.

WARE, C. 2012. *Information Visualization: Perception for Design* 3rd Ed., Morgan Kaufmann, San Francisco, CA.

WARE, C. AND KNIGHT, W. 1995. Using visual texture for information display. *ACM Trans. Graph. 14*, 1, 3–20.

WATSON, A. B. AND AHUMADA, JR., A. J. 2005. A standard model for foveal detection of spatial contrast. *J. Vision 5*, 9, 717–740.

WYSZECKI, G. AND STILES, W. S. 1982. *Color Science: Concepts and Methods, Quantitative Data and Formulae* 2nd Ed., Wiley, New York.

YOST, B., HACIAHMETOGLU, Y., AND NORTH, C. 2007. Beyond visual acuity: The perceptual scalability of information visualizations for large displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (*CHI'07*), ACM, New York, 101–110.

ZHANG, X. AND WANDELL, B. A. 1997. A spatial extension of CIELAB for digital color-image reproduction. *J. Soc. Inf. Display 5*, 1, 61–63.

Received July 2011; revised July 2012; accepted July 2012