Large Datasets at a Glance: Combining Textures and Colors in Scientific Visualization

Christopher G. Healey and James T. Enns

Abstract— **This paper presents a new method for using texture and color to visualize multivariate data elements arranged on an underlying height field. We combine simple texture patterns with perceptually uniform colors to increase the number of attribute values we can display simultaneously. Our technique builds multicolored perceptual texture elements (or pexels) to represent each data element. Attribute values encoded in an element are used to vary the appearance of its pexel. Texture and color patterns that form when the pexels are displayed can be used to rapidly and accurately explore the dataset. Our pexels are built by varying three separate texture dimensions: height, density, and regularity. Results from computer graphics, computer vision, and human visual psychophysics have identified these dimensions as important for the formation of perceptual texture patterns. The pexels are colored using a selection technique that controls color distance, linear separation, and color category. Proper use of these criteria guarantees colors that are equally distinguishable from one another. We describe a set of controlled experiments that demonstrate the effectiveness of our texture dimensions and color selection criteria. We then discuss new work that studies how texture and color can be used simultaneously in a single display. Our results show that variations of height and density have no effect on color segmentation, but that random color patterns can interfere with texture segmentation. As the difficulty of the visual detection task increases, so too does the amount of color on texture interference increase. We conclude by demonstrating the applicability of our approach to a real-world problem, the tracking of typhoon conditions in Southeast Asia.**

Keywords—**Color, color category, experimental design, human vision, linear separation, multivariate dataset, perception, pexel, preattentive processing, psychophysics, scientific visualization, texture, typhoon**

I. Introduction

THIS paper investigates the problem of visualizing multivariate data elements arrayed across an underlying height field. We seek a flexible method of displaying effectively large and complex datasets that encode multiple data values at a single spatial location. Examples include visualizing geographic and environmental conditions on topographical maps, representing surface locations, orientations, and material properties in medical volumes, or displaying rigid and rotational velocities on the surface of a three-dimensional object. Currently, features like hue, intensity, orientation, motion, and isocontours are used to represent these types of datasets. We are investigating the simultaneous use of perceptual textures and colors for multivariate visualization. We believe an effective combination of these features will increase the number of data values that can be shown at one time in a single display. To do this, we must first design methods for building texture and color patterns that support the rapid, accurate, and effective visualization of multivariate data elements.

We use multicolored perceptual texture elements (or pexels) to represent values in our dataset. Our texture elements are built by varying three separate texture dimensions: height, density, and regularity. Density and regularity have been identified in the computer vision literature as being important for performing texture classification [39], [40], [50]. Moreover, results from psychophysics have shown that all three dimensions are encoded in the low-level human visual system [1], [28], [51], [58]. Our pexels are colored using a technique that supports rapid, accurate, and consistent color identification. Three selection criteria are used to choose appropriate colors: color distance, linear separation, and named color category. All three criteria have been identified as important for measuring perceived color difference [3], [4], [14], [31], [60].

One of our real-world testbeds is the visualization of simulation results from studies being conducted in the Department of Zoology. Researchers are designing models of how they believe salmon feed and move in the open ocean. These simulated salmon are placed in a set of known environmental conditions, then tracked to see if their behavior mirrors that of the real fish. A method is needed for visualizing the simulation system. This method will be used to display both static (*e.g.,* environmental conditions for a particular month and year) and dynamic results (*e.g.,* a real-time display of environmental conditions as they change over time, possibly with the overlay of salmon locations and movement). We have approached the problems of dataset size and dimensionality by trying to exploit the power of the low-level human visual system. Research in computer vision and human visual psychophysics provides insight on how the visual system analyzes images. One of our goals is to select texture and color properties that will allow rapid visual exploration, while at the same time minimizing any loss of information due to interactions between the visual features being used.

Fig. 1 shows an example of our technique applied to the oceanographic dataset: environmental conditions in the northern Pacific Ocean are visualized using multicolored pexels. In this display, color represents open-ocean plankton density, height represents ocean current strength (taller for stronger), and density represents sea surface temperature (denser for warmer). Fig. 1 is only one frame from a much larger time-series of historical ocean conditions. Our choice of visual features was guided by experimental re-

C. G. Healey is with the Department of Computer Science, North Carolina State University, Raleigh, NC 27695-7534. E-mail: healey@csc.ncsu.edu.

J. T. Enns is with the Department of Psychology, University of British Columbia, Vancouver, British Columbia, Canada, V6T 1Z4. E-mail: jenns@psych.ubc.ca.

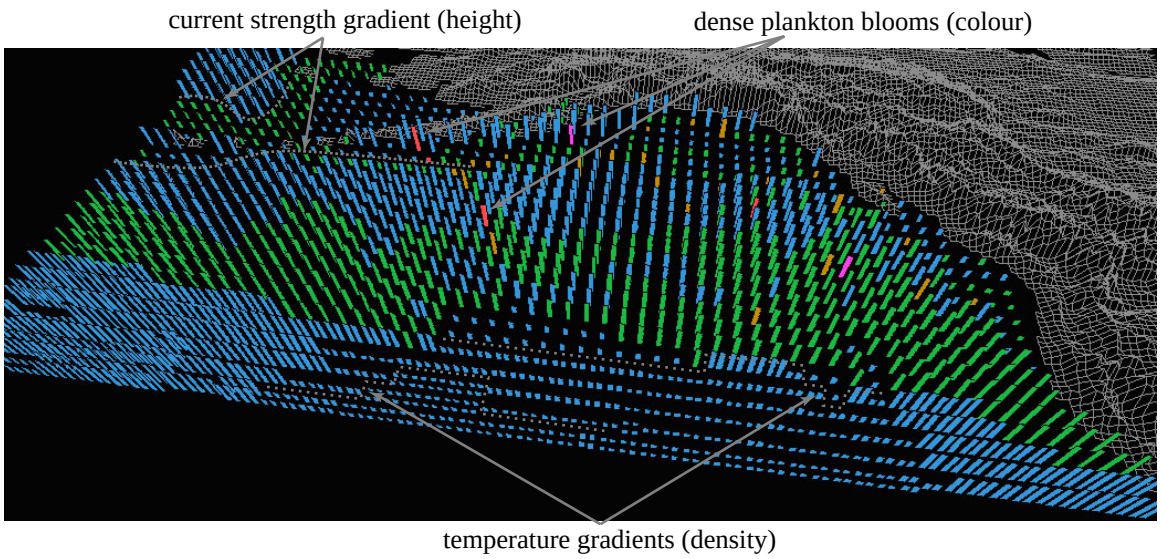

Fig. 1. Color, height, and density used to visualize open-ocean plankton density, ocean current strength, and sea surface temperature, respectively; low to high plankton densities represented with blue, green, brown, red, and purple, stronger currents represented with taller pexels, and warmer temperatures represented with denser pexels

sults that show how different color and texture properties can be used in combination to represent multivariate data elements.

Work described in this paper is an extension of earlier texture and color studies reported in [22], [23], [25]. We began our investigation by conducting a set of controlled experiments to measure user performance and identify visual interference that may occur during visualization. Individual texture and color experiments were run in isolation. The texture experiments studied the perceptual salience of and interference between the perceptual texture dimensions height, density, and regularity. The color experiments measured the effects of color distance, linear separation, and named color category on perceived color difference. Positive results from both studies led us to conduct an additional set of experiments that tested the combination of texture and color in a single display. Results in the literature vary in their description of this task: some researchers have reported that random color variation can interfere significantly with a user's ability to see an underlying texture region [8], [9], [49], while others have found no impact on performance [53], [58]. Our investigation extends this earlier work on two-dimensional texture patterns into an environment that displays height fields using perspective projections. To our knowledge, no one has studied the issue of color on texture or texture on color interference during visualization. Results from our experiments showed that we could design an environment in which color variations caused a small but statistically reliable interference effect during texture segmentation. The strength of this effect depends on the difficulty of the analysis task: tasks that are more difficult are more susceptible to color interference. Texture variation, on the other hand, caused no interference during color segmentation. We are using these results to build a collection of pexels that allow a viewer to visually explore a multivariate dataset in a rapid, accurate, and relatively effortless fashion.

We begin with a description of results from computer vision, computer graphics, and psychophysics that discuss methods for identifying and controlling the perceptual properties of texture and color. Next, we describe an area of human psychophysics concerned with modeling control of visual attention in the low-level human visual system. We discuss how the use of visual stimuli that control attention can lead to significant advantages during visualization. Section 4 gives an overview of the experiments we used to build and test our perceptual texture elements. In Section 5, we discuss how we chose to select and test our perceptual colors. Following this, we describe new experiments designed to study the combined use of texture and color. Finally, we report on practical applications of our research in Section 7, then discuss avenues for future research in Section 8.

II. Related Work

Researchers from many different areas have studied texture and color in the context of their work. Before we discuss our own investigations, we provide an overview of results in the literature that are most directly related to our interests.

A. Texture

The study of texture crosses many disciplines, including computer vision, human visual psychophysics, and computer graphics. Although each group focuses on separate problems (texture segmentation and classification in computer vision, modeling the low-level human visual system in psychophysics, and information display in graphics) they each need ways to describe accurately the texture patterns being classified, modeled, or displayed. [41] describes two

general classes of texture representation: statistical models that use convolution filters and other techniques to measure variance, inertia, entropy, or energy, and perceptual models that identify underlying perceptual texture dimensions like contrast, size, regularity, and directionality. Our current texture studies focus on the perceptual features that make up a texture pattern. In our work we demonstrate that we can use texture attributes to assist in visualization, producing displays that allow users to rapidly and accurately explore their data by analyzing the resulting texture patterns.

Different methods have been used to identify and investigate the perceptual features inherent in a texture pattern. Bela Julész [27], [28] has conducted numerous psychophysical experiments to study how a texture's first, second, and third-order statistics affect discrimination in the low-level visual system. This led to the texton theory [29], which proposes that early vision detects three types of features (or textons, as Julész called them): elongated blobs with specific visual properties (*e.g.,* hue, orientation, or length), ends of line segments, and crossing of line segments. Other psychophysical researchers have studied how visual features like color, orientation, and form can be used to rapidly and accurately segment collections of elements into spatially coherent regions [7], [8], [52], [58], [59].

Work in psychophysics has also been conducted to study how texture gradients are used to judge an object's shape. Cutting and Millard discuss how different types of gradients affect a viewer's perception of the flatness or curvature of an underlying 3D surface [13]. Three texture gradients were tested: perspective, which refers to smooth changes in the horizontal width of each texture element, compression, which refers to changes in the height to width ratio of a texture element, and density, which refers to changes in the number of texture elements per unit of solid visual angle. For most surfaces the perspective and compression gradients decrease with distance, while the density gradient increases. Cutting and Millard found that viewers use perspective and density gradients almost exclusively to identify the relative slant of a flat surface. In contrast, the compression gradient was most important for judging the curvature of undulating surfaces. Later work by Aks and Enns on overcoming perspective foreshortening in early vision also discussed the effects of texture gradients on the perceived shape of an underlying surface [1].

Work in computer vision is also interested in how viewers segment images, in part to try to develop automated texture classification and segmentation algorithms. Tamura et al. and Rao and Lohse identified texture dimensions by conducting experiments that asked observers to divide pictures depicting different types of textures (Brodatz images) into groups [39], [40], [50]. Tamura et al. used their results to propose methods for measuring coarseness, contrast, directionality, line-likeness, regularity, and roughness. Rao and Lohse used multidimensional scaling to identify the primary texture dimensions used by their observers to group images: regularity, directionality, and complexity. Haralick built grayscale spatial dependency matrices to identify features like homogeneity, contrast, and linear dependency [21]. These features were used to classify satellite images into categories like forest, woodlands, grasslands, and water. Liu and Picard used Wold features to synthesize texture patterns [35]. A Wold decomposition divides a 2D homogeneous pattern (*e.g.,* a texture pattern) into three mutually orthogonal components with perceptual properties that roughly correspond to periodicity, directionality, and randomness. Malik and Perona designed computer algorithms that use orientation filtering, nonlinear inhibition, and computation of the resulting texture gradient to mimic the discrimination ability of the low-level human visual system [37].

Researchers in computer graphics are studying methods for using texture to perform tasks such as representing surface shape and extent, displaying flow patterns, identifying spatially coherent regions in high-dimensional data, and multivariate visualization. Schweitzer used rotated discs to highlight the shape and orientation of a three-dimensional surface [47]. Grinstein et al. created a system called EXVIS that uses "stick-men" icons to produce texture patterns that show spatial coherence in a multivariate dataset [19]. Ware and Knight used Gabor filters to construct texture patterns; attributes in an underlying dataset are used to modify the orientation, size, and contrast of the Gabor elements during visualization [57]. Turk and Banks described an iterated method for placing streamlines to visualize two-dimensional vector fields [54]. Interrante displayed texture strokes to help show three-dimensional shape and depth on layered transparent surfaces; principal directions and curvatures are used to orient and advect the strokes across the surface [26]. Salisbury et al. used texturing techniques to build computer-generated pen-and-ink drawings that convey a realistic sense of shape, depth, and orientation [46]. Finally, Laidlaw described two methods for displaying a 2D diffuse tensor image with seven values at each spatial location [32]. The first method used the shape of normalized ellipsoids to represent individual tensor values. The second used techniques from oil painting [38] to represent all seven values simultaneously via multiple layers of varying brush strokes.

Visualization techniques like EXVIS [19] are sometimes referred to as "glyph-based" methods. Glyphs are graphical icons with visual features like shape, orientation, color, and size that are controlled by attributes in an underlying dataset. Glyph-based techniques range from representation via individual icons to the formation of texture and color patterns through the overlay of many thousands of glyphs. Initial work by Chernoff suggested the use of facial characteristics to represent information in a multivariate dataset [6], [10]. A face is used to summarize a data element; individual data values control features in the face like the nose, eyes, eyebrows, mouth, and jowls. Foley and Ribarsky have created a visualization tool called Glyphmaker that can be used to build visual representations of multivariate datasets in an effective, interactive fashion [16]. Glyphmaker uses a glyph editor and glyph binder to create glyphs, to arrange them spatially, and to bind attributes to their visual

properties. Levkowitz described a prototype system for combining colored squares to produce patterns to represent an underlying multivariate dataset [33]. Other techniques such as the normalized ellipsoids of Laidlaw [32], the Gabor elements of Ware [57], or even the pexels described in this paper might also be classified as glyphs, although we prefer to think of them as texture-based visualization methods.

B. Color

As with texture, color has a rich history in the areas of computer graphics and psychophysics. In graphics, researchers have studied issues related to color specification, color perception, and the selection of colors for information representation during visualization. Work in psychophysics describes how the human visual system mediates color perception.

A number of different color models have been built in computer graphics to try to support the unambiguous specification of color [60]. These models are almost always three-dimensional, and are often divided into a device-dependent class, where a model represents the displayable colors of a given output device, and a device-independent class, where a model provides coordinates for colors from the visible color spectrum. Common examples of device-dependent models include monitor RGB and CMYK. Common examples of device-independent models include CIE XYZ, CIE LUV, and CIE Lab. Certain models were designed to provide additional functionality that can be used during visualization. For example, both CIE LUV and CIE Lab provide rough perceptual uniformity; that is, the Euclidean distance between a pair of colors specified in these models roughly corresponds to perceived color difference. These models also provide a measure of isoluminance, since their L-axis is meant to correspond to perceived brightness.

Previous work has also addressed the issue of constructing color scales for certain types of data visualization. For example, Ware and Beatty describe a simple color visualization technique for displaying correlation in a five-dimensional dataset [56]. Ware has also designed a method for building continuous color scales that control color surround effects [55]. The color scales use a combination of luminance and hue variation that allows viewers to determine the value associated with a specific color, and to identify the spatial locations of peaks and valleys (*i.e.,* to see the shape) in a 2D distribution of an attribute's values. Controlling color surround ensures a small, near-constant perceptual error effect from neighboring colors across the entire range of the color scale. Robertson described user interface techniques for visualizing the range of colors a display device can support using perceptual color models [44]. Rheingans and Tebbs have built a system that allows users to interactively construct a continuous color scale by tracing a path through a 3D color model [43]. This technique allows users to vary how different values of an attribute map onto the color path. Multiple colors can be used to highlight areas of interest within an attribute, even when those areas constitute only a small fraction of the attribute's full range of allowable values. Levkowitz and Herman designed a locally optimal color scale that maximizes the just-noticeable color difference between neighboring pairs of colors [34]. The result is a significantly larger number of just-noticeably different colors in their color scales, compared to standard scales like red-blue, rainbow, or luminance.

Recent work at the IBM Thomas J. Watson Research Center has focused on a rule-based visualization tool [45]. Initial research addressed the need for rules that take into account how a user perceives visual features like hue, luminance, height, and so on. These rules are used to guide or restrict a user's choices during data-feature mapping. The rules use various metadata, for example, the visualization task being performed, the visual features being used, and the spatial frequency of the data being visualized. A specific example of one part of this system is the colormap selection tool from the IBM Visualization Data Explorer [5]. The selection tool uses perceptual rules and metadata to limit the choice of colormaps available to the user.

Finally, psychophysicists have been working to identify properties that affect perceived color difference. Two important discoveries include the linear separation [3], [4], [14] and color category [31] effects. The linear separation theory states that if a target color can be separated from all the other background colors being displayed with a single straight line in color space, it will be easier to detect (*i.e.,* its perceived difference from all the other colors will increase) compared to a case where it can be formed by a linear combination of the background colors. D'Zmura's initial work on this phenomena [14] showed that a target color could be rapidly identified in a sea of background elements uniformly colored one of two colors (*e.g.,* an orange target could be rapidly identified in a sea of red elements, or in a sea of yellow elements). The same target, however, was much more difficult to find when the background elements used both colors simultaneously (*e.g.,* an orange target could not be rapidly identified in a sea of red and yellow elements). This second case is an example of a target color (orange) that is a linear combination of its background colors (red and yellow). The color category effect suggests that the perceived difference between a pair of colors increases when the two colors occupy different named color regions (*i.e.,* one lies in the "blue" region and one lies in the "purple" region, as opposed to both in blue or both in purple). We believe both results may need to be considered to guarantee perceptual uniformity during color selection.

C. Combined Texture and Color

Although texture and color have been studied extensively in isolation, much less work has focused on their combined use for information representation. An effective method of displaying color and texture patterns simultaneously would increase the number of attributes we can represent at one time. The first step towards supporting this goal is the determination of the amount of visual interference that occurs between these features during visualization.

Experiments in psychophysics have produced interesting but contradictory answers to this question. Callaghan reported asymmetric interference of color on form during texture segmentation: a random color pattern interfered with the identification of a boundary between two groups of different forms, but a random form pattern had no effect on identifying color boundaries [8], [9]. Triesman, however, showed that random variation of color had no effect on detecting the presence or absence of a target element defined by a difference in orientation (recall that directionality has been identified as a fundamental perceptual texture dimension) [53]. Recent work by Snowden [49] recreated the differing results of both Callaghan and Triesman. Snowden ran a number of additional experiments to test the effects of random color and stereo depth variation on the detection of a target line element with a unique orientation. As with Callaghan and Triesman, results differed depending on the target type. Search for a single line element was rapid and accurate, even with random color or depth variation. Search for a spatial collection of targets was easy only when color and depth were fixed to a constant value. Random variation of color or depth produced a significant reduction in detection accuracy. Snowden suggests that the visual system wants to try to group spatially neighboring elements with common visual features, even if this grouping is not helpful for the task being performed. Any random variation of color or depth interferes with this grouping process, thereby forcing a reduction in performance.

These results leave unanswered the question of whether color variation will interfere with texture identification during visualization. Moreover, work in psychophysics studied two-dimensional texture segmentation. Our pexels are arrayed over an underlying height field, then displayed in 3D using a perspective projection. Most of the research to date has focused on color on texture interference. Less work has been conducted to study how changes in texture dimensions like height, density, or regularity will affect the identification of data elements with a particular target color. The question of interference in this kind of height-field environment needs to be addressed before we can recommend methods for the combined use of color and texture.

III. Perceptual Visualization

An important requirement for any visualization technique is a method for rapid, accurate, and effortless visual exploration. We address this goal by using what is known about the control of human visual attention as a foundation for our visualization tools. The individual factors that govern what is attended in a visual display can be organized along two major dimensions: bottom-up (or stimulus driven) versus top-down (or goal directed).

Bottom-up factors in the control of attention include the limited set of features that psychophysicists have identified as being detected very quickly by the human visual system, without the need for search. These features are often called preattentive, because their detection occurs rapidly and accurately, usually in an amount of time independent of the total number of elements being displayed. When applied properly, preattentive features can be used to perform different types of exploratory analysis. Examples include searching for data elements with a unique visual feature, identifying the boundaries between groups of elements with common features, tracking groups of elements as they move in time and space, and estimating the number of elements with a specific feature. Preattentive tasks can be performed in a single glance, which corresponds to 200 milliseconds (ms) or less. As noted above, the time required to complete the task is independent of the number of data elements being displayed. Since the visual system cannot choose to refocus attention within this timeframe, users must complete their task using only a "single glance" at the image.

Fig. 2 shows examples of both types of target search. In Fig. 2a-2d the target, a red circle, is easy to find. Here, the target contains a preattentive feature unique from the background distracters: color (red versus blue) or shape (circle versus square). This unique feature is used by the low-level visual system to rapidly identify the presence or absence of the target. Unfortunately, an intuitive combination of these results can lead to visual interference. Fig. 2e and 2f simulate a two-dimensional dataset where one attribute is encoded with color (red or blue), and the other is encoded with shape (circle or square). Although these features worked well in isolation, searching for a red circle target in a sea of blue circles and red squares is significantly more difficult. In fact, experiments have shown that search time is directly proportional to the number of elements in the display, suggesting that viewers are searching small subgroups of elements (or even individual elements themselves) to identify the target. In this example the low-level visual system has no unique feature to search for, since circular elements (blue circles) and red elements (red squares) are also present in the display. The visual system cannot integrate preattentively the presence of multiple visual features (circular and red) at the same spatial location. This is a very simple example of a situation where knowledge of preattentive vision would have allowed us to avoid displays that actively interfere with our analysis task.

In spite of the perceptual salience of the target in Fig. 2a-2d, bottom-up influences cannot be assumed to operate independently of the current goals and attentional state of the observer. Recent studies have demonstrated that many of the bottom-up factors only influence perception when the observer is engaged in a task in which they are expected or task-relevant (see the review by [15]). For example, a target defined as a color singleton will "pop out" of a display only when the observer is looking for targets defined by color. The same color singleton will not influence perception when observers are searching exclusively for luminance defined targets. Sometimes observers will fail completely to see otherwise salient targets in their visual field, either because they are absorbed in the performance of a cognitively-demanding task [36], there are a multitude of other simultaneous salient visual events [42], or because the salient event occurs during an eye movement or other change in viewpoint [48]. Therefore, the control of atten-

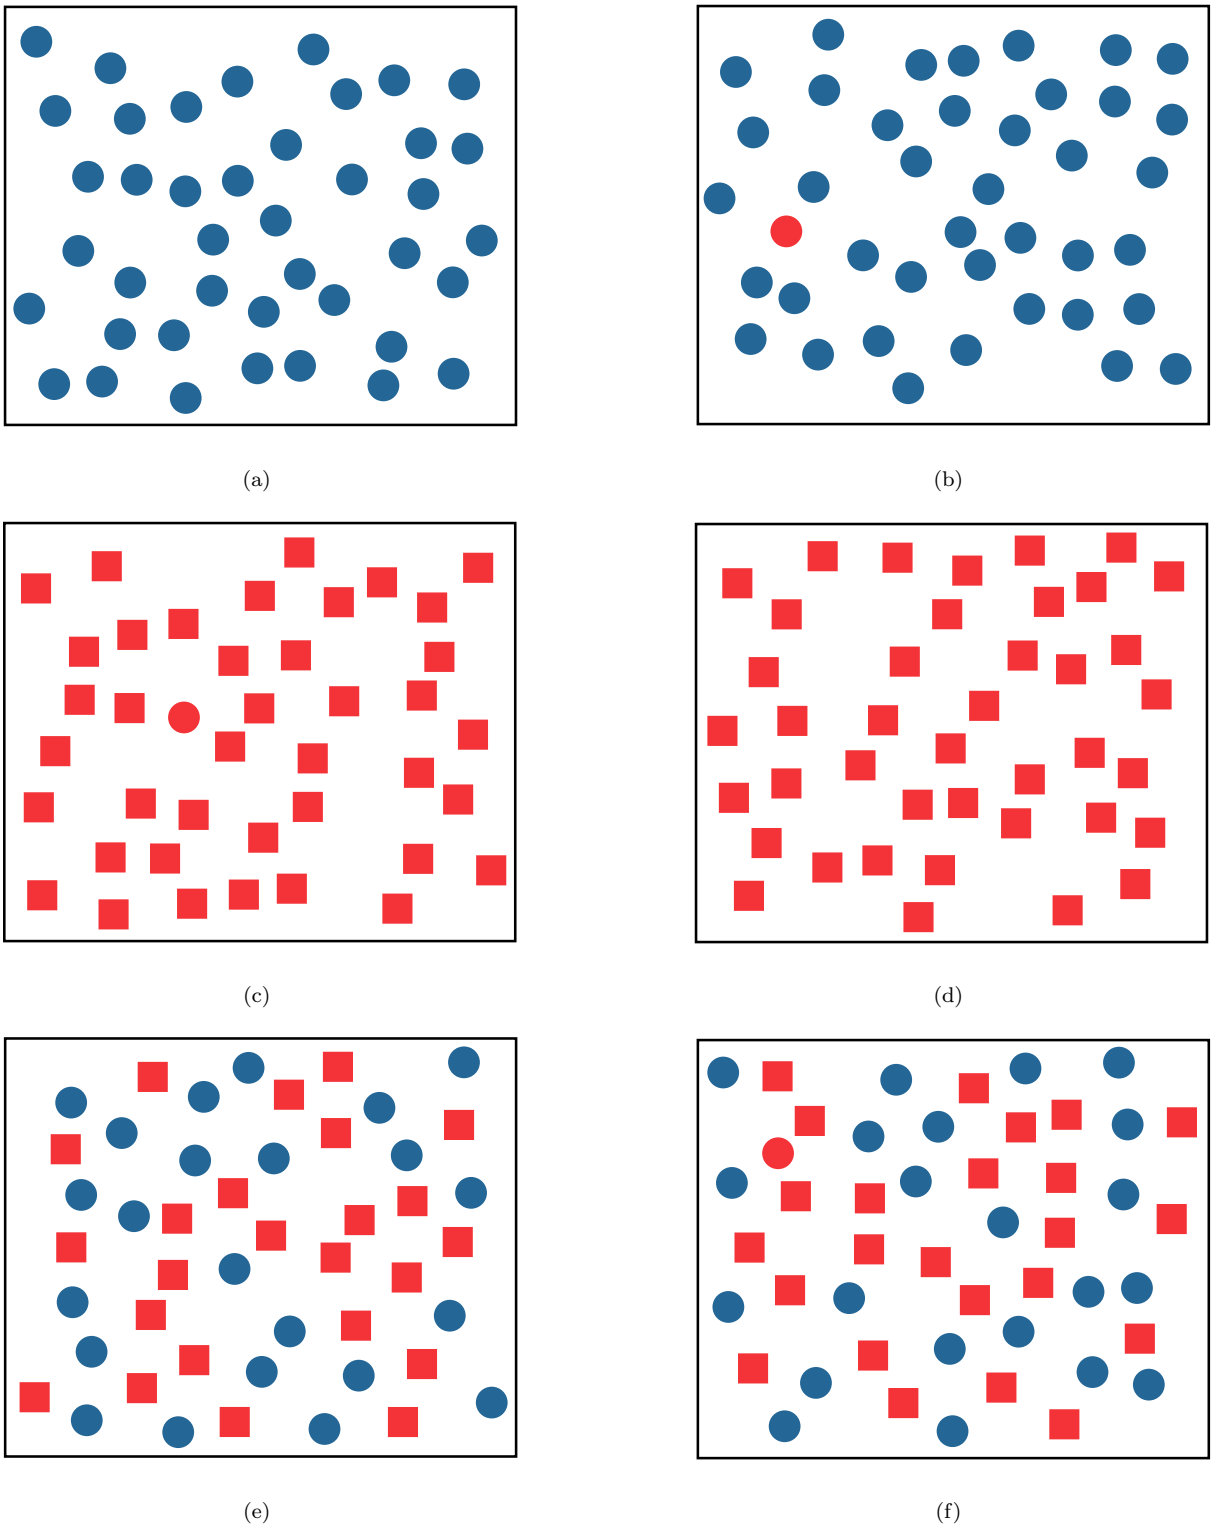

Fig. 2. Examples of target search: (a, b) identifying a red target in a sea of blue distracters is rapid and accurate, target absent in (a), present in (b); (c, d) identifying a red circular target in a sea of red square distracters is rapid and accurate, target present in (c), absent in (d); (e, f) identifying the same red circle target in a combined sea of blue circular distracters and red square distracters is significantly more difficult, target absent in (e), present in (f)

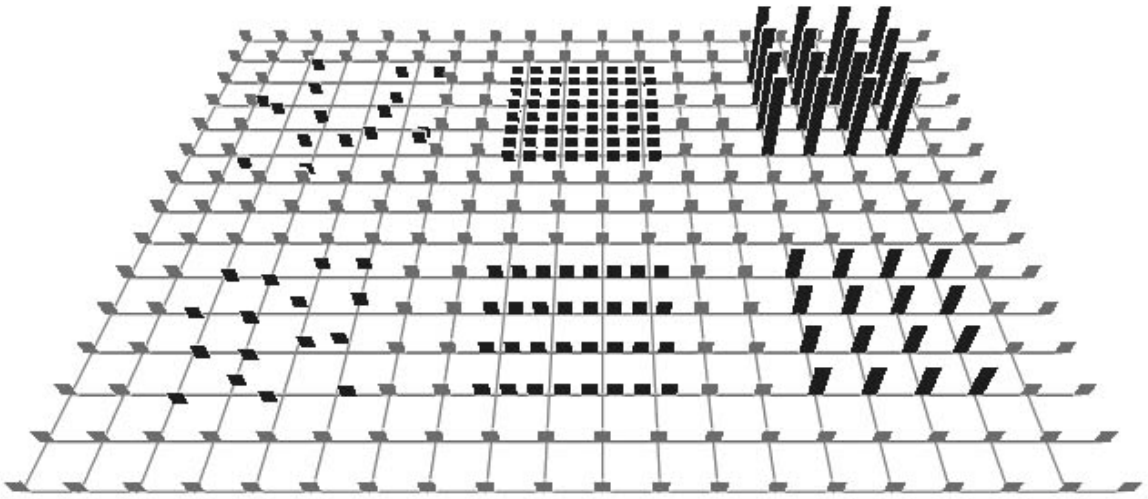

Fig. 3. A background array of short, sparse, regular pexels; the lower and upper groups on the left contain irregular and random pexels, respectively; the lower and upper groups in the center contain dense and very dense pexels; the lower and upper groups to the right contain medium and tall pexels

tion must always be understood as an interaction between bottom-up and top-down mechanisms.

Our research is focused on identifying relevant results in the vision and psychophysical literature, then extending these results and integrating them into a visualization environment. Tools that make use of preattentive vision offer a number of important advantages during multivariate visualization:

1. Visual analysis is rapid, accurate, and relatively effortless since preattentive tasks can be completed in 200 ms or less. We have shown that tasks performed on static displays extend to a dynamic environment where data frames are shown one after another in a movie-like fashion [24] (*i.e.,* tasks that can be performed on an individual display in 200 ms can also be performed on a sequence of displays shown at five frames a second).
2. The time required for task completion is independent of display size (to the resolution limits of the display). This means we can increase the number of data elements in a display with little or no increase in the time required to analyze the display.
3. Certain combinations of visual features cause interference patterns that mask information in the low-level visual system. Our experiments are designed to identify these situations. This means our visualization tools can be built to avoid data-feature mappings that might interfere with the analysis task.

Since preattentive tasks are rapid and insensitive to display size, we believe visualization techniques that make use of these properties will support high-speed exploratory analysis of large, multivariate datasets. Care must be taken, however, to ensure that we choose data-feature mappings that maximize the perceptual salience of all the data attributes being displayed.

We are currently investigating the combined use of two important and commonly used visual features: texture and color. Previous work in our laboratory has identified methods for choosing perceptual textures and colors for multivariate visualization. Results from vision and psychophysics on the simultaneous use of both features are mixed: some researchers have reported that background color patterns mask texture information, and vice-versa, while others claim that no interference occurs. Experiments reported in this paper are designed to test for color-texture interactions during visualization. A lack of interference would suggest that we could combine both features to simultaneously encode multiple attributes. The presence of interference, on the other hand, would place important limitations on the way in which visual attributes should be mapped onto data attributes. Visualization tools based on these findings would then be able to display textures with the appropriate mapping of data dimensions to visual attributes.

IV. Pexels

One of the main goals of multivariate visualization is to display multiple attribute values at a single spatial location without overwhelming the user's ability to comprehend the resulting image. Researchers in vision, psychophysics, and graphics have been studying how the visual system analyzes texture patterns. We wanted to know whether perceptual texture dimensions could be used to represent multivariate data elements during visualization. To this end, we designed a set of perceptual texture elements, or pexels, that support the variation of three separate texture dimensions: density, regularity, and height. Density and regularity have been identified in the literature as primary texture dimensions [39], [40], [50]. Although height might not be considered an "intrinsic textural cue", we note that height is one aspect of element size, and that size is an important property of a texture pattern. Results from psychophysical experiments have shown that differences in

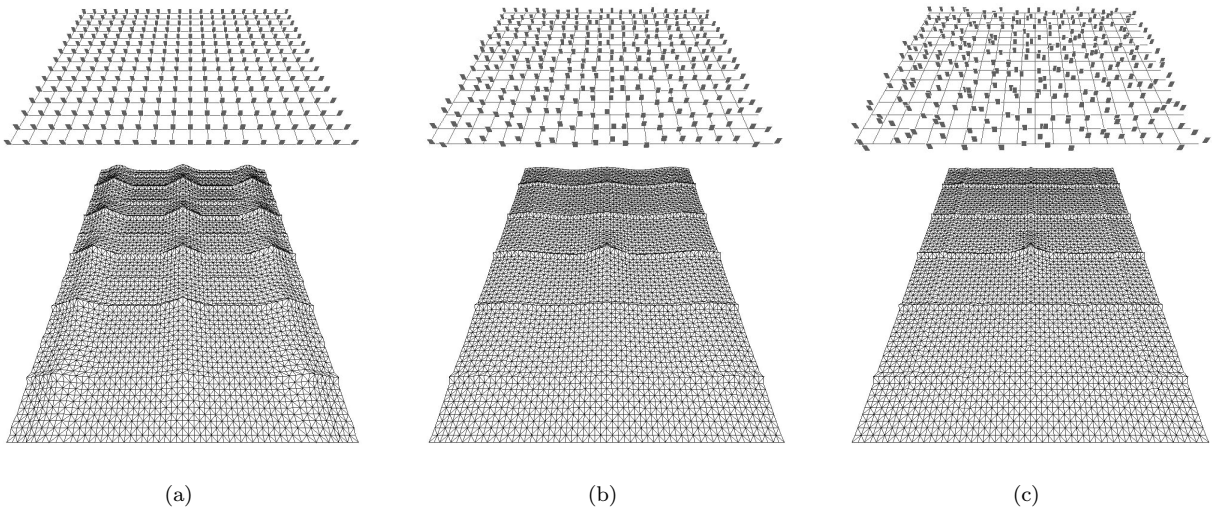

Fig. 4. Three displays of pexels with different regularity and a 5×3 patch from the center of the corresponding autocorrelation graphs: (a) a completely regular display, resulting in sharp peaks of height 1.0 at regular intervals in the autocorrelation graph; (b) a display with irregularly-spaced pexels, peaks in the graph are reduced to a maximum height between 0.3 and 0.4; (c) a display with randomly-spaced pexels, resulting in a completely flat graph except at (0,0) and where underlying grid lines overlap

height are detected preattentively [51], moreover, viewers properly correct for perspective foreshortening when they perceive that elements are being displayed in 3D [1]. We wanted to build three-dimensional pexels that "sit up" on the underlying surface. This allows for the possibility of applying various orientations (another important texture dimension) to a pexel.

Our pexels look like a collection of one or more upright paper strips. Each element in the dataset is represented by a single pexel. The user maps attributes in their dataset to density (which controls the number of strips in a pexel), height, and regularity. The attribute values for a particular element can then control the appearance of its pexel. When all the pexels for a particular data frame are displayed, they form texture patterns that represent the underlying structure of the dataset.

Fig. 3 shows an example of regularity, density, and height varied across three discrete values. Each pexel in the original array (shown in gray) is short, sparse, and regular. The lower and upper patches on the left of the array (shown in black) contain irregular and random pexels, respectively. The lower and upper patches in the middle of the array contain dense and very dense pexels. The lower and upper patches on the right contain medium and tall pexels.

A. Ordering Texture Dimensions

In order to use height, density, and regularity during visualization, we needed an ordinal ranking for each dimension. Height and density both have a natural ordering: shorter comes before taller, and sparser comes before denser.

Although viewers can often order regularity intuitively, we required a more formal method for measurement. We chose to use autocorrelation to rank regularity. This technique measures the second-order statistic of a texture pattern. Psychophysicists have reported that a change in regularity produces a corresponding change in a texture's second order statistic [27], [28], [30]. Intuitively, autocorrelating an image shifts two copies of the image on top of one another, to see how closely they can be matched. If the texture is made up of a regular, repeating pattern it can be shifted in various ways to exactly overlap with itself. As more and more irregularity is introduced into the texture, the amount of overlap decreases, regardless of where we shift the copies. Consider two copies of an image A and B, each with a width of N and a height of M pixels. The amount of autocorrelation that occurs when A is overlaid onto B at offset (t, u) is:

$$C(t,u) = \frac{1}{K} \sum_{x=1}^{N} \sum_{y=1}^{M} (A[x,y] - \overline{A})(B[x+t, y+u] - \overline{B}) \quad (1)$$

$$K = NM \sqrt{\sigma^2(A)} \sqrt{\sigma^2(B)} \quad (2)$$

$$\overline{A} = \frac{1}{NM} \sum_{x=1}^{N} \sum_{y=1}^{M} A[x,y] \quad (3)$$

$$\sigma^2(A) = \frac{1}{NM} \sum_{x=1}^{N} \sum_{y=1}^{M} (A[x,y] - \overline{A})^2 \quad (4)$$

with $\overline{B}$ and $\sigma^2(B)$ computed in a similar fashion. Elements in A that do not overlap with B are wrapped to the opposite side of B (*i.e.,* elements in A lying above the top of B wrap back to the bottom, elements lying below the bottom of B wrap back to the top, similarly for elements to the left or right of B).

As a practical example, consider Fig. 4a (pexels on a regular underlying grid), Fig. 4b (pexels on an irregular grid),

and Fig. 4c (pexels on a random grid). Irregular and random pexels are created by allowing each strip in the pexel to walk a random distance (up to fixed maximum) in a random direction from its original anchor point. Autocorrelation was computed on the orthogonal projection of each image. A 5×3 patch from the center of the corresponding autocorrelation graph is shown beneath each of the three grids. Results in the graphs mirror what we see in each display, that is, as randomness increases, peaks in the autocorrelation graph decrease in height. In Fig. 4a peaks of height 1.0 appear at regular intervals in the graph. Each peak represents a shift that places pexels so they exactly overlap with one another. The rate of increase towards each peak differs in the vertical and horizontal directions because the elements in the graph are rectangles (*i.e.,* taller than they are wide), rather than squares. In Fig. 4b, the graph has the expected sharp peak at (0,0). It also has gentle peaks at shift positions that realign the grid with itself. The peaks are not as high as for the regular grid, because the pexels no longer align perfectly with one another. The sharp vertical and horizontal ridges in the graph represent positions where the underlying grid lines exactly overlap with one another (the grid lines showing the original position of each pexel are still regular in this image). The height of each gentle peak ranges between 0.3 and 0.4. Increasing randomness reduces again the height of the peaks in the correlation graph. In Fig. 4c, no peaks are present, apart from (0,0) and the sharp ridges that occur when the underlying grid lines overlap with one another. The resulting correlation values suggests that this image is "more random" than either of its predecessors. Informal tests with a variety of regularity patterns showed a near-perfect match between user-chosen rankings and rankings based on our autocorrelation technique. Autocorrelation on the perspective projections of each grid could also be computed. The tall peaks and flattened results would be similar to those in Fig. 4, although the density of their spacing would change near the top of the image due to perspective compression and foreshortening.

B. Pexel Salience and Interference

We conducted experiments to test the ability of each texture dimension to display effectively an underlying data attribute during multivariate visualization. To summarize, our experiments were designed to answer the following three questions:

1. Can the perceptual dimensions of density, regularity, and height be used to show structure in a dataset through the variation of a corresponding texture pattern?
2. How can we use the dataset's attributes to control the values of each perceptual dimension?
3. How much visual interference occurs between each of the perceptual dimensions when they are displayed simultaneously?

C. Experiments

We designed texture displays to test the detectability of six different target types: taller, shorter, denser, sparser, more regular, and more irregular. For each target type, a number of parameters were varied, including exposure duration, texture dimension salience, and visual interference. For example, during the "taller" experiment, each display showed a 20×15 array of pexels rotated 45° about the X-axis. Observers were asked to determine whether the array contained a group of pexels that were taller than all the rest. The following conditions varied:

- *target-background pairing:* some displays showed a medium target in a sea of short pexels, while others showed a tall target in a sea of medium pexels; this allowed us to test whether some target defining attributes were generally more salient than others,
- *secondary texture dimension:* displays contained either no background variation (every pexel was sparse and regular), a random variation of density across the array, or a random variation of regularity across the array; this allowed us to test for background interference during target search,
- *exposure duration:* displays were shown for 50, 150, or 450 ms; this allowed us to test for a reduction in performance when exposure duration was decreased, and
- *target patch size:* target groups were either 2×2 pexels or 4×4 pexels in size; this allowed us to test for a reduction in performance for smaller target patches.

The heights, densities, and regularities we used were chosen through a set of pilot studies. Two patches were placed side-by-side, each displaying a pair of heights, densities, or regularities. Viewers were asked to answer whether the patches were different from one another. Response times for correct answers were used as a measure of performance. We tested a range of values for each dimension, although the spatial area available for an individual pexel during our experiments limited the maximum amount of density and irregularity we were able to display. The final values we chose could be rapidly and accurately differentiated in this limited setting.

The experiments that tested the other five target types (shorter, denser, sparser, regular, and irregular) were designed in a similar fashion, with one exception. Experiments testing regularity had only one target-background pairing: a target of regular pexels in a sea of random pexels (for the regular experiment), or a target of random pexels in a sea of regular pexels (for the irregular experiment). The pilot studies used to select values for each dimension showed that users had great difficulty discriminating an irregular patch from a random patch. This was due in part to the small spatial area available to each pexel.

Our pilot studies produced experiments that tested three separate heights (short, medium, and tall), three separate densities (sparse, dense, and very dense) and two separate regularities (regular and random). Examples of two display types (taller and regular) are shown in Fig. 5. Both displays include target pexels. Fig. 5a contains a 2×2 target group of medium pexels in a sea of short pexels. The density of each pexel varies across the array, producing an underlying density pattern that is clearly visible. This display type simulates two dimensional data elements being

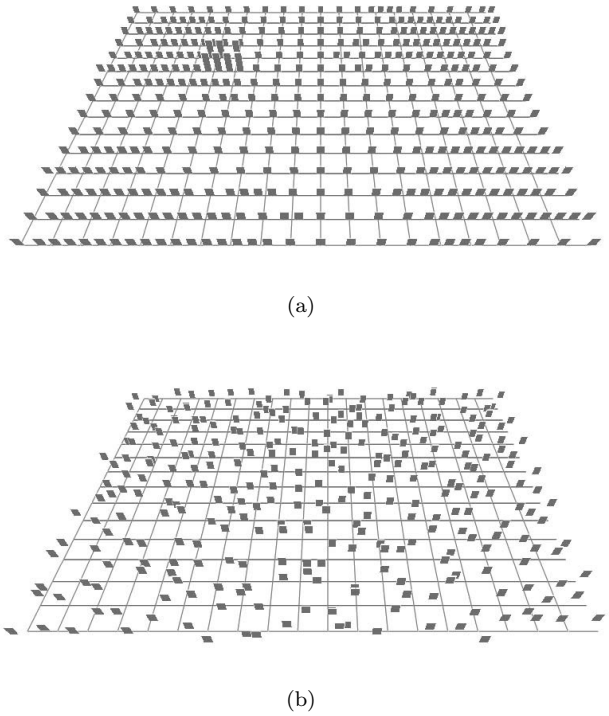

Fig. 5. Two display types from the taller and regular pexel experiments: (a) a target of medium pexels in a sea of short pexels with a background density pattern (2×2 target group located left of center); (b) a target of regular pexels in a sea of irregular pexels with no background texture pattern (2×2 target group located 3 grid steps right and 7 grid steps up from the lower-left corner of the array)

visualized with height as the primary texture dimension and density as the secondary texture dimension. Recall that the number of paper strips in a pexel depends on its density. Since three of the target pexels in Fig. 5a are dense, they each display two strips. The remaining pexel is sparse, and therefore displays a only single strip. Fig. 5b contains a 2×2 target group of regular pexels in a sea of random pexels, with a no background texture pattern. The taller target in Fig. 5a is very easy to find, while the regular target in Fig. 5b is almost invisible.

D. Results

Detection accuracy data were analyzed using a multifactor analysis of variance (ANOVA). A complete description of our analysis and statistical findings is available in [22], [23], [25]. In summary, we found:

1. Taller target regions were identified rapidly (*i.e.,* 150 ms or less) with very high accuracy (approximately 93%); background density and regularity patterns produced no significant interference.
2. Shorter, denser, and sparser targets were more difficult to identify than taller targets, although good results were obtained at both 150 and 450 ms (82.3%, 94.0%, and 94.7% for shorter, denser, and sparser targets with no background variation at 150 ms). This was not surprising, since similar results have been documented by [51] and [1] using displays of texture on a two-dimensional plane.
3. Background variation in non-target attributes produced small, but statistically significant, interference effects. These effects tended to be largest when target detectability was lowest. For example, density and regularity interfered with searching for shorter targets; height and regularity interfered with searching for sparser targets; but only height interfered with the (easier to find) denser targets.
4. Irregular target regions were difficult to identify at 150 and 450 ms, even with no secondary texture pattern (approximately 76%). Whether this accuracy level is sufficiently high will depend on the application. In contrast, regular regions were invisible under these conditions; the percentage of correct responses approached chance (*i.e.,* 50%) in every condition.

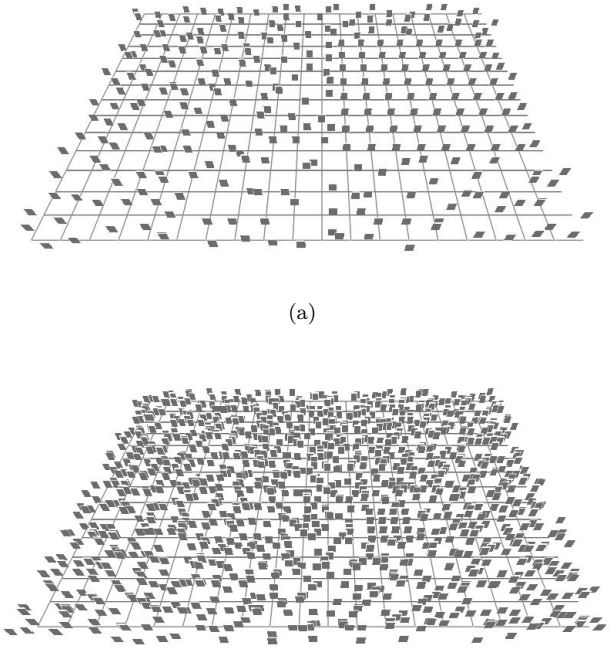

Fig. 6. Two displays with a regular target, both displays should be compared with the target shown in Fig. 5b: (a) larger target, an 8×8 target in a sea of sparse, random pexels; (b) denser background, a 2×2 target in a sea of dense, random pexels (target group located right of center)

Our poor detection results for regularity were unexpected, particularly since vision algorithms that perform texture classification use regularity as one of their primary decision criteria [35], [39], [40], [50]. We confirmed that our results were not due to a difference in our definition of regularity; the way we produced irregular patches matches the methods described by [20], [28], [30], [39], [40], [50]. It may be that regularity is important for classifying different textures, but not for the type of texture segmentation that we are performing. Informal post-experiment investigations showed that we could improve the salience of a reg-

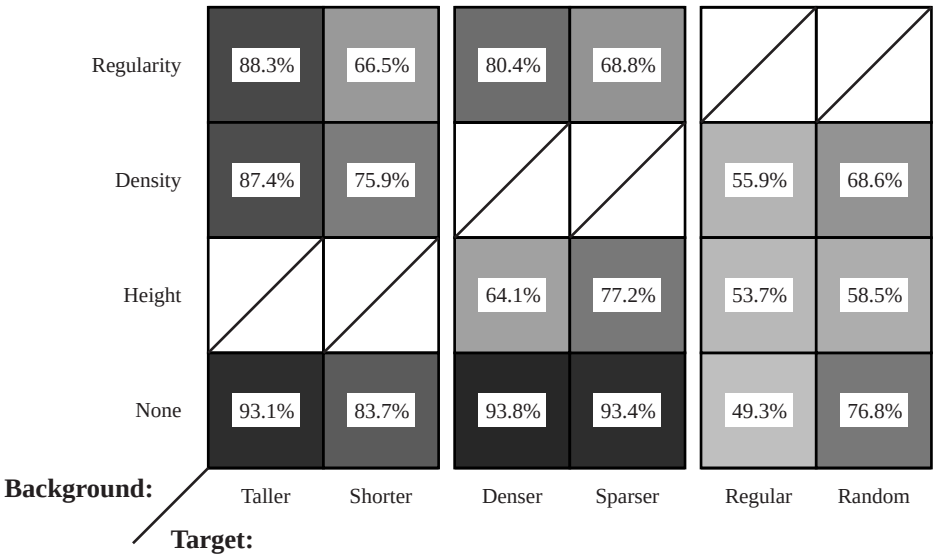

Fig. 7. A table showing the percentage of correct responses for each target-background pairing; target type along the horizontal axis, background type along the vertical axis; darker squares represent pairings with a high percentage of correct responses; blank entries with diagonal slashes indicate target-background pairings that do not exist

ular (or irregular) patch by increasing its size (Fig. 6a), or by increasing the minimum pexel density to be very dense (Fig. 6b). However, neither of these solutions is necessarily useful. There is no way to guarantee that data values will cluster together to form the large spatial regions needed for regularity detection. If we constrain density to be very dense across the array, we lose the ability to vary density over an easily identifiable range. This reduces the dimensionality of our pexels to two (height and regularity), producing a situation that is no better than when regularity is difficult to identify. Because of this, we normally choose to display an attribute with low importance using regularity. While differences in regularity cannot be detected consistently by the low-level visual system, in many cases users may be able to see the changes when areas of interest in the dataset are identified and analyzed in a focused or attentive fashion.

Fig. 7 shows average subject performance as a table representing each target-background pair. Target type varies along the horizontal axis, while background type varies along the vertical axis. Darker squares represent target-background pairings with highly accurate subject performance. The number in the center of each square reports the percentage of correct responses averaged across all subjects.

V. Perceptual Colors

In addition to our study of pexels, we have examined methods for choosing multiple individual colors. These experiments were designed to select a set of n colors such that:

1. Any color can be detected preattentively, even in the presence of all the others.
2. The colors are equally distinguishable from one another; that is, every color is equally easy to identify.

We also tested for the maximum number of colors that can be displayed simultaneously, while still satisfying the above requirements. Background research suggested that we needed to consider three separate selection criteria: color distance, linear separation, and color category.

A. Color Distance

Perceptually balanced color models are often used to measure perceived color difference between pairs of colors. Examples include CIE LUV, CIE Lab, Munsell, and the Optical Society of America Uniform Color Space. We used CIE LUV to measure color distance. Colors are specified in this model using three axes: L^*, u^*, and v^*. L^* encodes luminance, while u^* and v^* encode chromaticity (u^* and v^* correspond roughly to the red-green and blue-yellow opponent color channels). CIE LUV provides two important properties. First, colors with the same L^* are isoluminant, that is, they have roughly the same perceived brightness. Second, the Euclidean distance between a pair of colors corresponds roughly to their perceived color difference. Given two colors x and y in CIE LUV space, the perceived difference measured in ΔE^* units is:

$$\Delta E^*_{xy} = \sqrt{(\Delta L^*_{xy})^2 + (\Delta u^*_{xy})^2 + (\Delta v^*_{xy})^2} \qquad (5)$$

Our techniques do not depend on CIE LUV; we could have chosen to use any perceptually balanced color model. We picked CIE LUV in part because it is reasonably well known, and in part because it is recommended by the Commission Internationale de L'Éclairage (CIE) as the appropriate model to use for CRT displays [11].

B. Linear Separation

Results from vision and psychophysics suggest that colors that are linearly separable are significantly easier to

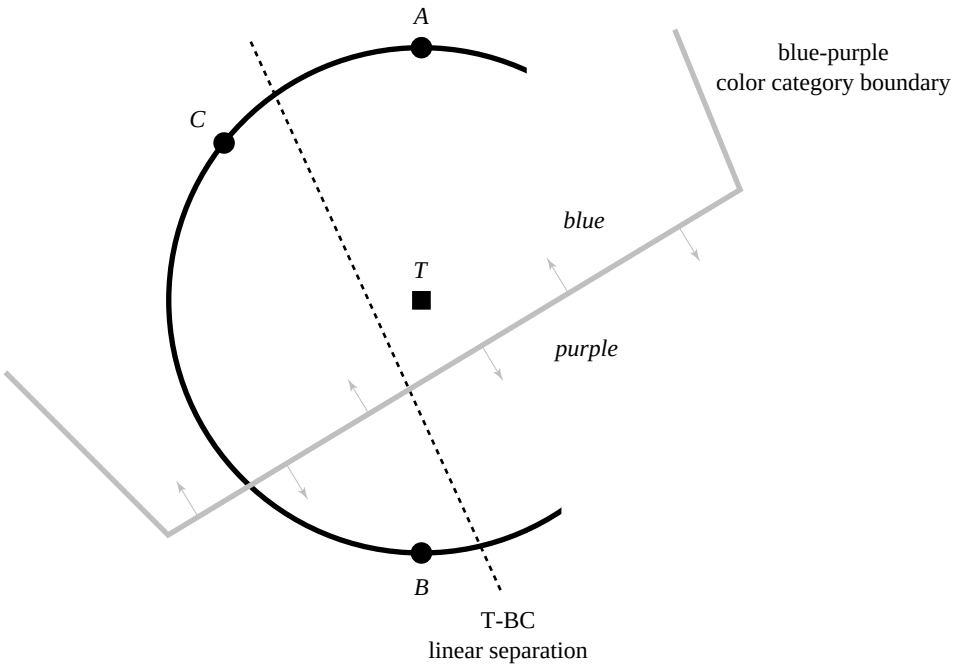

Fig. 8. A small, isoluminant patch within the CIE LUV color model, showing a target color T and three background distracter colors A, B, and C; note that T is collinear with A and B, but can be separated by a straight line from B and C; note also that T, A, and C occupy the "blue" color region, while B occupies the "purple" color region

distinguish from one another. Initial work on this problem was reported in [14]. These results were subsequently confirmed and strengthened by [3], [4] who showed that a perceptually balanced color model could not be used to overcome the linear separation effect.

As an example, consider a target color T and three background distracter colors A, B, and C shown in CIE LUV space in Fig. 8. Since the Euclidean distances TA, TB, and TC are equal, the perceived color difference between T and A, B, and C should also be roughly equal. However, searching for a target element colored T in a sea of background elements colored A and B is significantly more difficult than searching for T in a sea of elements colored B and C. Experimental results suggest that this occurs because T is collinear with A and B, whereas T can be separated by a straight line in color space from B and C. Linear separation increases perceived color difference, even when a perceptual color model is used to try to control that difference.

C. Color Category

Recent work reported by Kawai et al. showed that, during their experiments, the named categories in which people place individual colors can affect perceived color difference [31]. Colors from different named categories have a larger perceived color difference, even when Euclidean distance in a perceptually balanced color model is held constant.

Consider again the target color T and two background distracter colors A and B shown in CIE LUV space in Fig. 8. Notice also that this region of color space has been divided into two named color categories. As before, the Euclidean distances TA and TB are equal, yet finding an element colored T in a sea of background elements colored A is significantly more difficult than finding T in a sea of elements colored B. Kawai et al. suggest this is because both T and A lie within a color category named "blue", while B lies within a different category named "purple". Colors from different named categories have a higher perceived color difference, even when a perceptual color model is used to try to control that difference.

D. Color Selection Experiments

Our first experiment selected colors by controlling color distance and linear separation, but not color category. The reasons for this were twofold. First, traditional methods for subdividing a color space into named color regions are tedious and time-consuming to run. Second, we were not convinced that results from [31] were important for our color selection goals. If problems occurred during our initial experiment, and if those problems could be addressed by controlling color category during color selection, this would both strengthen the results of [31] and highlight their applicability to the general color selection task.

We selected colors from the boundary of a maximum-radius circle embedded in our monitor's gamut. The circle was located on an isoluminant slice through the CIE LUV color model. Previous work reported in [7], [9] showed that a random variation of luminance can interfere with the identification of a boundary between two groups of differently colored elements. Holding the perceived luminance of each color constant guaranteed variations in performance would not be the result of a random luminance effect. Fig. 9 shows an example of selecting five colors about the circumference of the maximum-radius circle inscribed within our

monitor's gamut at $L^* = 61.7$. Since colors are located equidistant around the circle, every color has a constant distance d to its two nearest neighbors, and a constant distance l to the line that separates it from all the other colors.

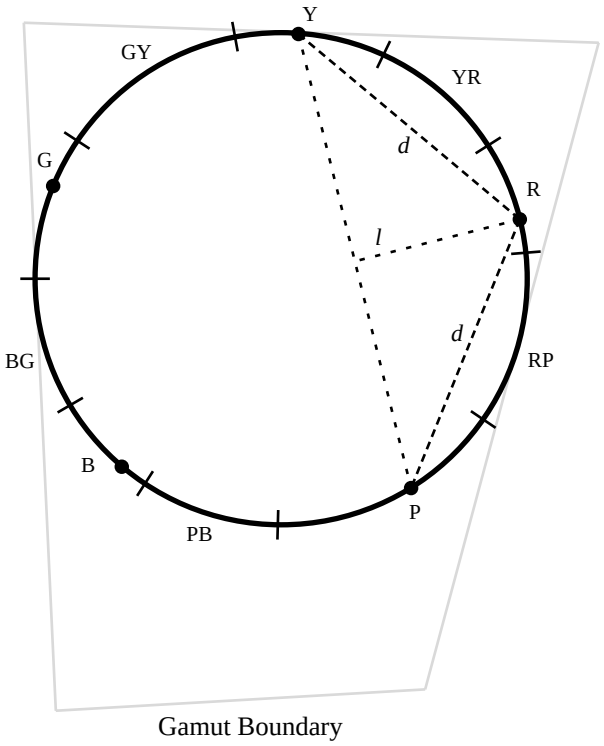

Fig. 9. Choosing colors from the monitor's gamut, the boundary of the gamut at $L^* = 61.7$ represented as a quadrilateral, along with the maximum inscribed circle centered at $(L^*, u^*, v^*) = (67.1, 13.1, -0.98)$, radius $70.5\Delta E^*$; five colors chosen around the circle's circumference; each element has a constant color distance d with its two neighbors, and a constant linear separation l from the remaining (non-target) elements; the circle's circumference has been subdivided into ten named categories, corresponding to the ten hue names from the Munsell color model

We split the experiment into four studies that displayed three, five, seven, and nine colors simultaneously. This allowed us to test for the maximum number of colors we could show while still supporting preattentive identification. Displays in each study were further divided along the following conditions:

- *target color:* each color being displayed was tested as a target, for example, during the three-color study some observers searched for a red target in a sea of green and blue distracters, others search for a blue target in a sea of red and green distracters, and the remainder searched for a green target in a sea of red and blue distracters; asymmetric performance (that is, good performance for some colors and poor performance for others) would indicate that constant distance and separation are not sufficient to guarantee equal perceived color difference, and
- *display size:* experiment displays contained either 17, 33, or 49 elements; any decrease in performance when display size increased would indicate that the search task is not preattentive.

At the beginning of an experiment session observers were asked to search a set of displays for an element with a particular target color. Observers were told that half the displays would contain an element with the target color, and half would not. They were then shown a sequence of experiment displays that contained multiple colored squares randomly located on an underlying 9×9 grid. Each display remained onscreen until the observer indicated via a keypress whether a square with the given target color was present or absent. Observers were told to answer as quickly as possible without making mistakes.

E. Results

Observers were able to detect all the color targets rapidly and accurately during both the three-color and five-color studies; the average error rate was 2.5%, while the average response times ranged from 459 to 661 ms (response times exceeded the normal preattentive limit of 200 ms because they include the time required for observers to enter their responses on the keyboard). Increasing the display size had no significant effect on response time.

Observers had much more difficulty identifying certain colors during the seven-color (Fig. 10a) and nine-color studies. Response times increased and accuracy decreased during both studies. More importantly, the time required to detect certain colors (*e.g.,* light green and dark green in the seven-color study) was directly proportional to display size. This indicates observers are searching serially through the display to find the target element. Other colors exhibited relatively flat response time curves. These asymmetric results suggest that controlling color distance and linear separation alone is not enough to guarantee a collection of equally distinguishable colors.

F. Color Category Experiments

We decided to try to determine whether named color categories could be used to explain the inconsistent results from our initial experiment. To do this, we needed to subdivide a color space (in our case, the circumference of our maximum radius circle) into named color regions. Traditional color naming experiments divide the color space into a fine-grained collection of color samples. Observers are then asked to name each of the samples. We chose to use a simpler, faster method designed to measure the amount of overlap between a set of named color regions. Our technique runs in three steps:

1. The color space is automatically divided into ten named color regions using the Munsell color model. The hue axis of the Munsell model is specified using the ten color names red, yellow-red, yellow, green-yellow, green, blue-green, blue, purple-blue, purple, and red-purple (or R, YR, Y, GY, G, BG, B, PB, P, and RP). Colors are converted to Munsell space, then assigned their hue name within that space (Fig. 9).
2. Representative colors from each of the ten named regions are selected. We chose the color at the center of each region to act as the representative color for that region.

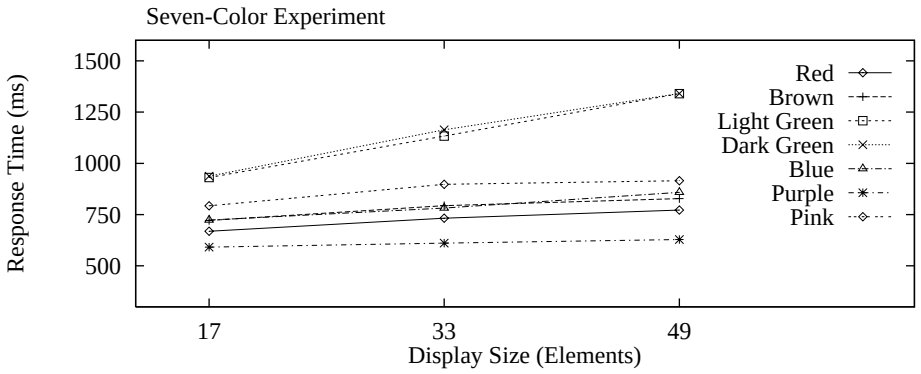

(a)

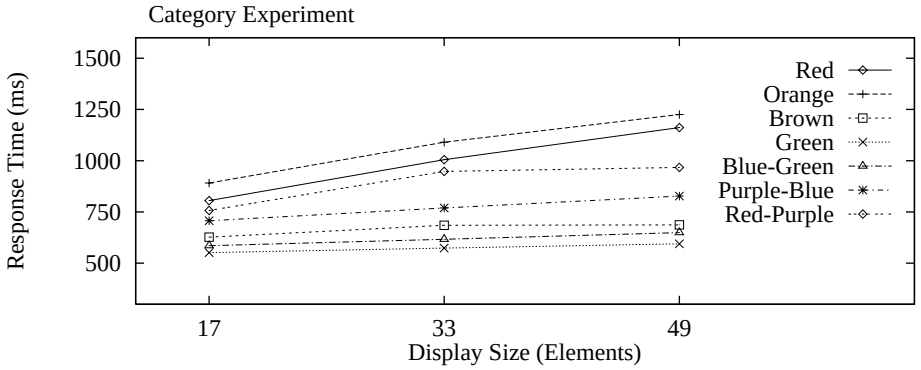

(b)

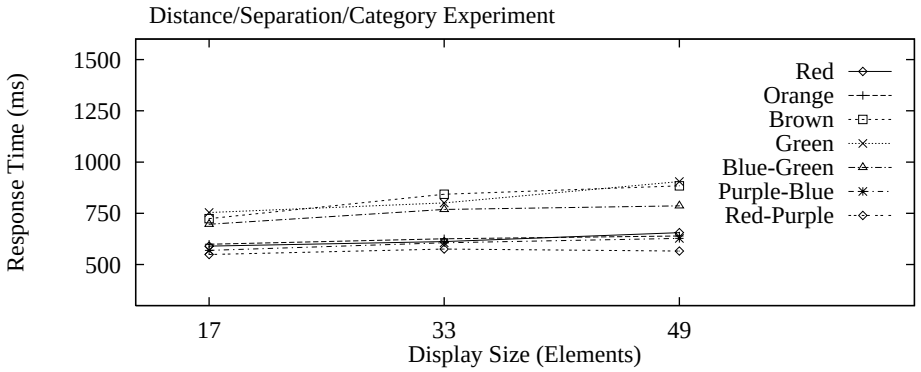

(c)

Fig. 10. Graphs showing averaged subject response times for three of the six studies: (a) response time as a function of display size (*i.e.,* total number of elements shown in each display) for each target from the seven-color study; (b) response times for each target from the color category study; (c) response times for each target from the combined distance-separation-category study

3. Observers are asked to name each of the representative colors. The amount of overlap between the names chosen for the representative colors for each region defines the amount of "category overlap" that exists between the regions.

Consider Table I, which lists the percentage of observers who selected a particular name for six of the representative colors. For example, the table shows that representative colors from P and R overlap only at the "pink" name. Their overlap is not that strong, since neither P nor R are strongly classified as pink. The amount of overlap is computed by multiplying the percentages for the common name, giving a P-R overlap of $5.2\% \times 26.3\% = 0.014$. A closer correspondence of user-chosen names for a pair of regions results in a stronger category similarity. For example, nearly all observers named the representative colors from the G and GY regions as "green". This resulted in an overlap of 0.973. Representative colors that overlap over multiple names are combined using addition, for example, YR and Y overlapped in both the "orange" and "brown"

TABLE I
A LIST OF SIX REPRESENTATIVE COLORS FOR THE COLOR REGIONS PURPLE, RED, YELLOW-RED, YELLOW, GREEN-YELLOW, AND GREEN, AND THE PERCENTAGE OF OBSERVERS WHO CHOSE A PARTICULAR NAME FOR EACH REPRESENTATIVE COLOR

	purple	*magenta*	*pink*	*red*	*orange*	*brown*	*yellow*	*green*
P	86.9%	2.6%	5.2%					
R			26.3%	71.0%				
YR				5.3%	86.8%	7.9%		
Y					2.6%	44.7%	47.4%	
GY								97.3%
G								100.0%

names, resulting in a YR-Y overlap of (86.8% × 2.6%) + (7.9% × 44.7%) = 0.058.

G. Color Category Results

When we compared the category overlap values against results from our seven and nine-color studies, we found that the amount of overlap between the target color and its background distracters provided a strong indication of performance. Colors that worked well as targets had low category overlap with all of their distracter colors. Colors that worked poorly had higher overlap with one or more of their distracter colors. A measure of rank performance to total category overlap produced correlation values of 0.821 and 0.762 for the seven and nine-color studies, respectively. This suggests that our measure of category overlap is a direct predictor of subject performance. Low category overlap between the target color and all of its background distracters produces relatively rapid subject performance. High category overlap between the target color and one or more background distracters results in relatively slow subject performance.

These results might suggest that color category alone can be used to choose a set of equally distinguishable colors. To test this, we selected seven new colors that all had low category overlap with one another, then reran the experiments. Results from this new set of colors were as poor as the original seven-color study (Fig. 10b). The seven new colors were located at the centers of their named categories, so their distances and linear separations varied. The colors with the longest response times had the smallest distances and separations. This suggests that we need to maintain at least a minimum amount of distance and separation to guarantee acceptable identification performance.

In our last experiment, we chose a final set of seven colors that tried to satisfy all three selection criteria. The categories in which the colors were located all had low overlap with one another. Colors were shifted within their categories to provide as large a distance and linear separation as possible. We also tried to maintain constant distances and linear separations for all the colors. Results from this final experiment were encouraging (Fig. 10c). Response times for each of the colors acting as a target were similar, with little or no effect from increased display size. The mean response error was also significantly lower than during the previous two seven-color experiments. We concluded that up to seven isoluminant colors can be displayed simultaneously while still allowing for rapid and accurate identification, but only if the colors satisfy proper color distance, linear separation, and color category guidelines.

VI. Combining Texture and Color

Previous work in our laboratory focused on selecting perceptual textures and colors in isolation. Clearly, we would like to use multicolored pexels during visualization. The ability to combine both features effectively would increase the number of attributes we can visualize simultaneously. Results in the literature are mixed on how this might be achieved. Some researchers have reported that task irrelevant variation in color has no effect on texture discrimination [51], [58], while others have found exactly this kind of interference [8], [9], [49]. Moreover, we are not aware of any studies that address whether there is interference from random variation in texture properties when discrimination is based on color. Experiments are therefore needed that examine possible interference effects in both directions, that is, effects of color variation on texture discrimination and effects of texture variation on color discrimination.

A. Experiments

In order to investigate these issues, we designed a new set of psychophysical experiments. Our two specific questions were:

1. Does random variation in pexel color influence the detection of a region of target pexels defined by height or density?
2. Does random variation in pexel height or density influence the detection of a region of target pexels defined by color?

We chose to ignore regularity, since it performed poorly as a target defining property during all phases of our original texture experiments [23], [25]. We chose three different colors using our perceptual color selection technique [22], [23]. Colors were initially selected in the CIE LUV color space, then converted to our monitor's RGB gamut. The three colors corresponded approximately to red (monitor RGB = 246, 73, 50), green (monitor RGB = 49, 144, 21) and blue (monitor RGB = 82, 109, 168). Our new experiments were constructed around a set of conditions similar to those used during the original texture experiments.

For color targets, we varied:

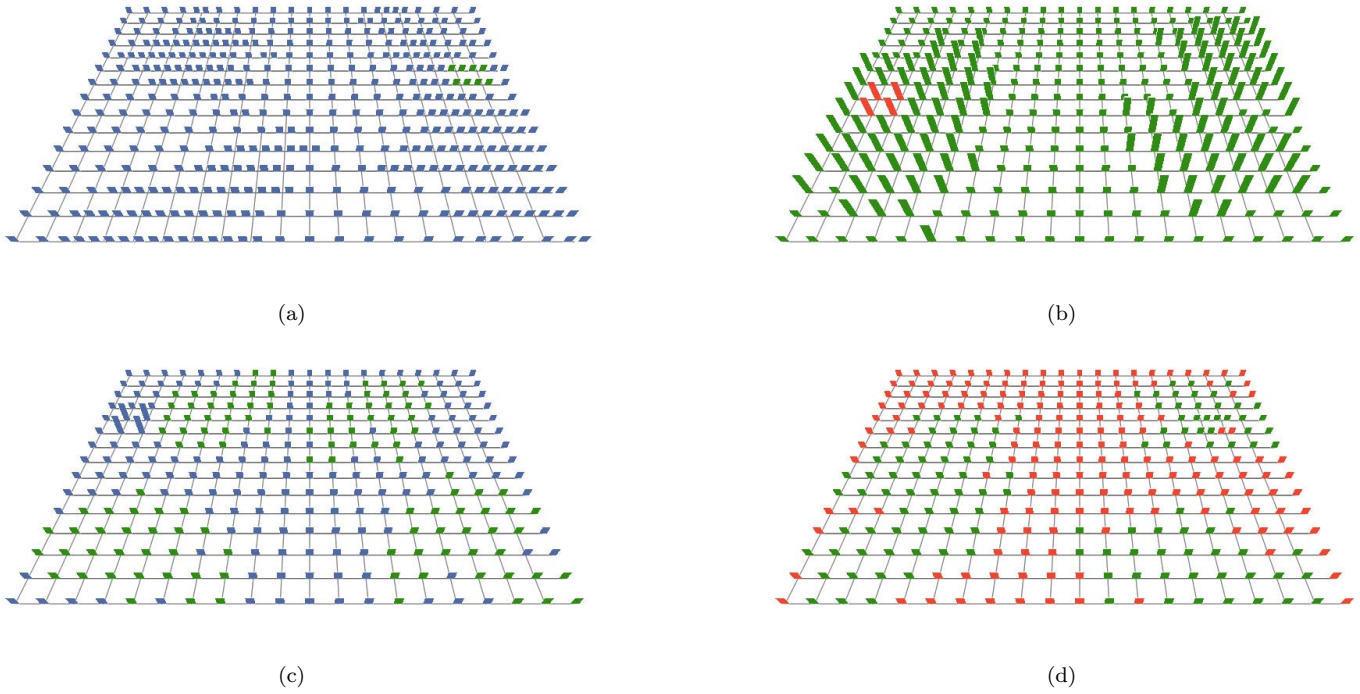

Fig. 11. Four displays from the combined color-texture experiments, printed colors may not match exactly on-screen colors used during our experiments: (a) a green target in a sea of blue pexels with background density variation; (b) a red target in a sea of green pexels with background height variation; (c) a tall target with background blue-green color variation; (d) a dense target with background green-red color variation

• *target-background pairing:* some displays contained a green target region in a sea of blue pexels, while others contained a red target region in a sea of green pexels (Fig. 11a and 11b); two different pairings were used to increase the generality of the results,
• *secondary dimension:* displays contained either no background variation (*e.g.,* every pexel was sparse and short), a random variation of density across the array, or a random variation of height across the array; this allowed us to test for interference from two different texture dimensions during target detection based on color,
• *exposure duration:* displays were shown for either 50, 150, or 450 ms; this allowed us to see how detection accuracy was influenced by the exposure duration of the display, and
• *target patch size:* target regions were either 2×2 pexels or 4×4 pexels in size. This allowed us to examine the influence of all the foregoing factors at both relatively difficult (2×2) and easy (4×4) levels of target detectability.

Two texture dimensions (height and density) were studied, and each involved two different target types: taller and shorter (for height) and denser and sparser (for density). For each type of target, we designed an experiment that tested a similar set of conditions. For example, in the taller experiment we varied:

• *target-background pairing:* half the displays contained a target region of medium pexels in a sea of short pexels, while the other half contained a target region of tall pexels in a sea of medium pexels; two different pairings were used to increase the generality of the results,
• *secondary dimension:* the displays contained pexels that were either a constant gray or that varied randomly between two colors; when color was varied, half the displays contained blue and green pexels, while the other half of the displays contained green and red pexels (Fig. 11c),
• *exposure duration:* displays were shown for 50, 150, or 450 ms, and
• *target patch size:* target groups were either 2×2 pexels or 4×4 pexels in size.

Fig. 11 shows examples of four experiment displays. Fig. 11a and 11b contain a green target in a sea of blue pexels, and a red target in a sea of green pexels, respectively. Density varies in the background in Fig. 11a, while height varies in Fig. 11b. Fig. 11c contains a tall target with a blue-green background color pattern. Fig. 11d contains a dense target with a green-red background color pattern. Any background variation that is present can pass through a target. This occurs in Fig. 11d, where part of the target is red and part is green. Note also that, as described for Fig. 5, the number of paper strips in an individual pexel depends on its density.

The colors we used during our experiments were chosen in CIE LUV color space. A simple set of formulas can be used to convert from CIE LUV to CIE XYZ (a standard device-independent color model), and from there to our monitor's color gamut. To move from LUV to XYZ:

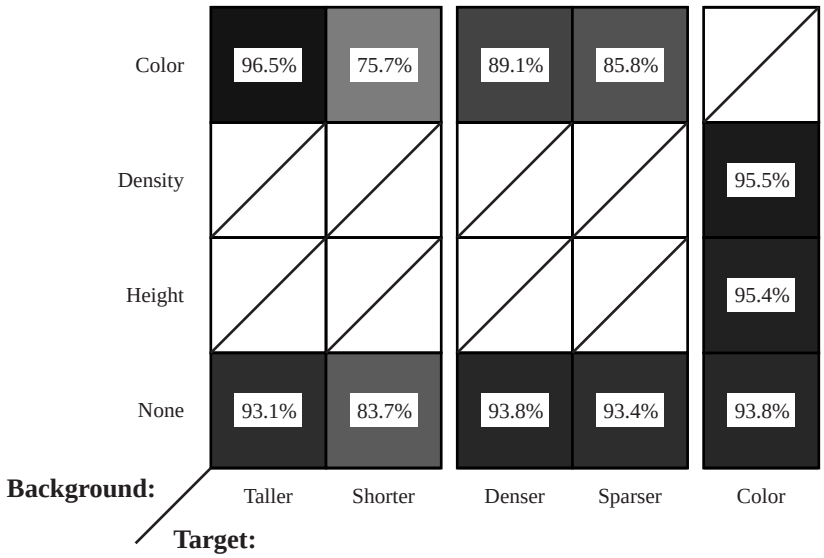

Fig. 12. A table showing the percentage of correct responses for each target-background pairing; target type along the horizontal axis, background type along the vertical axis; darker squares represent pairings with a high percentage of correct responses; results for taller, shorter, denser, and sparser with no background variation are from the original texture experiments; blank entries with diagonal slashes indicate target-background pairings that did not exist during the combined color-texture experiments

$$\begin{aligned} Y &= \left(\frac{L^* + 16}{116}\right)^3 Y_w \\ X &= \frac{9u'}{4v'}Y \\ Z &= \frac{3}{v'}Y - 5Y - \frac{3u'}{4v'}Y \end{aligned} \tag{6}$$

where L^*, u', and v' are used to specify a color in CIE LUV (u' and v' are simple respecifications of u^* and v^*), and Y_w represents the luminance of a reference white point. We then built a conversion matrix to map colors from CIE XYZ into our monitor's color gamut. This was done by obtaining the chromaticities of our monitor's red, green, and blue triads, then measuring the luminance of the monitor's maximum intensity red, green, and blue with a spot photometer. These values are needed to convert colors from a device-independent space (*i.e.,* CIE XYZ) into device-dependent coordinates (*i.e.,* our monitors RGB color space). All of our experiments were displayed on a Sony Trinitron monitor with CIE XYZ chromaticities $(x_r, y_r) = (0.625, 0.340)$, $(x_g, y_g) = (0.280, 0.595)$, and $(x_b, y_b) = (0.155, 0.070)$. The luminances of maximum intensity red, green, and blue were $Y_r = 5.5$, $Y_g = 16.6$, $Y_b = 2.8$. This produced an XYZ to monitor RGB conversion matrix of:

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 0.131 & 0.057 & 0.021 \\ -0.044 & 0.081 & 0.002 \\ 0.003 & 0.008 & 0.033 \end{bmatrix}^{-1} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \tag{7}$$

For a complete description of how the conversion formulas are built, we refer the reader to any of [17], [18], [60].

Ten users participated as observers in each of the two color and four texture experiments. Each observer had normal or corrected acuity. Observers who completed the color experiments were also tested for color blindness [12]. Observers were provided with an opportunity to practice before each experiment. This helped them become familiar with the task and the duration of the displays. Before each testing session began, observers were told that half the displays would contain a target, and half would not. We used a Macintosh computer with an 8-bit color lookup table to run our experiments. Responses (either "target present" or "target absent") for each display an observer was shown were recorded for later analysis.

B. Results

Mean percentage target detection accuracy was the measure of performance. Observer responses were collected, averaged, and analyzed using multi-factor ANOVA. In summary, we found:

1. Color targets were detected rapidly (*i.e.,* at 150 ms) with very high accuracy (96%). Background variation in height and density produced no interference effects in this detection task.
2. Detection accuracy for targets defined by density or height were very similar to results from our original texture experiments [23], [25]. When there was no background variation in color, excellent detection accuracy was obtained for density defined targets (*i.e.,* denser and sparser targets) at 150 ms (94%). Height defined targets (*i.e.,* taller and shorter) were detected somewhat less accurately at 150 ms (88%) but were highly detectable at 450 ms (93%). As we had also found previously, taller targets were generally easier to detect than shorter targets, and denser targets were easier than sparser targets.
3. In all four texture experiments, background variation in color produced a small but significant interference effect, averaging 6% in overall accuracy reduction.

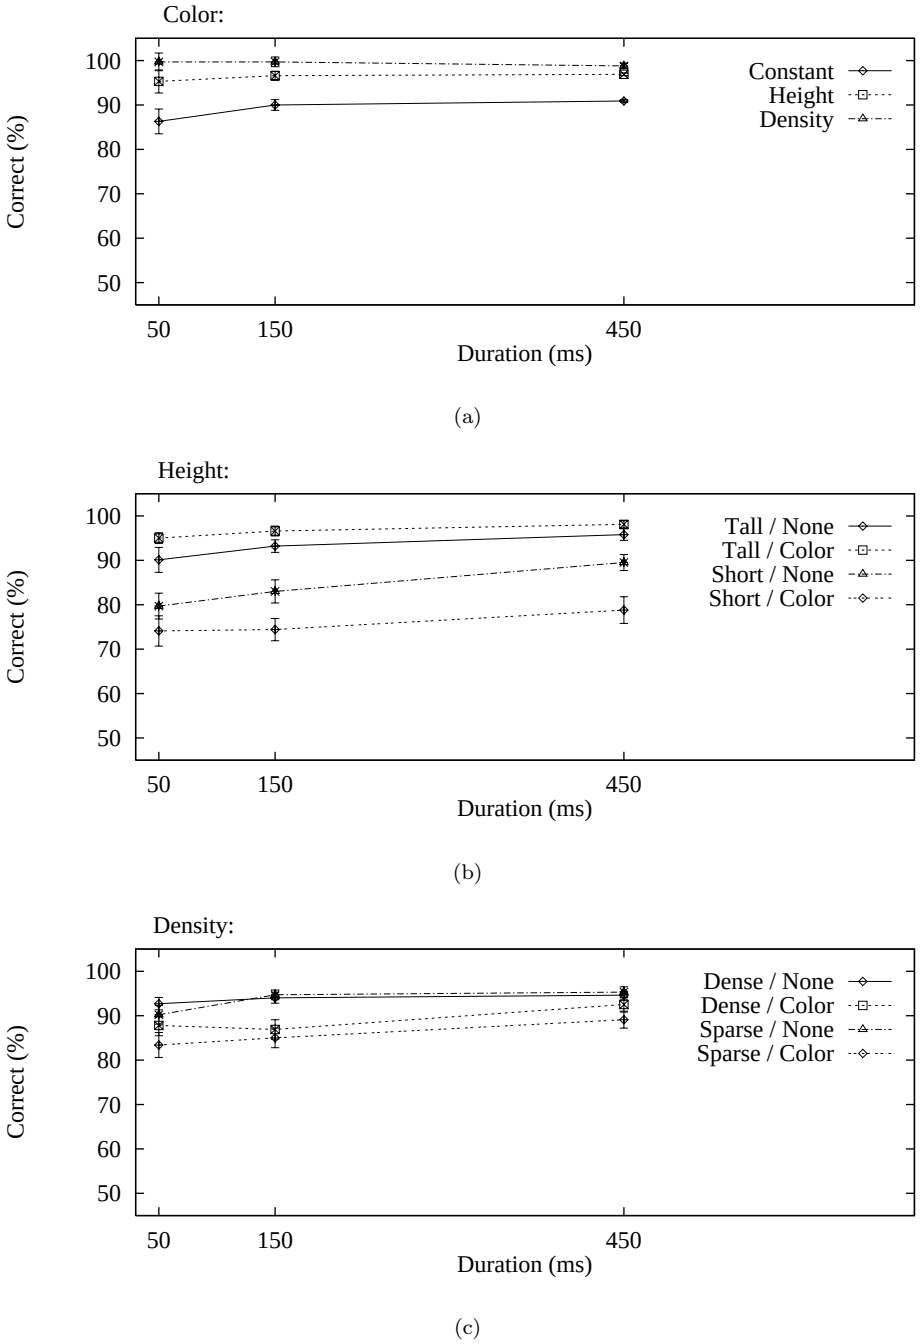

Fig. 13. Graphs showing averaged subject results for color, height, and density trials: (a) results for color trials, horizontal axis plots exposure duration, vertical axis plots percentage of correct responses, each line corresponds to one of the three different background conditions (no variation, height variation, or density variation); (b) results for height trials; (c) results for density trials

4. The absolute reduction in accuracy due to color interference depended on the difficulty of the detection task. Specifically, color interfered more with the less visible target values (shorter and sparser targets yielded a mean accuracy reduction of 8%) than with the more visible targets (taller and denser targets yield a mean accuracy reduction of 4%).

Fig. 12 shows average subject performance as a table representing each target-background pair. Target type varies along the horizontal axis, while background type varies along the vertical axis. Darker squares represent target-background pairings with highly accurate subject performance. The number in the center of each square reports the percentage of correct responses averaged across all subjects.

Target regions defined by a particular pexel color were identified rapidly and accurately in all cases. At a 150 ms exposure duration mean accuracy was approximately 96%. The small increase in accuracy from shorter to longer exposure durations was significant, $F(2, 36) = 41.03, p < .001$.

However, variation in the background height or density of pexels caused no significant reduction in performance (mean accuracy of 95.3% for constant background, 96.6% for varying height, and 96.9% for varying density; see also the graph in Fig. 13a). In fact, the graphs in Figure 13a report that absolute performance was slightly better for conditions with background variations of height or density. We suspect that geometric regularity in the texture pattern may produce a gestalt or configurational effect that interferes with target detection based on color. If so, this would be similar to previous reports in the psychophysical literature [2] showing inhibitory effects of gestalt grouping on target detection.

Detection accuracy for targets defined by texture properties were very similar to results from our previous texture experiments [22], [23]. Both kinds of targets benefited from longer exposure durations (density, $F(2,58) = 9.24, p < .001$; height, $F(2,58) = 10.66, p < .001$), with small but significant increases in accuracy with each increase in duration. With regard to the four kinds of targets, denser and taller target regions were easiest to identify, followed by sparser and shorter target regions (Fig. 13b and 13c). However, only the difference between taller versus shorter targets was statistically significant, $F(1,29) = 67.14, p < .001$. These effects were not unexpected, since they have been reported in other psychophysical studies [1], [51]. In the target present displays, accuracy for shorter targets seemed to be compromised even more than usual because of occlusion: a group of shorter pexels was often partially occluded by a group of taller pexels placed in front of them. A group of taller pexels, on the other hand, tended to stand out among the shorter pexels that surrounded them. Sparser targets suffer from a different problem: the need for a minimum amount of physical space to become perceptually salient. Since dense targets "add information to" their target region, rather than "take information away", they were less susceptible to this problem. This asymmetry contributed to a significant target type by region size interaction, $F(1,29) = 11.14, p < .01$. This was reflected in a dramatic reduction in the performance gap between dense and sparse targets when 2×2 and 4×4 target patches are compared. In displays with 2×2 target regions and background color variation, dense targets outperform sparse targets by approximately 7%. For 4×4 target regions, however, dense and sparse displays were nearly equal in accuracy (less than 1% difference).

For targets defined by texture, random color variation tended to interfere with detection, causing accuracy to be lower for both denser and sparser targets in the density displays ($F(1,29) = 9.12, p < .01$) and by interacting with target type in the case of height ($F(1,29) = 10.61, p < .01$, see also Fig. 13b and 13c). This interaction resulted from color variation having a greater influence on accuracy for short targets ($F(1,15) = 6.73, p < .03$), which were generally more difficult to see, than for tall targets, which were detected with uniformly high accuracy (greater than 90%). These results suggest that color interference can be limited when color and texture are combined, but only in cases where the detection task is relatively effortless prior to the addition of color variation. As can be seen in Fig. 13b and 13c, the interference effect of color variation tends to be greatest when the target detection task is most difficult.

Several other miscellaneous effects were worthy of note. Detection accuracy was generally higher on displays with a target present than when no target was present (color, $F(1,18) = 37.32, p < .001$; density, $F(1,29) = 5.09, p < .04$; height, $F(1,29) = 6.64, p < .02$). This was a small difference overall (an average of 4%) but it reflected a slight bias on the part of users to guess "target present" when they were uncertain what they had seen. Large target regions (4×4) were generally easier to identify than small regions (2×2) (color, $F(1,18) = 15.38, p < .001$; density, $F(1,29) = 94.24, p < .001$; height, $F(1,29) = 24.78, p < .001$), due to the greater visibility associated with a larger target region.

Taken together, these results are consistent with studies based on textures arrayed in a two-dimensional plane and reported in the psychophysical literature. As described by [49], we found that color produces a small but statistically reliable interference effect during texture segmentation. Moreover, we found color and texture form a "feature hierarchy" that produces asymmetric interference: color variation interferes with an observer's ability to see texture regions based on height or density, but variation in texture has no effect on region detection based on color. This is similar to reports by [8], [9], who reported asymmetric color on shape interference in a boundary detection task involving two-dimensional textures.

VII. Practical Applications

Although our theoretical results provide a solid design foundation, it is equally important to ensure that these results can be applied to real-world data. Our initial goal was a technique for visualizing multivariate data on an underlying height field. We decided to test our perceptual visualization technique by analyzing environmental conditions on a topographic map. Specifically, we visualized typhoons in the Northwest Pacific Ocean during the summer and fall of 1997.

A. Visualizing Typhoons

The names "typhoon" and "hurricane" are region-specific, and refer to the same type of weather phenomena: an atmospheric disturbance characterized by low pressure, thunderstorm activity, and a cyclic wind pattern. Storms of this type with windspeeds below 17m/s are called "tropical depressions". When windspeeds exceed 17m/s, they become "tropical storms". This is also when storms are assigned a specific name. When windspeeds reach 33m/s, a storm becomes a typhoon (in the Northwest Pacific) or a hurricane (in the Northeast Pacific and North Atlantic).

We combined information from a number of different sources to collect the data that we needed. A U.S. Navy elevation dataset[1] was used to obtain land elevations at

[1] http://grid2.cr.usgs.gov/dem/

ten minute latitude and longitude intervals. Land-based weather station readings collected from around the world and archived by the National Climatic Data Center[2] provided daily measurements for eighteen separate environmental conditions. Finally, satellite archives made available by the Global Hydrology and Climate Center[3] contained daily open-ocean windspeed measurements at thirty minute latitude and longitude intervals. The National Climatic Data Center defined the 1997 typhoon season to run from August 1 to October 31. Each of our datasets contained measurements for this time period.

We chose to visualize three environmental conditions related to typhoons: windspeed, pressure, and precipitation. All three values were measured on a daily basis at each land-based weather station, but only daily windspeeds were available for open-ocean positions. In spite of the missing open-ocean pressure and precipitation, we were able to track storms as they moved across the Northwest Pacific Ocean. When the storms made landfall the associated windspeed, sea-level pressure, and precipitation were provided by weather stations along their path.

Based on our experimental results, we chose to represent windspeed, pressure, and precipitation with height, density, and color, respectively. Localized areas of high windspeed are obvious indicators of storm activity. We chose to map increasing windspeed to an increased pexel height. Although our experimental results showed statistically significant interference from background color variation, the absolute effect was very small. Taller and denser pexels were easily identified in all other cases, suggesting there should be no changes in color interference due to an increase in task difficulty. Windspeed has two important boundaries: 17m/s (where tropical depressions become tropical storms) and 33m/s (where storms become typhoons). We mirrored these boundaries with height discontinuities. Pexel height increases linearly from 0-17m/s. At 17m/s, height approximately doubles, then continues linearly from 17-33m/s. At 33m/s another height discontinuity is introduced, followed by a linear increase for any windspeeds over 33m/s.

Pressure is represented with pexel density. Since our results showed it was easier to find dense pexels in a sea of sparse pexels (as opposed to sparse in dense), an increase in pressure is mapped to a decrease in pexel density (*i.e.,* very dense pexels represent the low pressure regions associated with typhoons). Three different texture densities were used to represent three pressure ranges. Pressure readings less than 996 millibars, between 996 and 1014 millibars, and greater than 1014 millibars produce very dense, dense, and sparse pexels, respectively.

Precipitation is represented with color. We used our perceptual color selection technique to choose five perceptually uniform colors. Daily precipitation readings of zero, 0–0.03 inches, 0.03–0.4 inches, 0.4–1.0 inches, and 1.0–10.71 inches were colored green, yellow, orange, red, and purple, respectively (each precipitation range had an equal number of entries in our typhoon dataset). Pexels on the open ocean or at weather stations where no precipitation values were reported were colored blue-green. Our experimental results showed no texture-on-color interference. Moreover, our color selection technique is designed to produce colors that are equally distinguishable from one another. Our mapping uses red and purple to highlight the high-precipitation areas associated with typhoon activity.

We should note that our data-feature mapping is designed to allow viewers to rapidly and accurately identify and track the locations of storms and typhoons as spatial collections of tall, dense, red and purple pexels. Our visualization system is not meant to allow users to determine exact values of windspeed, pressure, and precipitation from an individual pexel. However, knowing the range of values that produce certain types of height, density, and color will allow a viewer to estimate the environmental conditions at a given spatial location.

We built a simple visualization tool that maps windspeed, pressure, and precipitation to their corresponding height, density, and color. Our visualization tool allows a user to move forwards and backwards through the dataset day-by-day. One interesting result was immediately evident when we began our analysis: typhoon activity was not represented by high windspeed values in our open-ocean dataset. Typhoons normally contain severe rain and thunderstorms. The high levels of cloud-based water vapor produced by these storms block the satellites that are used to measure open-ocean windspeeds. The result is an absence of any windspeed values within a typhoon's spatial extent. Rather than appearing as a local region of high windspeeds, typhoons on the open-ocean are displayed as a "hole", an ocean region without any windspeed readings (see Fig. 14b and 14d). This absence of a visual feature (*i.e.,* a hole in the texture field) is large enough to be salient in our displays, and can be preattentively identified and tracked over time. Therefore, users have little difficulty finding storms and watching them as they move across the open ocean. When a storm makes landfall, the weather stations along the storm's path provide the proper windspeed, as well as pressure and precipitation. Weather stations measure windspeed directly, rather than using satellite images, so high levels of cloud-based water vapor cause no loss of information.

Fig. 14 shows windspeed, pressure, and precipitation around Japan, Korea, and Taiwan during August 1997. Fig. 14a looks north, and displays normal summer conditions across Japan on August 7, 1997. Fig. 14b, looking northeast, tracks typhoon Amber (one of the region's major typhoons) approaching along an east to west path across the Northwest Pacific Ocean on August 27, 1997. Fig. 14c shows typhoon Amber one day later as it moves through Taiwan. Weather stations within the typhoon show the expected strong winds, low pressure, and high levels of rainfall. These results are easily identified as tall, dense, red and purple pexels. Compare these images to Fig. 14d and 14e, where windspeed was mapped to regularity, pressure to height, and precipitation to density (a mapping without color that our original texture experiments predict will

[2]http://www.ncdc.noaa.gov/ol/climate/online/gsod.html
[3]http://ghrc.msfc.nasa.gov/ghrc/list.html

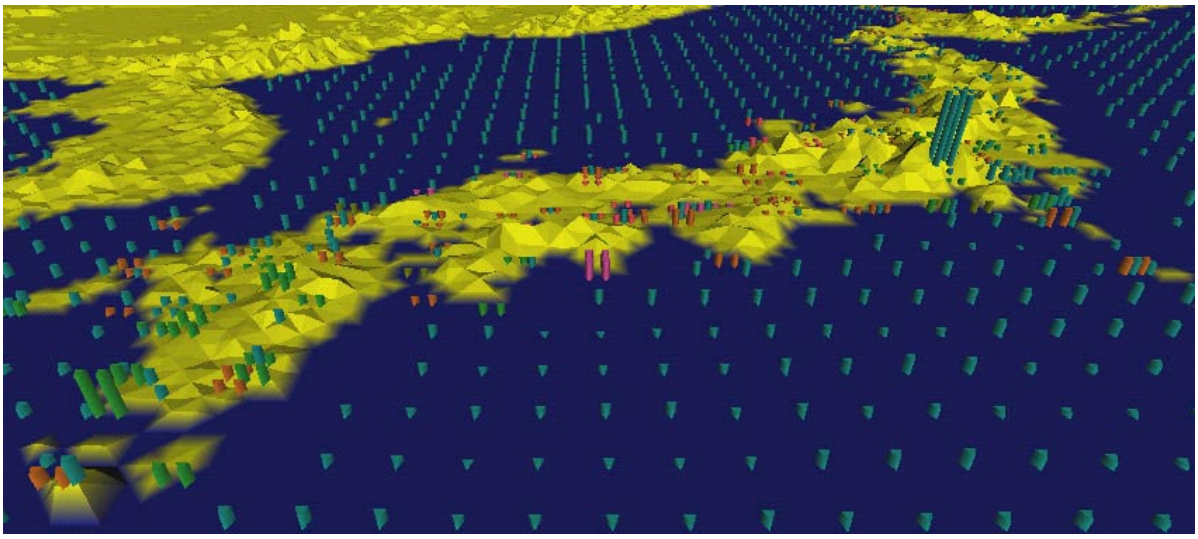

(a)

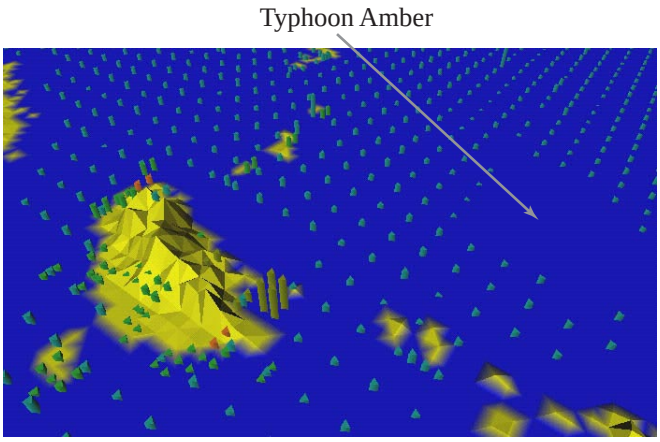

(b)

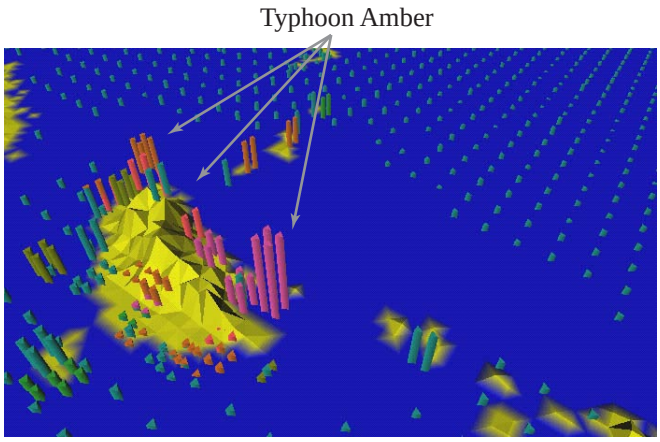

(c)

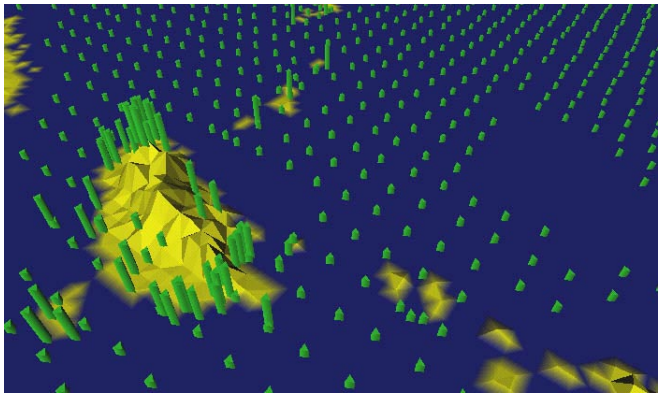

(d)

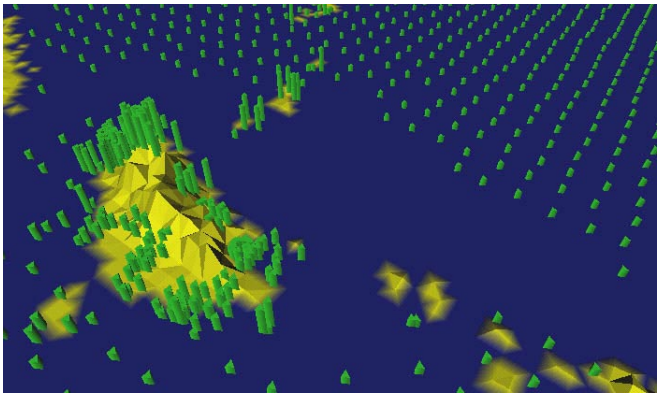

(e)

Fig. 14. Typhoon conditions across Southeast Asia during the summer of 1997: (a) August 7, 1997, normal weather conditions over Japan; (b) August 27, 1997, typhoon Amber approaches the island of Taiwan from the southeast; (c) August 28, 1997, typhoon Amber strikes Taiwan, producing tall, dense pexels colored orange, red, and purple (representing high precipitation); (d, e) the same data as in (b) and (c) but with windspeed represented by regularity, pressure by height, and precipitation by density

perform poorly). Although viewers can identify areas of lower and higher windspeed (*e.g.,* on the open ocean and over Taiwan), it is difficult to identify *a change* in lower or higher windspeeds (e.g., the change in windspeed as typhoon Amber moves onshore over Taiwan). In fact, viewers often searched for an increase in density that represents an increase in precipitation, rather than an increase in irregularity; pexels over Taiwan become noticeably denser between Fig. 14d and 14e.

VIII. Conclusions and Future Work

This paper describes a method for combining perceptual textures and colors for multivariate data visualization. Our pexels are built by varying three perceptual texture dimensions: height, density, and regularity. Our perceptual colors are selected by controlling the color distance, linear separation, and color category of each color. Both experimental and real-world results showed that colored pexels can be used to rapidly, accurately, and effortlessly analyze large, multi-element displays. Care must be taken, however, to ensure that the data-feature mapping builds upon the fundamental workings of the low-level human visual system. An ad-hoc mapping will often introduce visual artifacts that actively interfere with a user's ability to perform their visual analysis task. Our initial texture experiments showed that taller, shorter, denser, and sparser pexels can be easily identified, but that certain background patterns must be avoided to ensure accurate performance. During our color selection experiments we found that color distance, linear separation, and color category must all be considered to ensure a collection of equally distinguishable colors. New results on the combined use of texture and color showed that background color variation causes a small but statistically significant interference effect during a search for targets based on height or density. The size of the effect is directly related to the difficulty of the visual analysis task; tasks that are more difficult result in more color interference. Variation of height and density, on the other hand, had no effect on identifying color targets. These results are similar to reports in the psychophysical literature [8], [9], [49], although to our knowledge no one has studied perceptual textures and colors displayed in 3D using perspective projections.

Our results were further validated when we applied them to real-world applications like typhoon visualization. Our tools were designed to satisfy findings from our experiments. For example, attributes were mapped in order of importance to height, density, and color. In cases where an subject analyzed height or density patterns, we tried to ensure an effortless search task (*i.e.,* looking for taller or denser rather than shorter or sparser) to minimize any color on texture interference that might occur.

One important area of future work is a comparison of our visualization techniques against other methods that might be used to represent information in our real-world applications. For example, it would be useful to test a user's ability to track storm activity in our visualization environment against other standard techniques for representing weather activity. Although we have yet to conducted these kinds of practical experiments, we hope to initiate them in the near future as part of our perceptual visualization studies.

We are now working to integrate our colored pexels with other visual features. One candidate is orientation; in fact, our pexels were initially designed to "stand up" off the underlying height field to support variation of orientation. Another visual property with significant potential is apparent motion. This technique can be used to make individual strips in a pexel "walk" within their spatial extent. It may be possible to tie direction and speed of motion to two underlying attribute values, thereby increasing the dimensionality of our visualization techniques. We are designing experiments to investigate the effectiveness of each of these features for encoding information. We will also study any interactions that occur when multiple texture, color, orientation, and motion dimensions are displayed simultaneously.

Acknowledgments

We would like to thank the National Climatic Data Center, and Sherry Harrison and the Global Hydrology Resource Center for generously providing typhoon-related weather data. We would also like to thank Jeanette Lum for coordinating and running our experiment sessions. This research was funded in part by the National Science and Engineering Research Council of Canada, and by the Office of Naval Research (Grant N00014-96-1120) and the Ballistic Missile Defense Organization through the Multiuniversity Research Initiative.

References

[1] Aks, D. J., and Enns, J. T. Visual search for size is influenced by a background texture gradient. *Journal of Experimental Psychology: Perception and Performance 22*, 6 (1996), 1467–1481.

[2] Banks, W. P., and Prinzmetal, W. Configurational effects in visual information processing. *Perception & Psychophysics 19* (1976), 361–367.

[3] Bauer, B., Jolicoeur, P., and Cowan, W. B. Visual search for colour targets that are or are not linearly-separable from distractors. *Vision Research 36* (1996), 1439–1446.

[4] Bauer, B., Jolicoeur, P., and Cowan, W. B. The linear separability effect in color visual search: Ruling out the additive color hypothesis. *Perception & Psychophysics 60*, 6 (1998), 1083–1093.

[5] Bergman, L. D., Rogowitz, B. E., and Treinish, L. A. A rule-based tool for assisting colormap selection. In *Proceedings Visualization '95* (Atlanta, Georgia, 1995), pp. 118–125.

[6] Bruckner, L. A. On Chernoff faces. In *Graphical Representation of Multivariate Data*, P. C. C. Wang, Ed. Academic Press, New York, New York, 1978, pp. 93–121.

[7] Callaghan, T. C. Dimensional interaction of hue and brightness in preattentive field segregation. *Perception & Psychophysics 36*, 1 (1984), 25–34.

[8] Callaghan, T. C. Interference and domination in texture segregation: Hue, geometric form, and line orientation. *Perception & Psychophysics 46*, 4 (1989), 299–311.

[9] Callaghan, T. C. Interference and dominance in texture segregation. In *Visual Search*, D. Brogan, Ed. Taylor & Francis, New York, New York, 1990, pp. 81–87.

[10] Chernoff, H. The use of faces to represent points in k-dimensional space graphically. *Journal of the American Statistical Association 68*, 342 (1973), 361–367.

[11] CIE. *CIE Publication No. 15, Supplement Number 2 (E-1.3.1): Official Recommendations on Uniform Color Spaces, Color-Difference Equations, and Metric Color Terms.* Commission Internationale de L'Èclairge, 1976.

[12] Coren, S., and Hakstian, A. R. Color vision screening without the use of technical equipment: Scale development and cross-validation. *Perception & Psychophysics 43* (1988), 115–120.

[13] Cutting, J. E., and Millard, R. T. Three gradients and the perception of flat and curved surfaces. *Journal of Experimental Psychology: General 113*, 2 (1984), 198–216.

[14] D'Zmura, M. Color in visual search. *Vision Research 31*, 6 (1991), 951–966.

[15] Egeth, H. E., and Yantis, S. Visual attention: Control, representation, and time course. *Annual Review of Psychology 48* (1997), 269–297.

[16] Foley, J., and Ribarsky, W. Next-generation data visualization tools. In *Scientific Visualization: Advances and Challenges*, L. Rosenblum, Ed. Academic Press, San Diego, California, 1994, pp. 103–127.

[17] Foley, J. D., van Dam, A., Feiner, S. K., and Hughes, J. F. *Computer Graphics: Principles and Practice*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1990.

[18] Glassner, A. S. *Principles of Digital Image Synthesis*. Morgan Kaufmann Publishers, Inc., San Francisco, California, 1995.

[19] Grinstein, G., Pickett, R., and Williams, M. EXVIS: An exploratory data visualization environment. In *Proceedings Graphics Interface '89* (London, Canada, 1989), pp. 254–261.

[20] Hallett, P. E. Segregation of mesh-derived textures evaluated by resistance to added disorder. *Vision Research 32*, 10 (1992), 1899–1911.

[21] Haralick, R. M., Shanmugam, K., and Dinstein, I. Textural features for image classification. *IEEE Transactions on System, Man, and Cybernetics SMC-3*, 6 (1973), 610–621.

[22] Healey, C. G. Choosing effective colours for data visualization. In *Proceedings Visualization '96* (San Francisco, California, 1996), pp. 263–270.

[23] Healey, C. G. Building a perceptual visualisation architecture. *Behaviour and Information Technology (in press)* (1998).

[24] Healey, C. G., Booth, K. S., and Enns, J. T. Real-time multivariate data visualization using preattentive processing. *ACM Transactions on Modeling and Computer Simulation 5*, 3 (1995), 190–221.

[25] Healey, C. G., and Enns, J. T. Building perceptual textures to visualize multidimensional datasets. In *Proceedings Visualization '98* (Research Triangle Park, North Carolina, 1998), pp. 111–118.

[26] Interrante, V. Illustrating surface shape in volume data via principle direction-driven 3d line integral convolution. In *SIGGRAPH 97 Conference Proceedings* (Los Angeles, California, 1997), T. Whitted, Ed., pp. 109–116.

[27] Julész, B. Textons, the elements of texture perception, and their interactions. *Nature 290* (1981), 91–97.

[28] Julész, B. A theory of preattentive texture discrimination based on first-order statistics of textons. *Biological Cybernetics 41* (1981), 131–138.

[29] Julész, B. A brief outline of the texton theory of human vision. *Trends in Neuroscience 7*, 2 (1984), 41–45.

[30] Julész, B., and Bergen, J. R. Textons, the fundamental elements in preattentive vision and perception of textures. *The Bell System Technical Journal 62*, 6 (1983), 1619–1645.

[31] Kawai, M., Uchikawa, K., and Ujike, H. Influence of color category on visual search. In *Annual Meeting of the Association for Research in Vision and Ophthalmology* (Fort Lauderdale, Florida, 1995), p. #2991.

[32] Laidlaw, D. H., Ahrens, E. T., Kremers, D., Avalos, M. J., Jacobs, R. E., and Readhead, C. Visualizing diffusion tensor images of the mouse spinal cord. In *Proceedings Visualization '98* (Research Triangle Park, North Carolina, 1998), pp. 127–134.

[33] Levkowitz, H. Color icons: Merging color and texture perception for integrated visualization of multiple parameters. In *Proceedings Visualization '91* (San Diego, California, 1991), pp. 164–170.

[34] Levkowitz, H., and Herman, G. T. Color scales for image data. *IEEE Computer Graphics & Applications 12*, 1 (1992), 72–80.

[35] Liu, F., and Picard, R. W. Periodicity, directionality, and randomness: Wold features for perceptual pattern recognition. In *Proceedings 12th International Conference on Pattern Recognition* (Jerusalem, Israel, 1994), pp. 1–5.

[36] Mack, A., and Rock, I. *Inattentional Blindness*. MIT Press, Menlo Park, California, 1998.

[37] Malik, J., and Perona, P. Preattentive texture discrimination with early vision mechanisms. *Journal of the Optical Society of America A 7*, 5 (1990), 923–932.

[38] Meier, B. J. Painterly rendering for animation. In *SIGGRAPH 96 Conference Proceedings* (New Orleans, Louisiana, 1996), H. Rushmeier, Ed., pp. 477–484.

[39] Rao, A. R., and Lohse, G. L. Identifying high level features of texture perception. *CVGIP: Graphics Models and Image Processing 55*, 3 (1993), 218–233.

[40] Rao, A. R., and Lohse, G. L. Towards a texture naming system: Identifying relevant dimensions of texture. In *Proceedings Visualization '93* (San Jose, California, 1993), pp. 220–227.

[41] Reed, T. R., and Hans Du Buf, J. M. A review of recent texture segmentation and feature extraction techniques. *CVGIP: Image Understanding 57*, 3 (1993), 359–372.

[42] Rensink, R. A., O'Regan, J. K., and Clark, J. J. To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science 8* (1997), 368–373.

[43] Rheingans, P., and Tebbs, B. A tool for dynamic explorations of color mappings. *Computer Graphics 24*, 2 (1990), 145–146.

[44] Robertson, P. K. Visualizing color gamuts: A user interface for the effective use of perceptual color spaces in data displays. *IEEE Computer Graphics & Applications 8*, 5 (1988), 50–64.

[45] Rogowitz, B. E., and Treinish, L. A. An architecture for rule-based visualization. In *Proceedings Visualization '93* (San Jose, California, 1993), pp. 236–243.

[46] Salisbury, M., Wong, M. T., Hughes, J. F., and Salesin, D. H. Orientable textures for image-based pen-and-ink illustration. In *SIGGRAPH 97 Conference Proceedings* (Los Angeles, California, 1997), T. Whitted, Ed., pp. 401–406.

[47] Schweitzer, D. Artificial texturing: An aid to surface visualization. *Computer Graphics (SIGGRAPH 83 Conference Proceedings) 17*, 3 (1983), 23–29.

[48] Simon, D. J., and Levin, D. T. Change blindness. *Trends in Cognitive Science 1* (1997), 261–267.

[49] Snowden, R. J. Texture segregation and visual search: A comparison of the effects of random variations along irrelevant dimensions. *Journal of Experimental Psychology: Human Perception and Performance 24*, 5 (1998), 1354–1367.

[50] Tamura, H., Mori, S., and Yamawaki, T. Textural features corresponding to visual perception. *IEEE Transactions on Systems, Man, and Cybernetics SMC-8*, 6 (1978), 460–473.

[51] Triesman, A. Preattentive processing in vision. *Computer Vision, Graphics and Image Processing 31* (1985), 156–177.

[52] Triesman, A. Search, similarity, and integration of features between and within dimensions. *Journal of Experimental Psychology: Human Perception & Performance 17*, 3 (1991), 652–676.

[53] Triesman, A., and Gormican, S. Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review 95*, 1 (1988), 15–48.

[54] Turk, G., and Banks, D. Image-guided streamline placement. In *SIGGRAPH 96 Conference Proceedings* (New Orleans, Louisiana, 1996), H. Rushmeier, Ed., pp. 453–460.

[55] Ware, C. Color sequences for univariate maps: Theory, experiments, and principles. *IEEE Computer Graphics & Applications 8*, 5 (1988), 41–49.

[56] Ware, C., and Beatty, J. C. Using colour dimensions to display data dimensions. *Human Factors 30*, 2 (1988), 127–142.

[57] Ware, C., and Knight, W. Using visual texture for information display. *ACM Transactions on Graphics 14*, 1 (1995), 3–20.

[58] Wolfe, J. M. Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review 1*, 2 (1994), 202–238.

[59] Wolfe, J. M., Yu, K. P., Stewart, M. I., Shorter, A. D., Friedman-Hill, S. R., and Cave, K. R. Limitations on the parallel guidance of visual search: Color × color and orientation × orientation conjunctions. *Journal of Experimental Psychology: Human Perception & Performance 16*, 4 (1990), 879–892.

[60] Wyszecki, G., and Stiles, W. S. *Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd Edition*. John Wiley & Sons, Inc., New York, New York, 1982.

Christopher G. Healey received a PhD in computer science in 1996 from the University of British Columbia, an MSc in 1992 from the University of British Columbia, and a BMath in 1990 from the University of Waterloo. Following graduation he completed a two-year postdoctoral fellowship in computer graphics with Dr. Carlo Séquin at the University of California, Berkeley. He is currently working as an assistant professor in the Department of Computer Science at North Carolina State University. His dissertation studied methods for displaying effectively large, multivariate datasets during scientific visualization. This work investigated techniques for exploiting the low-level human visual system for information representation. His current research focuses on the use of visual features like color, texture, and apparent motion for visually exploring multivariate data. He is also investigating automated data-feature mapping techniques and data management issues in an effort to design a flexible, robust perceptual visualization architecture.

James T. Enns received a PhD in psychology from Princeton University (1984) and a BA (honours) from the University of Winnipeg (1980). Following graduation, he was appointed an assistant professor at Dalhousie University, before moving to the Unviersity of British Columbia in 1987, where he is now appointed as a professor in both the Department of Psychology and the Graduate Program in Neuroscience. A central focus of his research over the years has been the role of attention in perception. This has included studies of how perception and attention change with development, how the visual world is represented outside the focus of attention, and how attention changes the perceptions that form the basis for consciousness. Along with the publication of these studies in *Science, Psychological Review, Perception & Psychophysics,* and *The Journal of Experimental Psychology,* he has edited two research volumes (*The Development of Attention,* 1990; *Attention, Development, & Psychopathology,* 1997) and coauthored two textbooks (*Analysis of Variance,* 1986; *Sensation & Perception*, fifth edition, 1999).