

Combining Perception and Impressionist Techniques for Nonphotorealistic Visualization of Multidimensional Data

Christopher G. Healey
North Carolina State University

1 Introduction

An important research problem in computer graphics is the visualization of multidimensional data, the conversion of a dataset D containing strings and numbers into a sequence of one or more images. Values in D represent m attributes $A = \{A_1, \dots, A_m\}$ recorded at n sample points e_i , that is, $D = \{e_1, \dots, e_n\}$ and $e_i = \{a_{i,1}, \dots, a_{i,m}\}$, $a_{i,j} \in A_j$. A data-feature mapping $M(V, \Phi)$ defines a visual feature $V_j \in V$ to use to display values from A_j ; it also defines $\phi_j : A_j \rightarrow V_j$, $\phi_j \in \Phi$ to map the domain of A_j to the range of displayable values in V_j . Visualization in this framework is the construction of a data-feature mapping M together with a viewer's interpretation of the images produced by M .

Although the need to visualize multiple layers of information simultaneously is well documented [43, 61, 63], progress towards this goal has been slow [56]. It has often proven difficult to construct methods that represent multidimensional data in a way is easy to *explore*, *analyze*, *verify* and *discover*. The desire to build fundamental techniques that are appropriate for a wide range of visualization environments further complicates this problem.

Previous work has studied methods for harnessing the low-level human visual system during visualization [4, 26, 52, 81]. Certain visual features (*e.g.*, hue, luminance, contrast, and motion) are detected very quickly by the visual system [14, 70, 83]; when combined properly, these same features can be used to construct multidimensional displays that can be *rapidly*, *accurately*, and *effortlessly* explored and analyzed by a viewer. The application of perception in aid of visualization has shown great promise, and has been explicitly cited as an important area of current and future research [63].

We have recently initiated a study of the use of artistic techniques for multidimensional visualization. This investigation was motivated in large part by work on nonphotorealistic rendering in computer graphics [10, 18, 27, 28, 39, 44, 65], and by the efforts of researchers like Interrante [30], Laidlaw [36, 37], and Ebert and Rheingans [13] to extend this work to a visualization environment. Certain movements and techniques in painting (*e.g.*, impressionism, expressionism, or watercolor) are characterized by a set of fundamental styles. If these styles can be identified and simulated on a computer, we believe they can then be applied to represent individual data attributes in a multidimensional dataset. Consider, for example, a dataset containing weather conditions. A "painting" made up of simulated brush strokes could be used to visualize this data. A brush stroke's color would represent temperature at a given spatial location; stroke direction and length would represent wind direction and strength; stroke density would represent pressure. The result is an image that looks like a painting, not of a real-world scene, but rather of the information contained in the underlying dataset.

Such a technique might initially seem difficult to control and test. An important insight is that many painterly styles correspond closely to perceptual features that are detected by the human visual system. In some sense this is not surprising. Artistic masters understood intuitively which properties of a painting would capture a viewer's gaze, and their styles naturally focused on harnessing these features. Moreover, certain movements used scientific studies of the visual system to help them understand how viewers would perceive their work (*e.g.*, the use of the perceptual color models of Chevreul [9] and Rood [55] in Impressionism). The overlap of artistic styles and perception offers two significant advantages. Most importantly, the body of knowledge on the use of perception during visualization can help us to predict how corresponding painterly styles might perform in the same environment. In addition, psychophysical experiments offer a method for designing controlled studies that can test the fundamental strengths and limitations of a given style, both in isolation and in combination with other styles being shown simultaneously in the same display.

<i>Feature</i>	<i>Author</i>
line (blob) orientation	Julész & Bergen (1983); Wolfe (1992)
length	Triesman & Gormican (1988)
width	Julész (1984)
size	Triesman & Gelade (1980)
curvature	Triesman & Gormican (1988)
number	Julész (1985); Trick & Pylyshyn (1994); Healey, Booth, & Enns (1996)
terminators	Julész & Bergen (1983)
intersection	Julész & Bergen (1983)
closure	Enns (1986); Triesman & Souther (1986)
color (hue)	Triesman & Gormican (1988); Nagy & Sanchez (1990); D’Zmura (1991); Healey (1997)
intensity	Beck et al. (1983); Triesman & Gormican (1988)
flicker	Julész (1971)
direction of motion	Nakayama & Silverman (1986); Driver & McLeod (1992)
binocular lustre	Wolfe & Franzel (1988)
stereoscopic depth	Nakayama & Silverman (1986)
3-D depth cues	Enns (1990)
lighting direction	Enns (1990)
texture	Healey & Enns (1998)

Table 1: A list of two-dimensional features that “pop out” during visual search, and a list of authors who describe preattentive tasks performed using the given feature.


2 Low-Level Human Vision

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 displayed¹. 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. Table 1 lists a number of preattentive features, and provides references that describe the tasks that can be performed using these features.

Fig. 1 shows examples of both types of target search. In Fig. 1a-1d 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. 1e and 1f 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


¹Although we now know that these visual features are influenced by the goals and expectations of the observer, the term preattentive is still useful because it conveys the relative ease with which these processes are completed.




(a)




(b)




(c)



(d)



(e)



(f)

Figure 1: 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)

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. 1a-1d, 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 [14]). 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 [41], there are a multitude of other simultaneous salient visual events [51], or because the salient event occurs during an eye movement or other change in viewpoint [62]. Therefore, the control of attention must always be understood as an interaction between bottom-up and top-down mechanisms.

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 [23] (*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.

Properties that are processed preattentively can be used to highlight important image characteristics. Experiments in both the cognitive psychology and scientific visualization domains have used various features to assist in performing the following visual tasks:

- *target detection*, where users attempt to rapidly and accurately detect the presence or absence of a “target” element that uses a unique visual feature within a field of distracter elements (Fig. 1),
- *boundary detection*, where users attempt to rapidly and accurately detect a texture boundary between two groups of elements, where all the elements in each group have a common visual feature (Fig. 2), and
- *counting and estimation*, where users attempt to count or estimate the number or percentage of elements in a display that have a unique visual feature.

Callaghan [7, 8] first reported the interference effects shown in Fig. 2. The visual system seems to prioritize features in order of importance. This means that the presence of visually “important” features can interfere with tasks that use lower priority features. In Fig. 2a, the vertical boundary defined by hue is detected preattentively, even though the shape of each element is random. In Fig. 2b, however, it is difficult to detect the horizontal boundary defined by form due to random hue variations. If hue were fixed to a constant value for each element, the form boundary could be detected preattentively. Callaghan explains this phenomena by suggesting that the visual system assigns a higher importance to hue than to form during boundary detection. Thus, a random hue interferes with form boundary detection, but a random form has no effect on hue boundary detection. A similar asymmetry exists between hue and intensity. Random hue has no effect on detecting boundaries defined by intensity. However, random intensity interferes with hue boundary detection. Callaghan concluded that intensity is more important than hue to the low-level visual system during boundary identification [6].

Researchers continue to expand preattentive processing in a number of exciting directions. To date, most of the features used in preattentive tasks have been relatively simple properties (*e.g.*, hue, orientation, line length, and size). Enns and Rensink,

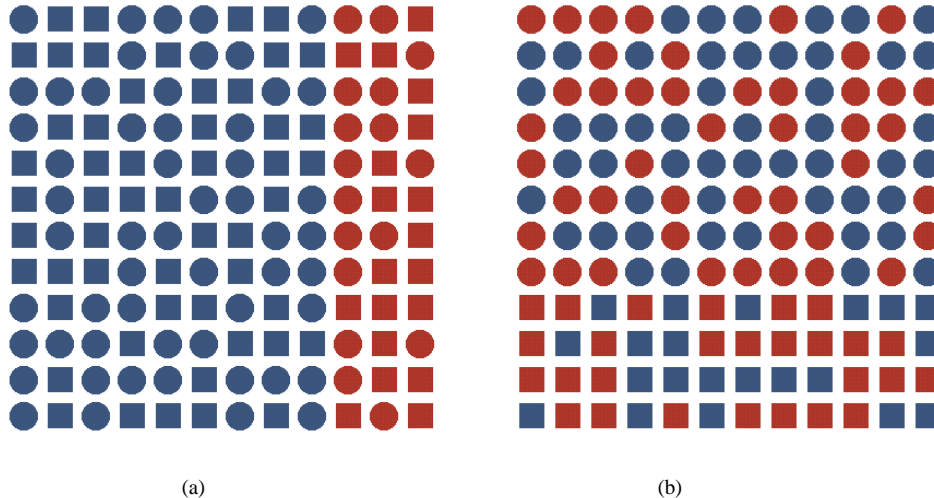


Figure 2: Region segregation by form and hue: (a) hue boundary is identified preattentively, even though form varies in the two regions; (b) random hue variations interfere with the identification of a region boundary based on form

however, have identified a class of three-dimensional elements that can be detected preattentively [15, 16]. They have shown that three-dimensional orientation, shape, and direction of lighting can be used to make elements “pop-out” of a visual scene (Fig. 3b and 3c). This is important, because it suggests that complex high-level concepts may be processed preattentively by the low-level visual system.

New tasks that can be performed preattentively are also being investigated. Research has recently been conducted on counting and estimation in preattentive processing. Varey describes experiments in which subjects were asked to estimate the relative frequency of white or black dots [76]. Her results showed that subjects could estimate in four different ways: “percentage” of white dots, “percentage” of black dots, “ratio” of black dots to white dots, and “difference” between the number of black and white dots. She also found that subjects consistently overestimated small proportions and underestimated large proportions. Estimation of relative frequency using hue and orientation was shown to be preattentive in experiments conducted in our laboratory [22, 24]. Moreover, our results showed that there was no feature interaction. Random orientation did not interfere with estimation of targets with a unique hue, and random hue did not interfere with estimation of targets with a unique orientation. This is important because it suggests that hue and orientation can be used to encode two independent data values in a single display without causing visual interference.

A number of scientists have proposed competing theories to explain how preattentive processing occurs, in particular Triesman’s feature integration theory [70], Julé’sz’ texton theory [34], Quinlan and Humphreys’ similarity theory [48], and Wolfe’s guided search theory [83]. Our interest is in the use of visual features that have already been shown to be preattentive. Results from psychology are extended, modified, tested, and then integrated into our visualization environment.

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.

2.1 Real-Time Preattentive Visualization

Most preattentive techniques are validated by studying a single data frame in isolation. This leads to an interesting question with important relevance to visualization. If a preattentive task can be performed on a single frame in 100 milliseconds, can the same task can be performed on a real-time sequence of frames displayed at ten frames per second? We hypothesized that important aspects of preattentive processing will extend to a real-time environment. A visualization tool that uses preattentive features will allow viewers to perform visual tasks such as grouping of similar data elements (boundary detection), detection of elements with a unique characteristic (target detection), and estimation of the number of elements with a given value or range of




Figure 3: Combination of simple components to form emergent features: (a) closure, a simple closed figure is seen; (b) three-dimensionality, the figure appears to have depth; (c) volume, a solid figure is seen

values, all in real-time on temporally animated data frames. We tested this hypothesis using experiments that simulated the use of our preattentive visualization tools in a real-time environment. Analysis of the experimental results supported our hypothesis for boundary and target detection using hue and shape. Moreover, interference properties previously reported for static frames were found to apply to a dynamic environment.

Our initial experiment addressed two general questions about the preattentive features hue and shape, and their use in our visualization tools:

- *Question 1:* Is it possible for subjects to detect a data frame with a horizontal boundary within a sequence of random frames? If so, what features allow this and under what conditions?
- *Question 2:* Do Callaghan's feature hierarchy effects apply to our real-time visualization environment? Specifically, does random hue interfere with form boundary detection within a sequence of frames? Does random form interfere with hue boundary detection within a sequence of frames?

Experimental results showed accurate boundary detection can be performed using either hue or form on sequences of frames displayed at ten frames per second. Moreover, feature hierarchy effects extended to a dynamic environment, specifically, hue dominated form during boundary detection. A random hue pattern masked form boundaries, while a random form pattern had no effect on hue boundary detection.

A corresponding set of experiments were run to test target detection, with similar results. While both hue and form targets can be detected preattentively in a real-time environment (at frame rates of ten to twenty frames per second), form targets were only visible when the background hue was held constant. Hue variation masked form targets. Form variation had no effect on the detection of hue targets.

We have built a number of visualization tools that allow users to perform exploratory analysis on their datasets in real-time. Experience from using these tools confirmed that our experimental results hold for these datasets and tasks.

3 Color Selection

Color is a powerful and often-used visual feature. Previous work has addressed the issue of choosing colors 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 [79]. Robertson, Ware, Rheingans and Tebbs, and Levkowitz and Herman discuss various methods for building effective color gamuts and colormaps [38, 52, 53, 78]. Recent work at the IBM Thomas J. Watson Research Center has focused on a rule-based visualization tool that considers how a user perceives visual features like hue, luminance, height, and so on [4, 54].

If we use color to represent our data, one important question to ask is: “How can we choose effective colors that provide good differentiation between data elements during the visualization task?” We addressed this problem by trying to answer three related questions:

- How can we allow rapid and accurate identification of individual data elements through the use of color?
- What factors determine whether a “target” element’s color will make it easy to find, relative to differently colored “non-target” elements?
- How many colors can we display at once, while still allowing for rapid and accurate target identification?

None of the currently-available color selection techniques were specifically designed to investigate the rapid and accurate identification of individual data elements based on color. Also, since the color gamut and colormap work uses continuous color scales to encode information, they have not addressed the question of how many colors we can effectively display at once, while still providing good differentiation between individual data elements.

We began by using the perceptually balanced CIE LUV color model to provide control over color distance and isoluminance. We also exploited two specific results related to color target detection: linear separation [12, 2] and color category [35]. These effects are controlled to allow for the rapid and accurate identification of color targets. Target identification is a necessary first step towards performing other types of exploratory data analysis. If we can rapidly and accurately differentiate elements based on their color, we can apply our results to other important visualization techniques like detection of data boundaries, the tracking of data regions in real-time, and enumeration tasks like counting and estimation [24, 71, 76].

3.1 CIE LUV

The CIE LUV color model was proposed by the Commission Internationale de L’Èclairge (CIE) in 1976 [85]. Colors are specified using the three dimensions L^* (which encodes luminance), u^* , and v^* (which together encode chromaticity). CIE LUV provides two useful properties for controlling perceived color difference. First, colors with the same L^* are isoluminant. Second, Euclidean distance and perceived color difference (specified in ΔE^* units) can be interchanged, since the color difference between two color stimuli x and y (positioned in CIE LUV at (L_x^*, u_x^*, v_x^*) (L_y^*, u_y^*, v_y^*) , respectively) is roughly:

$$\Delta E_{xy}^* = \sqrt{(\Delta L_{xy}^*)^2 + (\Delta u_{xy}^*)^2 + (\Delta v_{xy}^*)^2} \tag{1}$$

3.2 Linear Separation

The linear separation effect has been described by both D’Zmura and Bauer et. al [2, 12]. D’Zmura was investigating how the human visual system finds a target color in a sea of background non-target colors. D’Zmura ran experiments that asked observers to determine the presence or absence of an orange target. Two groups of differently colored non-target elements were also present in each display (*e.g.*, in one experiment half the non-targets in each display were colored green and half were colored red). Results showed that when the target could be separated by a straight line from its non-targets in color space (Fig. 4, target T and non-targets A and C), the time required to determine the target’s presence or absence was constant, and independent of the total number of elements being displayed. This suggests detection occurs preattentively in the low-level

visual system. When the target was collinear with its non-targets (Fig. 4, target T and non-targets A and B), the time required to identify the target was linearly proportional to the number of elements being displayed. Observers had to search serially through each display to determine whether the target was present or absent. Bauer et. al strengthened D’Zmura’s results by showing that perceptually balanced color models cannot be used to overcome the linear separation effect. Bauer et. al also replicated their findings in three additional color regions: green, blue, and green-yellow. This suggests linear separation applies to colors from different parts of the visible color domain.




Figure 4: 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

3.3 Color Category

Kawai et. al [35] reported results that suggest that the time required to identify a color target depends in part on the named color regions occupied by the target and its non-targets. Kawai et. al tested displays that contained a uniquely color target and a constant number of uniformly colored non-targets. They divided an isoluminant, perceptually balanced color slice into named color regions. Their results showed that search times decreased dramatically whenever the non-target was moved outside the target’s color region. For example, finding a target colored T in a set of non-targets colored A was significantly more difficult than than finding T in a set of non-targets colored B (Fig. 4). Since the target–non-target distances \overline{TA} and \overline{TB} are equal, there was an expectation of perceptual balance that should have been provided by the underlying color model. This expectation was not met. Kawai et. al suggests the difference in performance is due to the fact that both T and A are located in the blue color region, but B is not.

3.4 Experiments

We ran a number of experiments to study the effects of color distance, linear separation, and color category. Subjects were shown displays that contained multiple colored squares. Each subject was asked to determine whether a target square with a particular color was present or absent in each display. The experiments were designed to test the following conditions:

- *selection criteria*, which selection criteria (color distance, linear separation, and color category) need to be considered to guarantee equally distinguishable colors,
- *simultaneous colors*, how many colors can we display at the same time, while still allowing users to rapidly and accurately determine the presence or absence of any particular color, and
- *display size*, is performance affected by the number of elements in a display.

We found that up to seven isoluminant colors can be displayed simultaneously, while still allowing for the rapid and accurate identification of any one of the seven. The time required to perform identification was independent of display size, suggesting that target detection is occurring preattentively. Our results also showed that all three selection criteria needed to be considered to guarantee consistent performance. When only some of the selection criteria were used (*e.g.* only distance and separation, or only category) the amount of time required to identify targets depended on the color of the target: some colors were very easy to identify, while other colors were very difficult. This asymmetry suggested that the colors we chose were not *equally* distinguishable, and therefore that the selection criteria being used were not sufficient to properly control perceived difference between the colors during target detection.

4 Visualizing CT Medical Scans

One practical application of our color selection technique is the use of color to highlight regions of interest during volume visualization [66]. Radiologists from the University Hospital at the University of British Columbia are studying methods for visualizing abdominal aneurisms. Traditional repair of an abdominal aneurism entails a major operation with an incision into the aneurism, evacuation of the clot contained within, placement of a synthetic graft, and wrapping of the graft with the remnants of the wall of the aneurism. Recently, a new treatment option, endovascular stenting, has been proposed and is currently undergoing clinical trials. This procedure does not require general anesthesia and can be done less invasively by simply placing a self-expanding stent graft via a catheter into the aneurism to stabilize it. Less fit patients are able to withstand the procedure, hospital stay is cut to 1 to 2 days, and post-operative recovery is shortened considerably.

After the operation computed tomography (CT) scans are used to obtain two-dimensional slices of a patient’s abdominal region. These slices are reconstructed to produce a three-dimensional volume. The volume is visualized by the radiologists to perform post-operative analysis. A two-pass segmentation step is used to strip out material in each CT slice that does not correspond to one of the regions of interest: the artery running through the abdomen, the aneurism, and the metal hooks (called tynes) used to embed the stent graft within the aneurism. The reconstructed volumes must show clearly each of these three regions.

Normally, greyscale is used to display reconstructed medical volumes. Changes in luminance are most effective for representing the high spatial frequency data contained in these kinds of datasets. For our application, however, one of the most important tasks is identifying the exact position of the tynes (which in turn identify the positions of each of the corresponding stent grafts). In our greyscale volume the location of tynes within the aneurism are obscured by the wall of the aneurism itself (Fig. 5a). Different levels of transparency were used to try to “see through” the aneurism, however, we could not find any appropriate value that showed the tyne locations within the artery, while at the same time providing an effective representation of the three-dimensional shape and extent of the wall of the aneurism. We decided that, for this application, it might be more appropriate to high the three regions of interest using color.

Although the radiologists had already chosen a set of colors based on context and aesthetic considerations, it did a poor job of showing the size and shape of the aneurism (Fig. 5b). We replaced their colors with three new ones using our color selection technique. The radiologists asked us to avoid greens and green-yellows, since these are associated with bile. We decided to use yellow to represent the artery, purple to represent the aneurism, and red to represent the tynes (Fig. 5c). These colors show clearly the location of all three regions of interest within the volume. For example, consider the large patches of yellow within the aneurism. These are areas of “low support” where the grafts in the lower part of the artery were not inserted far enough to mesh with their upstream partner. Although not dangerous, these are exactly the kinds of features the radiologists want to identify during post-operative visualization.

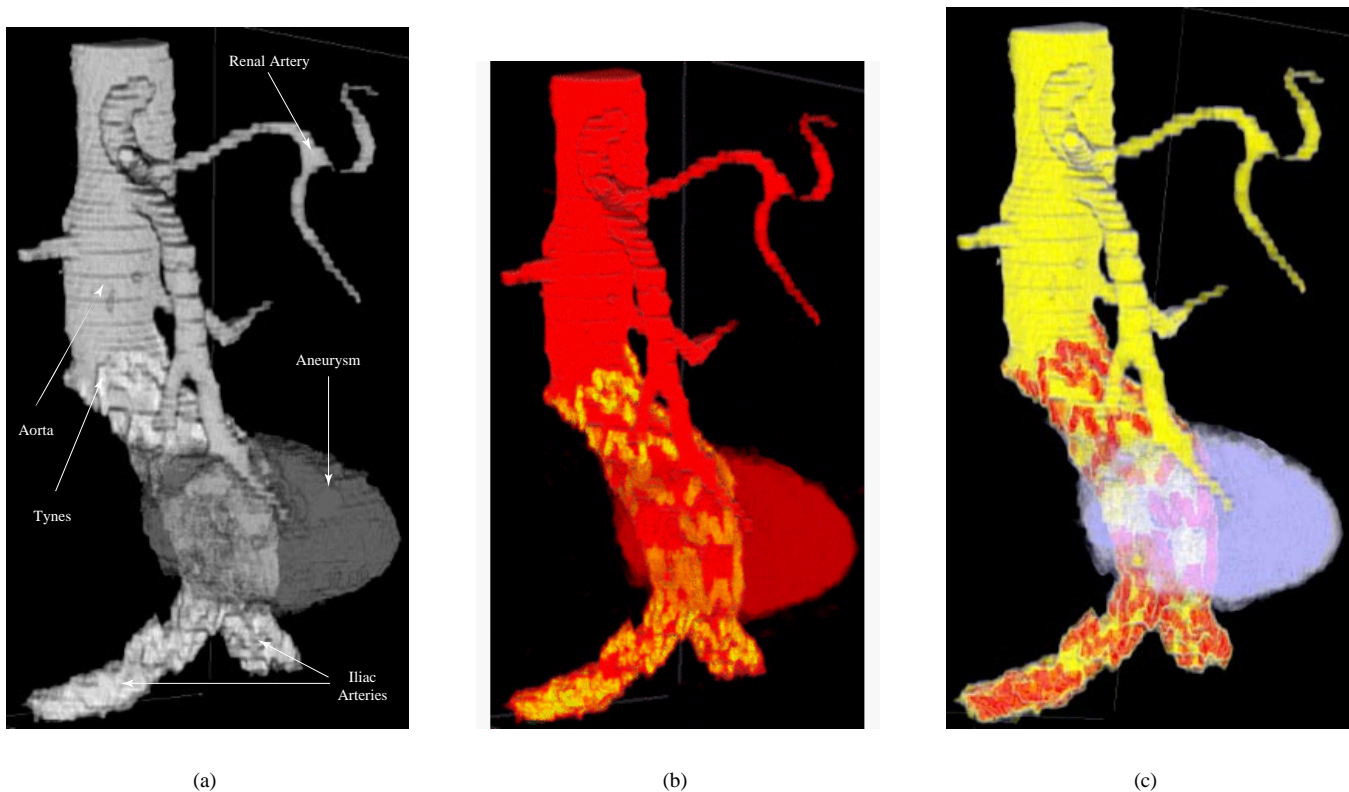


Figure 5: A reconstructed CT volume showing an abdominal aneurism: (a) greyscale hides the location of the tynes within the aneurism; (b) a color scale chosen by the radiologists obscures the shape and extent of the wall of the aneurism; (c) three colors chosen with our perceptual color selection technique

5 Perceptual Textures

One of the important issues in scientific visualization is designing methods for representing multiple values (or attributes) at a single spatial location. Although it is possible to assign different visual features to each attribute, simply “stacking” multiple features together will most likely lead to displays that are unintelligible.

Rather than choosing multiple individual features (*i.e.*, color, shape, size, orientation, line length), we decided to try using a single visual feature that could be decomposed into a set of fundamental parts or dimensions. We chose to investigate texture for this purpose.

Texture has been studied extensively in the computer vision, computer graphics, and cognitive psychology communities. Although each group focuses on separate tasks (texture identification and texture segmentation in computer vision, displaying information with texture patterns in computer graphics, and modeling the low-level human visual system in cognitive psychology) they each need ways to describe precisely the textures being identified, classified, or displayed.

Researchers have used different methods to study the perceptual features inherent in a texture pattern. Bela Julész [32] conducted numerous experiments that investigated how a texture’s first, second, and third-order statistics affect discrimination in the low-level human visual system. This led to the texton theory [33], which suggests that early vision detects three types of features (or textons, as Julész called them): elongated blobs with specific visual properties (*e.g.*, hue, orientation, and width), ends of line segments, and crossings of line segments. Tamura et al. [67] and Rao and Lohse [49, 50] identified texture dimensions by conducting experiments that asked subjects to divide pictures depicting different types of textures (Brodatz images) into groups. 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 subjects to group images: regularity, directionality, and complexity. Haralick et al. [19] built greyscale spatial dependency matrices to identify features like homogeneity, contrast, and linear dependency. These features were used to classify satellite images into categories like forest, woodlands, grasslands, and water. Liu and Picard [40] used Wold features to synthesize texture patterns. 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 [42] 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. We used these results to choose the perceptual texture dimensions we wanted to investigate during our experiments.

Work in computer graphics has studied methods for using texture patterns to display information during visualization. Schweitzer [59] used rotated discs to highlight the orientation of a three-dimensional surface. Pickett and Grinstein [17] built “stick-men” icons to produce texture patterns that show spatial coherence in a multidimensional dataset. Ware and Knight [80, 81] 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. Turk and Banks [75] described an iterated method for placing streamlines to visualize two-dimensional vector fields. Interrante [29] 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. Finally, Salisbury et al. [57] and Wikenbach and Salesin [82] used texturing techniques to build computer-generated pen-and-ink drawings that convey a realistic sense of shape, depth, and orientation. We built upon these results to try to develop an effective method for displaying multidimensional data through the use of texture.

5.1 Pexels

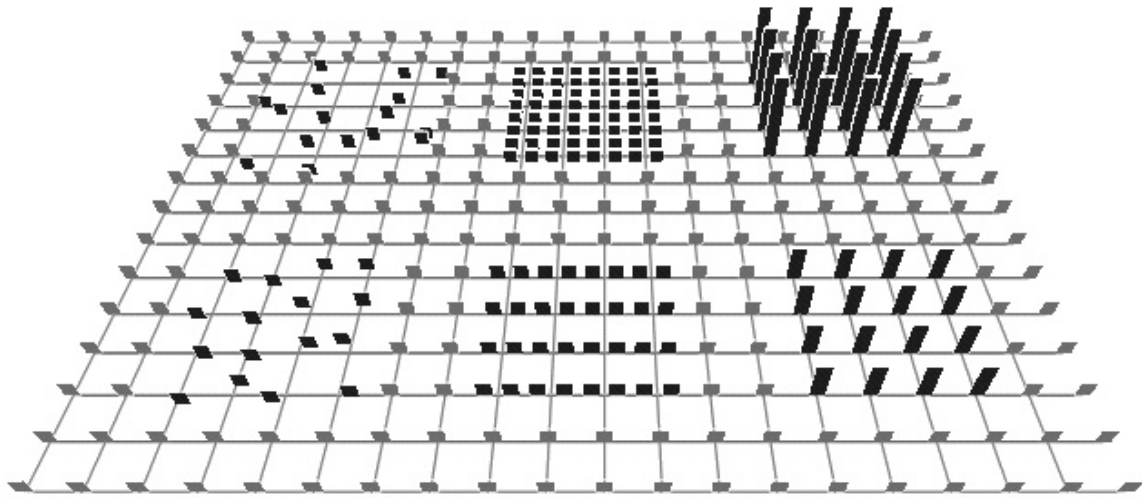
We wanted to design a technique that will allow users to visualize multidimensional datasets with perceptual textures. To this end, we used a method similar to Ware and Knight to build our displays. Each data element is represented with a single perceptual texture element, or pexel. Our visualization environment consists of a large number of elements arrayed across a three-dimensional surface (*e.g.*, a topographical map or the surface of a three-dimensional object). Each element contains one or more attributes to be displayed. Attribute values are used to control the visual appearance of a pexel by modifying its texture dimensions. Texture patterns formed by groups of spatially neighboring pexels can be used to visually analyze the dataset.

We chose to study three perceptual dimensions: density, regularity, and height. Density and regularity have been identified in the literature as primary texture dimensions [49, 50, 67]. Although height might not be considered an “intrinsic textural cue”, we note that height is one aspect of element size, and that element size is an important property of a texture pattern. Moreover, results from cognitive vision have shown that differences in height are detected preattentively by the low-level visual system [1, 70]. We wanted to build three-dimensional pexels that “sit up” on the underlying surface. This allows the possibility of applying various orientations (another important perceptual dimension) to a pexel. Because of this, we chose height as our third texture dimension.

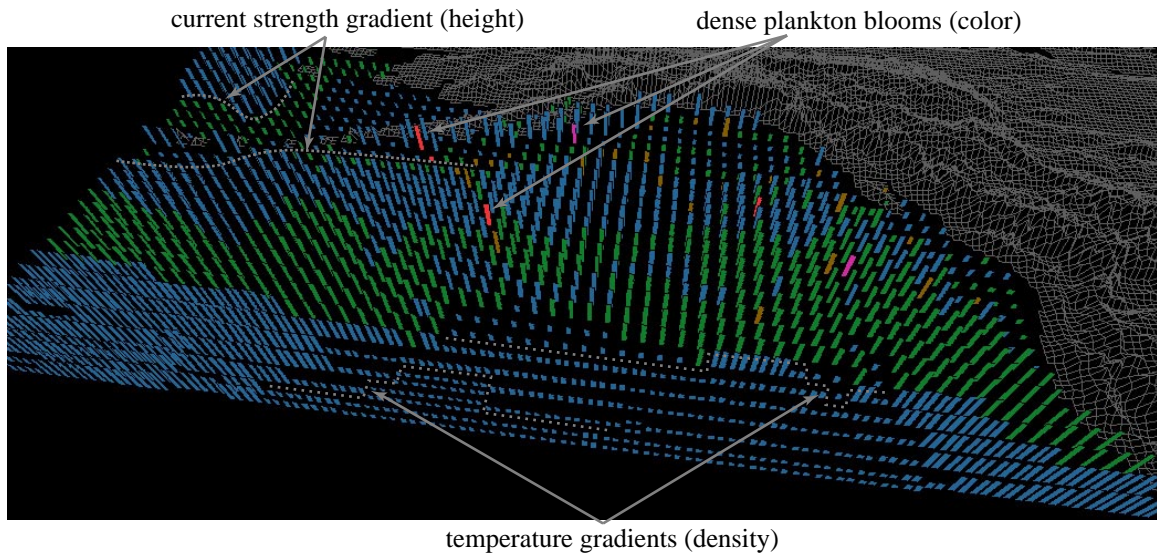
In order to support variation of height, density, and regularity, we built pexels that look like a collection of paper strips. The user maps attributes in the dataset to the density (which controls the number of strips in a pexel), height, and regularity of each pexel. Examples of each of these perceptual dimensions are shown in Fig. 6a. Fig. 6b 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. 6b is only one frame from a much larger time-series of historical ocean conditions. Our choice of visual features was guided by experimental results that show how different color and texture properties can be used in combination to represent multivariate data elements.

5.2 Experiments

In order to test our perceptual dimensions and the interactions that occur between them during visualization, we ran a set of psychophysical experiments. Our experiments were designed to investigate a user’s ability to rapidly and accurately identify target pexels defined by a particular height, density, or regularity. Users were asked to determine whether a small group of pexels with a particular type of texture (*e.g.*, a group of taller pexels, as in Fig. 7a) was present or absent in a 20×15 array.



(a)



(b)

Figure 6: Pexel examples: (a) 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; (b) 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




Figure 7: 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 8 grids step right and 2 grid steps up from the lower-left corner of the array)

Conditions like target pexel type, exposure duration, target group size, and background texture dimensions differed for each display. This allowed us to test for preattentive task performance, visual interference, and a user preference for a particular target type. In all cases, user accuracy was used to measure performance.

Each experimental display contained a regularly-spaced 20×15 array of pexels rotated 45° about the X-axis (Fig. 7). All displays were monochromatic (*i.e.*, grey and white), to avoid variations in color or intensity that might mask the underlying texture pattern. Grid lines were drawn at each row and column, to ensure users perceived the pexels as lying on a tilted 3D plane. After a display was shown, users were asked whether a group of pexels with a particular target value was present or absent. In order to avoid confusion, each user searched for only one type of target pexel: taller, shorter, sparser, denser, more regular, or more irregular. The appearance of the pexels in each display was varied to test for preattentive performance, visual interference, and feature preference. For example, the following experimental conditions were used to investigate a user's ability to identify taller pexels (similar conditions were used for the shorter, denser, sparse, regular, and irregular experiments):

- two target-background pairings: a target of medium pexels in a sea of short pexels, and a target of tall pexels in a sea of medium pexels; different target-background pairings allowed us to test for a subject preference for a particular target type,
- three display durations: 50 msec, 150 msec, and 450 msec; we varied exposure duration to test for preattentive performance, specifically, does the task become more difficult during shorter exposures,
- three secondary texture dimensions: none (every pexel is sparse and regular), density (half the pexels are randomly chosen to be sparse, half to be dense), and regularity (half the pexels are regular, half are random); we added a "background" texture feature to test for visual interference, that is, does the task become more difficult when a secondary texture dimension appears at random spatial locations in the display, and
- two target group sizes: 2×2 pexels and 4×4 pexels; we used different target group sizes to see how large a group of pexels was needed before the target could be detected by a viewer.

Our results suggest that pexels can be used to represent multidimensional data elements, but only if specific data-feature mappings are chosen. Some dimensions were more salient than others, and interference occurred when certain types of pexels were displayed. Specifically, we found that:

- taller pexels can be identified at preattentive exposure durations (*i.e.*, 150 msec or less) with very high accuracy (approximately 93%); background density and regularity patterns produce no significant interference,
- shorter, denser, and sparser pexels are more difficult to identify than taller pexels, although good results are possible at both 150 and 450 msec; height, regularity, and density background texture patterns cause interference for all three target types,

- irregular pixels are difficult to identify, although reasonable accuracy (approximately 76%) is possible at 150 and 450 msec with no background texture pattern, and
- regular pixels cannot be accurately identified; the percentage of correct results approached chance (*i.e.*, 50%) for every condition.

These results show that height and density can be used to form texture patterns that can be identified preattentively. Regularity, however, can only be used as a secondary dimension. While differences in regularity cannot be detected consistently by the low-level visual system, in many cases users will be able to see changes in regularity when areas of interest in a dataset are identified and analyzed in a focused or attentive fashion.

6 Orientation

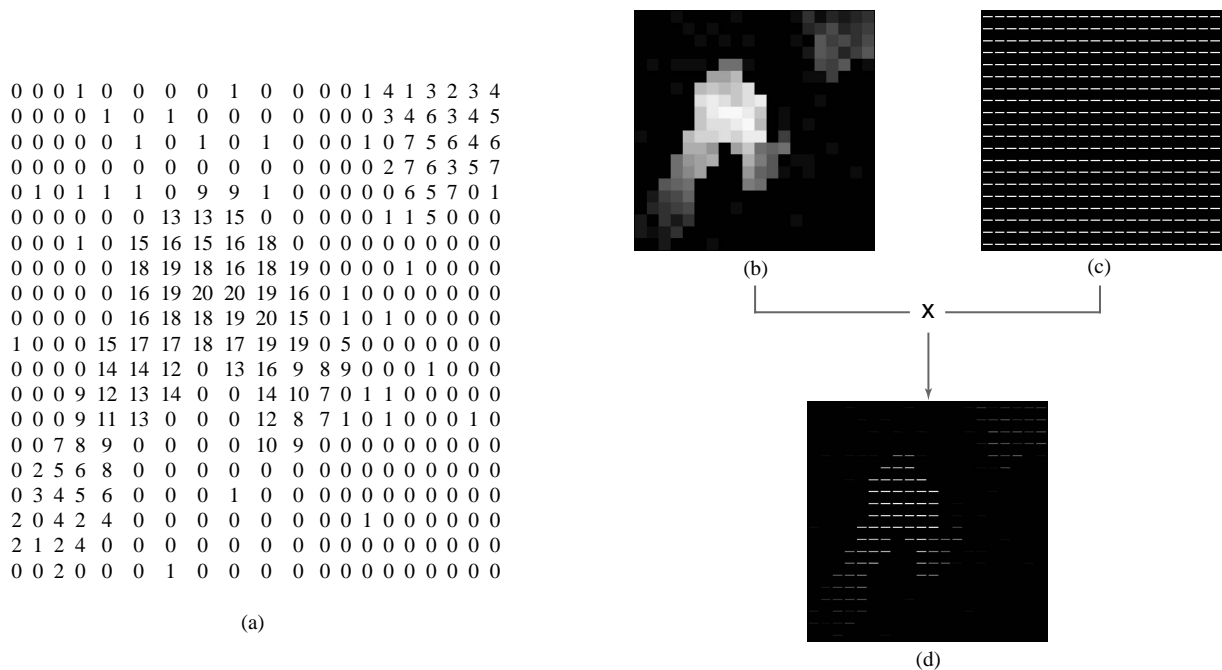


Figure 8: (a) A 20×20 patch of values from a scalar field; (b) the patch represented by greyscale swatches; (c) a collection of slivers oriented 0° at each data value location; (d) the greyscale map and slivers and combined to produce the final sliver layer

A follow-up study was recently conducted to test another perceptual texture dimension: orientation. This work was motivated in part by the desire to construct an alternative method for visualizing multiple overlapping data fields. The well-known method of selecting m visual features to represent each of the m data attributes has a number of inherent limitations:

- *dimensionality*: as the number of attributes n in the dataset grows, it becomes more and more difficult to find additional visual features to represent them.
- *interference*: different visual features will often interact with one another, producing visual interference; these interference effects must be controlled or eliminated to guarantee effective exploration and analysis.
- *attribute-feature matching*: different visual features are best suited to a particular type of attribute and analysis task; an effective visualization technique needs to respect these preferences.

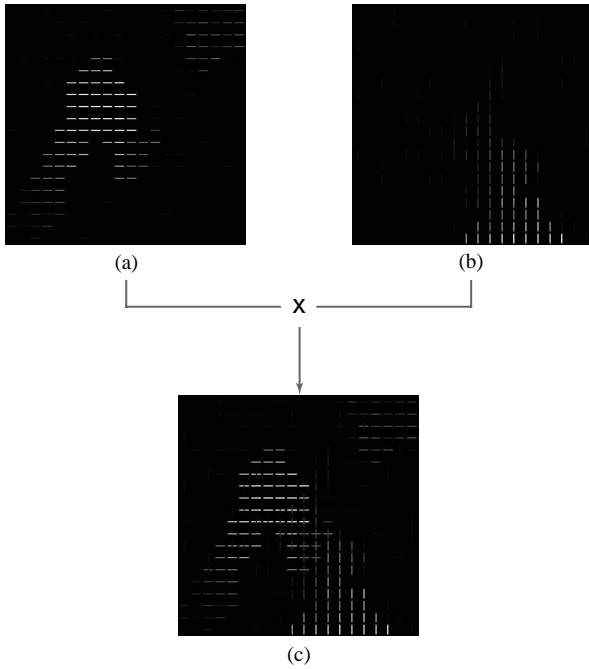


Figure 9: (a,b) two scalar fields represented with 0° and 90° , respectively; (c) both fields displayed in a single image, overlapping values show as elements that look like plus signs

ers are overlaid to produce the single image shown in Fig. 9c. This image allows the viewer to locate values in each individual field, while at the same time identifying important interactions between the fields. The use of thin, well separated slivers is key to allowing values from multiple fields to show through in a common spatial location. A viewer can use these images to:

- determine which fields are prominent in a region,
- determine how strongly a given field is present,
- estimate the relative weights of the field values in the region, and
- locate regions where all the fields have low, medium, or high values.

6.1 Experiments

Our initial experiments investigated a viewer’s ability to distinguish sliver textures with different 2D orientations (*i.e.*, slivers with different rotations embedded in the XY-plane). Each trial contained a 20×20 grid of rectangles rotated bg . A randomly selected 2×2 patch of target rectangles was then rotated tg . In half the trials $bg = tg$ (*i.e.*, the target was absent, Fig. 10b). In the other half, $bg \neq tg$ (*i.e.*, a target was present, Fig. 10a and 10c). We tested background orientations ranging from 0 - 45° and 45 - 90° in 5° steps. For each background, every orientation was tested as a target (*i.e.*, ten target rotations $0, 5, \dots, 45^\circ$ were tested for each background in the range 0 - 45° ; ten target rotations $45, 50, \dots, 90^\circ$ were tested for each background in the range 45 - 90°).

In summary, our results showed:

Multidimensional datasets can often be viewed as a collection of m scalar fields that overlap spatially with one another. Rather than using m visual features to represent these fields, we can use only two: orientation and luminance. For each scalar field (representing attribute A_j) we select a constant orientation o_j ; at various spatial locations where a value $a_{i,j} \in A_j$ exists, we place a corresponding sliver texture oriented at o_j . The luminance of the sliver texture depends on $a_{i,j}$: the maximum $a_{max-j} \in A_j$ produces a white (full luminance) sliver, while the minimum $a_{min-j} \in A_j$ produces a black (zero luminance) sliver. A perceptually-balanced luminance scale running from black to white is used to select a luminance for an intermediate values $a_{i,j}$, $a_{min-j} < a_{i,j} < a_{max-j}$.

Fig. 8a shows a uniformly-sampled patch from a hypothetical scalar field. Values in the field are represented as greyscale swatches in Fig. 8b. A constant orientation of 0° is used to represent values in the field (slivers rotated 0° a replaced at the spatial locations for each reading in the field, shown in Fig. 8c). Blending these two representations together produces the final image (Fig. 8d), a layer of variable-luminance slivers showing the positions and values of all the data in the original field.

Multiple scalar fields are displayed by composing their sliver layers together. Fig. 9a-b shows two separate sliver layers representing two scalar fields. The first field uses slivers oriented 0° ; the second uses slivers oriented 90° . When a viewer visualizes both fields simultaneously, the sliver lay-




Figure 10: An example of three experiment displays: (a) a 10° target in a 0° background (target is five steps right and eight steps up from the lower left corner of the array); (b) a 30° background with no target; (c) a 65° target in a 55° background (target is six steps left and seven steps up from the lower right corner of the array)

1. A target oriented $\pm 15^\circ$ or more from its background elements is rapidly and accurately distinguishable.
2. Errors for backgrounds oriented 0° or 90° were significantly lower than for other backgrounds (*e.g.*, tilted targets are easier to see in backgrounds of 0° or 90°).
3. Errors for targets oriented 0° or 90° were significantly higher than for other targets, suggesting an asymmetry (good as a background, bad as a target) for these orientations.

7 Scanning Electron Microscope Images


The application for which this technique was originally developed is the display of multiple data fields from a scanning electron microscope (SEM). Each field represents the concentration of a particular element (oxygen, silicon, carbon, and so on) across a surface. Physicists studying mineral samples need to determine what elements make up each part of the surface and how those elements mix. By allowing the viewer to see the relative concentrations of the elements in a given area, our technique enables recognition of composites more easily than side-by-side comparison, especially for situations where there are complex amalgams of materials.

Fig. 11a shows sliver layers representing eight separate elements: calcium (15°), copper (30°), iron (60°), magnesium (75°), manganese (105°), oxygen (120°), sulphur (150°), and silicon (165°). The orientations for each layer were chosen to ensure no two layers have an orientation difference of less than 15° . Fig. 11b shows the eight layers blended together to form a single image. Fig. 11c changes the orientations of silicon and oxygen to 90° and 180° , respectively, to investigate a potential interaction between the two (the presence of silicon oxide in the upper right, upper left, and lower left where regions of “plus sign” textures appear).


Our results suggest up to 15 orientations in the 0 - 180° range can be rapidly and accurately differentiated. The greyscale ramp used to assign a luminance to each sliver is also constructed to be perceptually linear. The result is an image that shows data values in each individual field, while at the same time highlighting important interactions between the fields.

8 Combining Color and Texture


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



(a)



(b)



(c)

Figure 11: (a) eight sliver layers representing calcium (15°), copper (30°), iron (60°), magnesium (75°), manganese (105°), oxygen (120°), sulphur (150°), and silicon (165°), (b) all eight layers blended into a single image; (c) silicon and oxygen re-oriented at 90° and 180° , respectively, to highlight the presence of silicon oxide (as a “plus sign” texture) in the upper right, upper left, and lower left corners of the image

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 [7, 8]. 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) [73]. Recent work by Snowden [64] 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 left unanswered the question of whether color variation would interfere with texture identification during visu-

alization. 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.

8.1 Experiments

In order to investigate the combined use of color and texture, 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 [25, 26]. We chose three different colors using our perceptual color selection technique [20, 25]. 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. We varied target-background pairings, exposure duration, the presence or absence of a secondary background feature, and target patch size from trial to trial to test user performance in a number of different circumstances.

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 [25, 26]. 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.
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%).

Our results showed an asymmetric interference effect, similar to those reported in the psychophysical literature. As described by [64], 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 [7, 8], who reported asymmetric color on shape interference in a boundary detection task involving two-dimensional textures. The amount of color on texture interference depended on the difficulty of the segmentation task. Targets that were harder to identify in isolation (*e.g.*, shorter and sparser targets) showed a much higher sensitivity to color interference, compared to targets that were easy to identify (*e.g.*, taller and denser targets). This suggests that color and texture can be combined in a single display, but only in cases where the texture targets have a strong perceptual salience.

9 Nonphotorealistic Visualization

For many years the areas of modeling and rendering in computer graphics have studied the problem of producing photorealistic images, images of graphical models that are indistinguishable from photographs of an equivalent real-world scene. Advances in areas like the simulation of global light transport, modeling of natural phenomena, and image-based rendering have made dramatic strides towards achieving this goal. At the same time, researchers have approached the issue of image generation from a completely different direction. Although photographs are common, there are many other compelling methods of visual discourse, for example, oil and watercolor paintings, pen and ink sketches, cel animation, and line art. In the proper situations, these types of pictures are often considered more effective, more appropriate, or even more expressive than an equivalent photograph. The study of methods that construct images of these types is known as *nonphotorealistic rendering*.

Our current interests lie specifically in nonphotorealistic rendering methods that use simulated brush strokes to produce images that look like paintings. Strassmann [65] constructed a “hairy brush”, a collection of bristles placed along a line segment; Japanese-style brush strokes are produced by applying ink to the brush, then moving it along a path over a simulated paper surface. Haberli [18] built a system that paints with a brush that a user strokes across an underlying image; the size, shape, color, location, and direction of brush strokes can all be varied. Hsu et al. [28] used vector-based skeletal strokes with variable thickness drawn along a parametric curve to produce interesting line-art images. Meier [44] attached particles to surfaces in a 3D scene, then rendered a brush stroke (with variable color, size, and direction) at each particle position to paint the scene. Litwinowicz [39] clipped strokes to object boundaries, then rendered them as lines or texture maps with variable length, thickness, direction, and color. Curtis et al. [10] built a fluid-flow simulation to model the interactions of brush, watercolor, and paper during the painting of watercolor images. Shiraishi and Yamaguchi [60] computed image moments on a target image to control the color, location, orientation, and size of the brush strokes in their painterly rendering. Hertzmann [27] used a multilayer painting technique, where each new layer contains finer details painted with smaller brush strokes; brush paths are modeled with splines, while the brush itself allows variation of length, size, opacity, placement, and color jitter. Interrante [30] discussed applying natural textures to visualize multidimensional datasets. Laidlaw [36, 37] extended the layered approach of Meier to visualize multidimensional data in a painterly fashion. Finally, Ebert and Rheingans [13] used nonphotorealistic techniques like silhouettes, sketch lines, and halos to highlight important features in a volumetric dataset.

9.1 Painterly Styles

We believe that fundamental properties of a nonphotorealistic image can be identified in part by studying the styles used by artists to construct their paintings. Our investigation of painterly styles is directed by two separate criteria. First, we are restricting our search to a particular movement in art known as Impressionism. Second, we attempt to pair each style with a corresponding visual feature that has been shown to be effective in a perceptual visualization environment. There are no technical reasons for why Impressionism was chosen over any other movement. In fact, we expect the basic theories behind our technique will extend easily to other types of artistic presentation. For our initial work, however, we felt it was important to narrow our focus to a set of fundamental goals in the the context of a single type of painting style.

The term “Impressionism” was attached to a small group of French artists (initially including Monet, Degas, Manet, Renoir, and Pissaro, and later Cézanne, Sisley, and Van Gogh, among others) who broke from the traditional schools of the time to approach painting from a new perspective. The Impressionist technique was based on a number of underlying principles [5, 58, 77], for example:

1. *Object and environment interpenetrate.* Outlines of objects are softened or obscured (*e.g.*, Monet’s water lilies); objects are bathed and interact with light; shadows are colored and movement is represented as unfinished outlines.
2. *Color acquires independence.* There is no constant hue for an object, atmospheric conditions and light moderate color across its surface; objects may reduce to swatches of color.
3. *Show a small section of nature.* The artist is not placed in a privileged position relative to nature; the world is shown as a series of close-up details.
4. *Minimize perspective.* Perspective is shortened and distance reduced to turn 3D space into a 2D image.

5. *Solicit a viewer's optics.* Study the retinal system; divide tones as separate brush strokes to vitalize color rather than graying with overlapping strokes; harness simultaneous contrast; use models from color scientists like Chevreul [9] or Rood [55].

Although these general characteristics are perhaps less precise than we might prefer, we can still draw a number of important conclusions. Properties of hue, luminance, and lighting were explicitly controlled and even studied in a scientific fashion by some of the Impressionists. Rather than using an “object-based” representation, the artists appear to be more concerned with subdividing a painting based on the interactions of light with color and other surface properties. Additional painterly styles can be identified by studying the paintings themselves. These styles often varied dramatically between individual artists, acting to define their unique painting techniques. Examples include:

- *path*: the direction a brush stroke follows; Van Gogh made extensive use of curved paths to define boundaries and shape in his paintings; other artists favored simpler, straighter strokes,
- *length*: the length of individual strokes on the canvas, often used to differentiate between contextually different parts of a painting,
- *density*: the number and size of strokes laid down in a fixed area of canvas,
- *coarseness*: the coarseness of the brush used to apply a stroke; a coarser brush causes visible bristle lines and surface roughness, and
- *weight*: the amount of paint applied during each stroke; heavy strokes highlight coarseness and produce ridges of paint that cause underhanging shadows when lit from the proper direction.

Although this collection of painterly styles provides a good starting point, it is by no means exhaustive. All of the styles we use are evaluated for effectiveness by identifying their perceptual characteristics, and by validating their ability to support visualization, discovery, analysis, and presentation in a real-world application environment.

A comparison of perceptual color and texture properties with painterly styles from Impressionist art reveals a strong correspondence between the two. Reduced to perceptual elements, color and texture are the precise properties that an artist varies in the application of colored pigments of paint to a canvas with a brush. From this perspective, color and lighting in Impressionism has a direct relationship to the use of hue and luminance in perceptual vision. Other painterly styles (*e.g.*, path, density, and length) have similar partners in perception (*e.g.*, orientation, contrast, and size). This close correspondence between perceptual features and many of the painterly styles we hope to apply is particularly advantageous. Since numerous controlled experiments on the use of perception have already been conducted, we have a large body of evidence to use to predict how we expect painterly styles to react in a multidimensional visualization environment.

9.2 Experiments

We conducted a set of psychophysical experiments to test our hypothesis that guidelines from human perception will extend to a painterly environment. Our experiments were designed to investigate a viewer’s ability to rapidly and accurately identify target brush strokes defined by a particular color or orientation. Background orientation, color, regularity, and density varied between displays. This allowed us to test for preattentive task performance, and for visual interference effects. The experimental results were then used to identify similarities and differences between painterly images and existing perceptual visualization techniques.

Each experimental display contained a 22×22 array of simulated brush strokes (Fig. 12). Viewers were asked to determine whether a small, 3×3 group of strokes with a particular target type was present or absent in each display. Displays were shown for 200 milliseconds, after which the screen was cleared; the system then waited for viewers to enter their answer (either “target present” or “target absent”).

The displays were equally divided into two groups: one studied a viewer’s ability to identify target strokes based on color, the other studied identification based on orientation. The appearance of the strokes in each display was varied to test for preattentive performance and visual interference. For example, Fig. 12a shows an orange target in a sea of pink strokes; all the strokes have a constant orientation of 30^{circ} ; they are sparsely packed, and are located in a completely regular, grid-like pattern.

Fig. 12b shows a green target in a sea of orange strokes; these strokes, however, have a random orientation, a dense packing, and an irregular placement. Fig. 12c shows a 45° orientation target in a sea of strokes rotated 30°; the strokes have a random background color, very dense packing, and irregular placement. Finally, Fig. 12d shows a 60° target in a 45° background; these strokes have a constant color, dense packing, and regular placement.




Figure 12: Example experiment displays: (a) orange target in pink strokes, constant 30° background orientation, sparse packing, regular placement; (b) green target in orange strokes, random background orientation, dense packing, irregular placement; (c) 45° target in 30° strokes, random background color, very dense packing, irregular placement; (d) 60° target in 45° strokes, constant background color, dense packing, regular placement

Viewer accuracy and response time were combined and tested for significance using a multi-factor analysis of variance (ANOVA). In summary, we found:

- color targets were easy to detect at a 200 millisecond preattentive exposure duration; a random orientation pattern had no interfering effect on performance,
- orientation targets were easy to detect when a constant color was displayed in the background; a random background color pattern caused a significant reduction in performance,
- background density had a significant effect on both color and orientation targets; denser displays produced an improvement in performance, and
- background regularity had a significant effect on both color and orientation targets; irregular displays caused a reduction in performance.

Our results match previous findings in both the psychophysical and the visualization literature, specifically: (1) color produces better performance than orientation during target identification, and (2) an asymmetric color on texture interference effect exists (random color patterns interfere with orientation identification, but not vice-versa). Both results have been shown to exist in experimental [8, 64] and real-world visualization settings [25, 26]. These results extend our previous work [26], which found a general color on texture interference pattern, but no corresponding texture on color effect. Overall, our results provide positive evidence to support the belief that perceptual findings will carry to a painterly visualization environment.

10 Visualizing Environmental Weather Data

Based on the results from our experiments, we decided to build a nonphotorealistic visualization system that varied brush stroke color, spatial density, direction (*i.e.*, orientation), and stroke placement (*i.e.*, regularity) to encode multiple data attributes (in addition to the two spatial values used to position each stroke). The presence of feature hierarchies suggest color should be used to represent the most important attribute, followed by texture properties. Our results further refine this to mapping color, direction, density, and placement in order of attribute importance (from most important to least important).

The brush strokes in our current prototype are similar to the ones shown during our experiments. They are constructed using a simple texture mapping scheme. This technique is common in nonphotorealistic rendering (e.g., in [18, 27, 39, 44]). Real painted strokes are captured and converted into texture maps. These textures are applied to an underlying polygon to produce an approximation of a collection of generic brush strokes. We currently use a small library of five representative brush stroke textures; one of the textures is randomly selected and applied when a stroke is rendered. A stroke’s color, direction, density, and placement are controlled by modifying its texture and by transforming the polygon it maps to.

One of the application testbeds for our nonphotorealistic visualization technique is a collection of monthly environmental and weather conditions collected and recorded by the Intergovernmental Panel on Climate Change. This dataset contains mean monthly surface climate readings in $\frac{1}{2}^\circ$ latitude and longitude steps for the years 1961 to 1990 (e.g., readings for January averaged over the years 1961-1990, readings for February averaged over 1961-1990, and so on). We chose to visualize values for mean temperature, windspeed, pressure, precipitation, and frost frequency (or *temp*, *wind pressure*, *precip*, and *frost*). Based on this order of importance, we built a data-feature mapping M that varies brush stroke color, density, orientation, and regularity. We divided density into two separate parts: *energy*, the number of strokes used to represent a data element e_i , and *coverage*, the percentage of e_i ’s screen space covered by its strokes. Both properties represent painterly styles. Energy describes the number and vitality of strokes in a fixed region of a painting (e.g., a few long, broad, lazy strokes or many small, short, energetic strokes). Coverage describes the amount of the underlying canvas, if any, that shows through the strokes. This produced the following attribute to visual feature pairings:

- *temp* → *color*: dark blue for low *temp* to bright pink for high *temp*,
- *wind* → *coverage*: low coverage for weak *wind* to full coverage for strong *wind*,
- *pressure* → *energy*: a single stroke, a 1×2 array of strokes, or a 2×2 array of strokes for low to high *pressure*,
- *precip* → *orientation*: upright (90° rotation) for light *precip* to flat (0° rotation) for heavy *precip*, and
- *frost* → *regularity*: regular for low *frost* frequency to irregular for high *frost* frequency.

Fig. 13 shows an example of applying M to data for February along the east coast of the continental United States. The top five images use a perceptual color ramp (running from dark blue and green for small values to bright red and pink for large values) to show the individual variation in *temp*, *pressure*, *wind*, *precip*, and *frost*. The result of applying M to construct a nonphotorealistic visualization of all five attributes is shown in the bottom image. Various color and texture patterns representing different weather phenomena are noted on this image. Changes in temperature are visible as a smooth blue-green to red-pink color variation running north to south over the map. Pressure gradients produce energy boundaries, shown as regions with different numbers of strokes packed into a unit area of screen space (e.g., higher energy strokes in Florida represent higher *pressure* readings). Windspeed modifies stroke coverage: weak *wind* values produce small strokes with a large amount of background showing through (e.g., north of the Great Lakes), while strong *wind* values produce larger strokes that completely fill their corresponding screen space (e.g., in central Texas and Kansas). Increases in rainfall are shown as an increasing stroke tilt running from upright (light *precip*) to flat (heavy *precip*). Finally, a higher frost frequency produces more irregularity (e.g., strokes in Florida and southern Texas are completely regular, while strokes in the northern states and Canada are highly irregular).

Fig. 14 uses the same mapping M to visualize weather conditions over the western United States for January and August. These visualizations provide a number of interesting insights into historical weather conditions for this part of the continent. In January (Fig. 14a) weak *wind* values (shown as small, low coverage strokes) highlight the locations of the Rocky Mountains, the Cascades, and the Sierra Nevada range. Typically heavy rainfall in the Pacific Northwest is represented by nearly flat strokes. Regions of severe cold east of the Rocky Mountains near Denver and in the northern plains and Canadian prairies appear as patches of dark green and blue strokes. Low pressure (i.e., a single low energy stroke for each data element) and high frost frequency (shown as irregular stroke placement) covers most of the map. Conditions in August (Fig. 14b) are markedly different. Most of the western United States is warm (a completely regular placement of strokes denoting little or no frost during this month). An area of intense heat, shown as bright pink strokes, is visible in southern California and Arizona. High pressure regions cover the coast, the south, and much of the central and northern plains. Little or no rainfall is evident. Finally, an area of weak *wind* values is visible as small, low coverage strokes in northern Washington, Idaho, and Montana.





Figure 13: Nonphotorealistic visualization of weather conditions over the eastern United States: (top row) perceptual color ramps (dark blue for low to bright pink for high) of mean temperature, pressure, windspeed, precipitation, and frost frequency in isolation; (bottom row) combined visualization of temperature (dark blue to bright pink for cold to hot), pressure (low to to high energy for low to high), windspeed (low to high coverage for weak to strong), precipitation (upright to flat for light to heavy), and frost frequency (regular to irregular for low to high)




Figure 14: Weather conditions over the western United States: (a) mean *temp, pressure, wind, precip* and *frost* for January; (b) mean conditions for August

11 Visualizing Typhoon Data

One of our current testbeds for using pexels to visualize multidimensional data is the analysis of environmental conditions related to typhoons. We used pexels to visualize typhoon activity in the Northwest Pacific Ocean during the summer and fall of 1997. 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 we needed. A U.S. Navy elevation dataset² was used to obtain land elevations at ten minute latitude and longitude intervals. Land-based weather station readings collected from around the world and archived by the National Climatic Data Center³ provided daily measurements for eighteen separate environmental conditions. Finally, satellite archives made available by the Global Hydrology and Climate Center⁴ 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:

²<http://grid2.cr.usgs.gov/dem/>

³<http://www.ncdc.noaa.gov/ol/climate/online/g sod.html>

⁴<http://ghrc.msfc.nasa.gov/ghrc/list.html>

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 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. 15b and 15d). 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. 15 shows windspeed, pressure, and precipitation around Japan, Korea, and Taiwan during August 1997. Fig. 15b, 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. 15c 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. 15d and 15e, where windspeed was mapped to regularity, pressure to height, and precipitation to density (a mapping without color that our original texture experiments predict will 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. 15d and 15e.

12 Oceanography Simulations


Our final example describes a set of oceanography simulations being jointly conducted at North Carolina State University and the University of British Columbia [21]. Researchers in oceanography are studying the growth and movement patterns of different species of salmon in the northern Pacific Ocean. Underlying environmental conditions like plankton density, sea surface temperature (SST), current direction, and current strength affect where the salmon live and how they move and grow [68]. For example, salmon like cooler water and tend to avoid ocean locations above a certain temperature. Since the salmon feed on plankton blooms, they will try to move to areas where plankton density is highest. Currents will “push” the salmon as they swim. Finally, SST, current direction, and current strength affect the size and location of plankton blooms as they form.




(a)



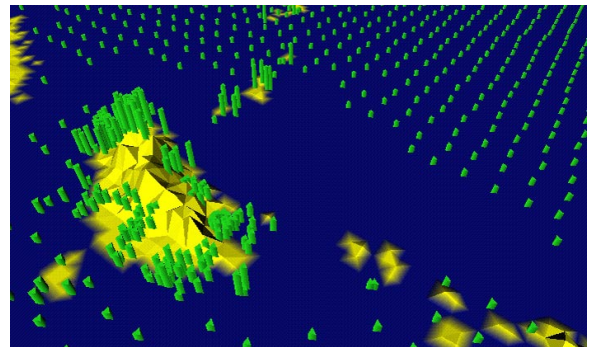
(b)



(c)



(d)



(e)

Figure 15: Typhoon conditions across Southeast Asia during the summer of 1997: (a) August 7, normal weather conditions over Japan; (b) August 27, typhoon Amber approaches the island of Taiwan from the southeast; (c) August 28, 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.

The oceanographers are designing models of how they believe salmon feed and move in the open ocean. These simulated salmon will be placed in a set of known environmental conditions, then tracked to see if their behavior mirrors that of the real fish. For example, salmon that migrate back to the Fraser River to spawn chose one of two routes. When the Gulf of Alaska is warm, salmon make landfall at the north end of Vancouver Island and approach the Fraser River primarily via a northern route through the Johnstone Strait (the upper arrow in Fig. 16). When the Gulf of Alaska is cold, salmon are distributed further south, make landfall on the west coast of Vancouver Island, and approach the Fraser River primarily via a southern route through the Juan de Fuca Strait (the lower arrow in Fig. 16). The ability to predict salmon distributions from prevailing environmental conditions would allow the commercial fishing fleet to estimate how many fish will pass through the Johnstone and Juan de Fuca straits. It would also allow more accurate predictions of the size of the salmon run, helping to ensure that an adequate number of salmon arrive at the spawning grounds.

In order to test their hypotheses, the oceanographers have created a database of SSTs and ocean currents for the region 35° north latitude, 180° west longitude to 62° north latitude, 120° west longitude (Fig. 16). Measurements within this region are available at 1° × 1° grid spacings. This array of values exists for each month for the years 1956 to 1964, and 1980 to 1989.

Partial plankton densities have also been collected and tabulated; these are obtained by ships that take readings at various positions in the ocean. We estimated missing plankton values using a set of knowledge discovery (KD) algorithms that we have specifically modified for use during visualization. Our KD algorithms identified month, SST, and current magnitude as the attributes used to estimate missing plankton values. Because of this, we restricted our initial visualization to a monthly time-series of plankton density, SST, and current magnitude.

Displaying the three attributes together allows the oceanographers to search for relationships between plankton density, current strength, and SST. Plankton is displayed using color; SST and current strength are displayed using texture. Colors for the five plankton ranges were chosen using our color selection technique [20]. Although other color scales were available (for example, by Ware [78]), our colors are specifically designed to highlight outliers, and to show clearly the boundaries between groups of elements with a common plankton density. We display the five plankton density ranges from low to high using blue (monitor RGB=36, 103, 151), green (monitor RGB=18, 127, 45), brown (monitor RGB=134, 96, 1), red (monitor RGB=243, 51, 55), and purple (monitor RGB=206, 45, 162),

For the underlying texture, we mapped current strength to height and SST to density. Our choices were guided by results we observed from our texture experiments:

- differences in height (specifically, taller elements) may be easier to detect, compared to differences in density or randomness,
- variation in height may mask differences in density or randomness; this appears to be due to the occlusion that occurs when tall pixels in the foreground hide short pixels in the background; this will be less important when users can control their viewpoint into the dataset (our visualization tool allows the user to interactively manipulate the viewpoint), and
- tightly spaced grids can support up to three easily distinguishable density patterns; placing more strips in a single pixel (e.g., arrays of 3 × 3 or 4 × 4 strips) will either cause the strips to overlap with their neighbors, or make each strip too



Figure 16: Map of the North Pacific; arrows represent possible salmon migration paths as they pass through the either Johnstone Strait (upper arrow) or the Strait of Juan de Fuca (lower arrow)

thin to easily identify.


Because there may be a feature preference for height over density, and because current strength was deemed “more important” than SST during knowledge discovery, we used height to represent currents and density to represent SSTs. The five ranges of current strength are mapped to five different heights. We do not use a linear mapping, rather the lower two ranges (corresponding to the weakest currents) are displayed using two types of short pexels, and the upper three ranges (corresponding to the strongest currents) are displayed using three types of tall pexels. This allows a user to rapidly locate boundaries between weak and strong currents, while still being able to identify each of the five ranges. For SSTs, the lower three ranges (corresponding to the coldest SSTs) are displayed with a pexel containing a single strip, while the upper two ranges (corresponding to the warmest SSTs) are displayed with pexels containing arrays of 2×1 and 2×2 strips, respectively. The densities we chose allow a user to see clearly the boundaries between cold and warm temperature regions. If necessary, users can change the range boundaries to focus on different SST gradients.

The oceanographers want to traverse their datasets in monthly and yearly steps. Experiments run in our laboratory have shown that preattentive tasks performed on static frames can be extended to a dynamic environment, where displays are shown one after another in a movie-like fashion [23]. Our visualization tool was designed to allow users to scan rapidly forwards and backwards through the dataset. This makes it easy to compare changes in the value and location of any of the environmental variables being displayed. The oceanographers can track seasonal changes in current strength, SST, and plankton density as they move month by month through a particular year. They can also see how interannual variability affects the environmental conditions and corresponding plankton densities for a particular month across a range of years.


Fig. 17 shows three frames from the oceanography dataset: February 1956, June 1956, and October 1956. Color shows the seasonal variation in plankton densities. Height and density allow the oceanographers to track current strengths and SSTs. In February (Fig. 17a), most plankton densities are less than 28 g/m^3 (*i.e.*, blue and green strips). Currents are low in the north-central Pacific; a region of weak currents also sits off the south coast of Alaska. Most of the ocean is cold (sparse pexels), although a region of higher temperatures can easily be seen as dense pexels in the south. In June (Fig. 17b) dense plankton blooms (red and purple strips) are present across most of the northern Pacific. The positions of the strong currents have shifted (viewing the entire dataset shows this current pattern is relatively stable for the months March to August). Warmer SSTs have pushed north, although the ocean around Alaska and northern British Columbia is still relatively cold. By October the plankton densities have started to decrease (green, brown, and red strips); few high or low density patches are visible. Current strengths have also decreased in the eastern regions. Overall a much larger percentage of the ocean is warm (*i.e.*, dense pexels). This is common, since summer temperatures will sometimes last in parts of the ocean until October or November.

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] 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.
- [3] BECK, J., PRAZDNY, K., AND ROSENFELD, A. A theory of textural segmentation. In *Human and Machine Vision*, J. Beck, K. Prazdny, and A. Rosenfeld, Eds. Academic Press, New York, New York, 1983, pp. 1–39.
- [4] 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.
- [5] BROWN, R. Impressionist technique: Pissarro’s optical mixture. In *Impressionism in Perspective*, B. E. White, Ed. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1978, pp. 114–121.
- [6] CALLAGHAN, T. C. Dimensional interaction of hue and brightness in preattentive field segregation. *Perception & Psychophysics* 36, 1 (1984), 25–34.
- [7] CALLAGHAN, T. C. Interference and domination in texture segregation: Hue, geometric form, and line orientation. *Perception & Psychophysics* 46, 4 (1989), 299–311.



(a)



(b)



(c)

Figure 17: Visualization of the oceanography datasets, color used to represent plankton density (blue, green, brown, red, and purple represent lowest to highest densities), height used to represent current strength, texture density used to represent SST: (a) February, 1956; (b) June, 1956; (c) October, 1956

- [8] 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.
- [9] CHEVREUL, M. E. *The Principles of Harmony and Contrast of Colors and Their Applications to the Arts*. Reinhold Publishing Corporation, New York, New York, 1967.
- [10] CURTIS, C. J., ANDERSON, S. E., SEIMS, J. E., FLEISCHER, K. W., AND SALESIN, D. H. Computer-generated watercolor. In *SIGGRAPH 97 Conference Proceedings* (Los Angeles, California, 1997), T. Whitted, Ed., pp. 421–430.
- [11] DRIVER, J., MCLEOD, P., AND DIENES, Z. Motion coherence and conjunction search: Implications for guided search theory. *Perception & Psychophysics* 51, 1 (1992), 79–85.
- [12] D’ZMURA, M. Color in visual search. *Vision Research* 31, 6 (1991), 951–966.
- [13] EBERT, D., AND RHEINGANS, P. Volume illustration: Non-photorealistic rendering of volume models. In *Proceedings Visualization 2000* (San Francisco, California, 2000), pp. 195–202.
- [14] EGETH, H. E., AND YANTIS, S. Visual attention: Control, representation, and time course. *Annual Review of Psychology* 48 (1997), 269–297.
- [15] ENNS, J. T. The promise of finding effective geometric codes. In *Proceedings Visualization ’90* (San Francisco, California, 1990), pp. 389–390.
- [16] ENNS, J. T., AND RENSINK, R. A. Sensitivity to three-dimensional orientation in visual search. *Psychology Science* 1, 5 (1990), 323–326.
- [17] GRINSTEIN, G., PICKETT, R., AND WILLIAMS, M. EXVIS: An exploratory data visualization environment. In *Proceedings Graphics Interface ’89* (London, Canada, 1989), pp. 254–261.
- [18] HABERLI, P. Paint by numbers: Abstract image representations. *Computer Graphics (SIGGRAPH 90 Conference Proceedings)* 24, 4 (1990), 207–214.
- [19] 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.
- [20] HEALEY, C. G. Choosing effective colours for data visualization. In *Proceedings Visualization ’96* (San Francisco, California, 1996), pp. 263–270.
- [21] HEALEY, C. G. One the use of perceptual cues and data mining for effective visualization of scientific datasets. In *Proceedings Graphics Interface ’98* (Vancouver, Canada, 1998), pp. 177–184.
- [22] HEALEY, C. G., BOOTH, K. S., AND ENNS, J. T. Harnessing preattentive processes for multivariate data visualization. In *Proceedings Graphics Interface ’93* (Toronto, Canada, 1993), pp. 107–117.
- [23] 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.
- [24] HEALEY, C. G., BOOTH, K. S., AND ENNS, J. T. High-speed visual estimation using preattentive processing. *ACM Transactions on Computer-Human Interaction* 3, 2 (1996), 107–135.
- [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] HEALEY, C. G., AND ENNS, J. T. Large datasets at a glance: Combining textures and colors in scientific visualization. *IEEE Transactions on Visualization and Computer Graphics* 5, 2 (1999), 145–167.
- [27] HERTZMANN, A. Painterly rendering with curved brush strokes of multiple sizes. In *SIGGRAPH 98 Conference Proceedings* (Orlando, Florida, 1998), M. Cohen, Ed., pp. 453–460.

- [28] HSU, S. C., AND LEE, I. H. H. Drawing and animation using skeltal strokes. In *SIGGRAPH 94 Conference Proceedings* (Orlando, Florida, 1994), A. Glassner, Ed., pp. 109–118.
- [29] 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.
- [30] INTERRANTE, V. Harnessing natural textures for multivariate visualization. *IEEE Computer Graphics & Applications* 20, 6 (2000), 6–11.
- [31] JULÉSZ, B. *Foundations of Cyclopean Perception*. University of Chicago Press, Chicago, Illinois, 1971.
- [32] JULÉSZ, B. A theory of preattentive texture discrimination based on first-order statistics of textons. *Biological Cybernetics* 41 (1981), 131–138.
- [33] JULÉSZ, B. A brief outline of the texton theory of human vision. *Trends in Neuroscience* 7, 2 (1984), 41–45.
- [34] 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.
- [35] 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.
- [36] LAIDLAW, D. H. Loose, artistic “textures” for visualization. *IEEE Computer Graphics & Applications* 21, 2 (2001), 6–9.
- [37] 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.
- [38] LEVKOWITZ, H., AND HERMAN, G. T. Color scales for image data. *IEEE Computer Graphics & Applications* 12, 1 (1992), 72–80.
- [39] LITWINOWICZ, P. Processing images and video for an impressionist effect. In *SIGGRAPH 97 Conference Proceedings* (Los Angeles, California, 1997), T. Whitted, Ed., pp. 407–414.
- [40] 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.
- [41] MACK, A., AND ROCK, I. *Inattentional Blindness*. MIT Press, Menlo Park, California, 1998.
- [42] 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.
- [43] MCCORMICK, B. H., DEFANTI, T. A., AND BROWN, M. D. Visualization in scientific computing—a synopsis. *IEEE Computer Graphics & Applications* 7, 7 (1987), 61–70.
- [44] MEIER, B. J. Painterly rendering for animation. In *SIGGRAPH 96 Conference Proceedings* (New Orleans, Louisiana, 1996), H. Rushmeier, Ed., pp. 477–484.
- [45] NAGY, A. L., AND SANCHEZ, R. R. Critical color differences determined with a visual search task. *Journal of the Optical Society of America A* 7, 7 (1990), 1209–1217.
- [46] NAGY, A. L., SANCHEZ, R. R., AND HUGHES, T. C. Visual search for color differences with foveal and peripheral vision. *Journal of the Optical Society of America A* 7 (1990), 1995–2001.
- [47] NAKAYAMA, K., AND SILVERMAN, G. H. Serial and parallel processing of visual feature conjunctions. *Nature* 320 (1986), 264–265.
- [48] QUINLAN, P. T., AND HUMPHREYS, G. W. Visual search for targets defined by combinations of color, shape, and size: An examination of task constraints on feature and conjunction searches. *Perception & Psychophysics* 41, 5 (1987), 455–472.

- [49] RAO, A. R., AND LOHSE, G. L. Identifying high level features of texture perception. *CVGIP: Graphical Models and Image Processing* 55, 3 (1993), 218–233.
- [50] 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.
- [51] 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.
- [52] RHEINGANS, P., AND TEBBS, B. A tool for dynamic explorations of color mappings. *Computer Graphics* 24, 2 (1990), 145–146.
- [53] 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.
- [54] ROGOWITZ, B. E., AND TREINISH, L. A. An architecture for rule-based visualization. In *Proceedings Visualization '93* (San Jose, California, 1993), pp. 236–243.
- [55] ROOD, O. N. *Modern Chromatics, with Applications to Art and Industry*. Appleton, New York, New York, 1879.
- [56] ROSENBLUM, L. J. Research issues in scientific visualization. *IEEE Computer Graphics & Applications* 14, 2 (1994), 61–85.
- [57] 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.
- [58] SCHAPIRO, M. *Impressionism: Reflections and Perceptions*. George Brazillier, Inc., New York, New York, 1997.
- [59] SCHWEITZER, D. Artificial texturing: An aid to surface visualization. *Computer Graphics (SIGGRAPH 83 Conference Proceedings)* 17, 3 (1983), 23–29.
- [60] SHIRAIISHI, M., AND YAMAGUCHI, Y. Image moment-based stroke placement. In *SIGGRAPH 99 Sketches & Applications* (Los Angeles, California, 1999), R. Kidd, Ed., p. 247.
- [61] SILBERSHATZ, A., STONEBRAKER, M., AND ULLMAN, J. D. The “Lagunita” report of the NSF invitational workshop on the future of database systems research. Tech. Rep. TR-90-22, Department of Computer Science, University of Texas at Austin, 1990.
- [62] SIMON, D. J., AND LEVIN, D. T. Change blindness. *Trends in Cognitive Science* 1 (1997), 261–267.
- [63] SMITH, P. H., AND VAN ROSENDALE, J. Data and visualization corridors report on the 1998 CVD workshop series (sponsored by DOE and NSF). Tech. Rep. CACR-164, Center for Advanced Computing Research, California Institute of Technology, 1998.
- [64] 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.
- [65] STRASSMANN, S. Hairy brushes. *Computer Graphics (SIGGRAPH 86 Proceedings)* 20, 4 (1986), 185–194.
- [66] TAM, R., HEALEY, C. G., AND FLAK, B. Volume visualization of abdominal aortic aneurysms. In *Proceedings Visualization '97* (Phoenix, Arizona, 1997), pp. 43–50.
- [67] 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.
- [68] THOMSON, K. A., INGRAHAM, W. J., HEALEY, M. C., LEBLOND, P. H., GROOT, C., AND HEALEY, C. G. Computer simulations of the influence of ocean currents on Fraser River sockeye salmon (*oncorhynchus nerka*) return times. *Canadian Journal of Fisheries and Aquatic Sciences* 51, 2 (1994), 441–449.

- [69] TRICK, L., AND PYLYSHYN, Z. Why are small and large numbers enumerated differently? A limited capacity preattentive stage in vision. *Psychology Review* 101 (1994), 80–102.
- [70] TRIESMAN, A. Preattentive processing in vision. *Computer Vision, Graphics and Image Processing* 31 (1985), 156–177.
- [71] TRIESMAN, A. Search, similarity, and integration of features between and within dimensions. *Journal of Experimental Psychology: Human Perception & Performance* 17, 3 (1991), 652–676.
- [72] TRIESMAN, A., AND GELADE, G. A feature-integration theory of attention. *Cognitive Psychology* 12 (1980), 97–136.
- [73] TRIESMAN, A., AND GORMICAN, S. Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review* 95, 1 (1988), 15–48.
- [74] TRIESMAN, A., AND SOUTHER, J. Illusory words: The roles of attention and top-down constraints in conjoining letters to form words. *Journal of Experimental Psychology: Human Perception & Performance* 14 (1986), 107–141.
- [75] TURK, G., AND BANKS, D. Image-guided streamline placement. In *SIGGRAPH 96 Conference Proceedings* (New Orleans, Louisiana, 1996), H. Rushmeier, Ed., pp. 453–460.
- [76] VAREY, C. A., MELLERS, B. A., AND BIRNBAUM, M. H. Judgments of proportions. *Journal of Experimental Psychology: Human Perception & Performance* 16, 3 (1990), 613–625.
- [77] VENTURI, L. Impressionist style. In *Impressionism in Perspective*, B. E. White, Ed. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1978, pp. 105–113.
- [78] WARE, C. Color sequences for univariate maps: Theory, experiments, and principles. *IEEE Computer Graphics & Applications* 8, 5 (1988), 41–49.
- [79] WARE, C., AND BEATTY, J. C. Using colour dimensions to display data dimensions. *Human Factors* 30, 2 (1988), 127–142.
- [80] WARE, C., AND KNIGHT, W. Orderable dimensions of visual texture for data display: Orientation, size, and contrast. In *Proceedings SIGCHI '92* (Monterey, California, 1992), pp. 203–209.
- [81] WARE, C., AND KNIGHT, W. Using visual texture for information display. *ACM Transactions on Graphics* 14, 1 (1995), 3–20.
- [82] WINKENBACH, G., AND SALESIN, D. H. Rendering free-form surfaces in pen-and-ink. In *SIGGRAPH 96 Conference Proceedings* (New Orleans, Louisiana, 1996), H. Rushmeier, Ed., pp. 469–476.
- [83] WOLFE, J. M. Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review* 1, 2 (1994), 202–238.
- [84] WOLFE, J. M., FRIEDMAN-HILL, S. R., STEWART, M. I., AND O'CONNELL, K. M. The role of categorization in visual search for orientation. *Journal of Experimental Psychology: Human Perception & Performance* 18, 1 (1992), 34–49.
- [85] 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.