

Perceptually Based Brush Strokes for Nonphotorealistic Visualization

CHRISTOPHER G. HEALEY and LAURA TATEOSIAN

North Carolina State University

and

JAMES T. ENNS and MARK REMPLE

The University of British Columbia

An important problem in the area of computer graphics is the visualization of large, complex information spaces. Datasets of this type have grown rapidly in recent years, both in number and in size. Images of the data stored in these collections must support rapid and accurate exploration and analysis. This article presents a method for constructing visualizations that are both effective and aesthetic. Our approach uses techniques from master paintings and human perception to visualize a multidimensional dataset. Individual data elements are drawn with one or more brush strokes that vary their appearance to represent the element's attribute values. The result is a *nonphotorealistic visualization* of information stored in the dataset. Our research extends existing glyph-based and nonphotorealistic techniques by applying perceptual guidelines to build an effective representation of the underlying data. The nonphotorealistic properties the strokes employ are selected from studies of the history and theory of Impressionist art. We show that these properties are similar to visual features that are detected by the low-level human visual system. This correspondence allows us to manage the strokes to produce perceptually salient visualizations. Psychophysical experiments confirm a strong relationship between the expressive power of our nonphotorealistic properties and previous findings on the use of perceptual color and texture patterns for data display. Results from these studies are used to produce effective nonphotorealistic visualizations. We conclude by applying our techniques to a large, multidimensional weather dataset to demonstrate their viability in a practical, real-world setting.

Categories and Subject Descriptors: H.1.2 [Models and Principles]: User/Machine Systems—*human factors, human information processing*; I.3.3 [Computer Graphics]: Picture/Image Generation—*display algorithms*; I.3.6 [Computer Graphics]: Methodology and Techniques—*interaction techniques*; J.5 [Arts and Humanities]—*fine arts*

General Terms: Experimentation, Human Factors, Performance

Additional Key Words and Phrases: Abstractionism, color, computer graphics, human vision, Impressionism, nonphotorealistic rendering, perception, psychophysics, scientific visualization, texture

1. INTRODUCTION

Visualization is the conversion of collections of strings and numbers (datasets) into images that are used to *explore, discover, validate, and analyze*. The term, “scientific visualization,” originated during an

This work was supported in part by the National Science Foundation (NSF) grant CISE-ACI-0083421, and by the Natural Sciences and Engineering Research Council of Canada.

Authors' addresses: C. G. Healey and L. Tateosian, Department of Computer Science, 173 Venture III Suite 165A, 900 Main Campus Drive #8207, North Carolina State University, Raleigh, NC 27695-8207; email: healey@csc.ncsu.edu; lgtateos@unity.ncsu.edu; J. T. Enns and M. Remple, Department of Psychology, 2136 Main Mall, University of British Columbia, Vancouver, B.C., Canada, V6T 1Z4; email: jenns@psych.ubc.ca; contactmir@hotmail.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.
© 2004 ACM 0730-0301/04/0100-0064 \$5.00

NSF panel on graphics and image processing [McCormick et al. 1987], although the field had a long and rich history prior to this meeting (e.g., cartography, or charts and graphs [MacEachren 1995; Slocum 1998; Tufte 1983, 1990, 1997]). A number of important research problems were identified during these initial discussions. In particular, panelists emphasized the need for ways to manage the overwhelming amount of data being generated. This is not only an issue of the total number of sample points or data elements stored in a dataset (i.e., its size). Each element may also encode multiple values representing multiple independent data attributes (i.e., its dimensionality). The challenge is to design methods to represent even some of this information together in a common display, without overwhelming a viewer's ability to make sense of the resulting images.

A follow-up report on advances in scientific visualization discussed new techniques in important application areas such as volume and flow visualization [Rosenblum 1994]. At the same time, the report noted that much less progress had been made towards application-independent methods for managing and displaying large, multidimensional datasets. Increasing information quality and quantity remains an open problem; this need was again emphasized during a recent DOE/NSF meeting on research directions in visualization [Smith and Van Rosendale 1998].

Work in our laboratories has studied various issues in scientific visualization for much of the last ten years. A large part of this effort has focused on multidimensional visualization, the need to visualize multiple layers of overlapping information simultaneously in a common display. We often divide this problem into two steps:

- (1) The design of a data-feature mapping M , a function that defines visual features (e.g., color, texture, or motion) to represent the data.
- (2) An analysis of a viewer's ability to use the images produced by M to explore and analyze the data.

A multidimensional dataset D represents 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 m visual features $V_j \in V$ to use to display values for each A_j ; it also defines a corresponding $\Phi_j : A_j \rightarrow V_j$ to map the domain of A_j to the range of displayable values in V_j . An effective M must generate images that allow viewers to "see" effortlessly within their data. The need to build fundamental techniques that are appropriate for a wide range of visualization environments further complicates this problem.

The guidelines used to design our M are based on the perceptual abilities of the low-level human visual system. Previous work has documented different methods for harnessing perception during visualization [Bergman et al. 1995; Grinstein et al. 1989; Healey 1996; Healey et al. 1996; Healey and Enns 1999; Rheingans and Tebbs 1990; Rogowitz and Treinish 1993; Ware 1988; 2000; Ware and Knight 1995; Weigle et al. 2000]. Certain visual features are detected very quickly by the visual system [Egeth and Yantis 1997; Mack and Rock 1998; Pomerantz and Pristach 1989; Rensink 2000; Simons 2000; Triesman 1985; Triesman and Gormican 1988; Wolfe 1994; Wolfe et al. 2000]; when combined properly, these same features can be used to construct multidimensional displays that support rapid, accurate, and effortless exploration and analysis. For example, properties of color and texture (e.g., luminance, hue, contrast, or regularity) are often used to represent different attributes in a dataset. The way that color and texture are mapped to the data attributes is controlled using results from psychophysical studies of our ability to distinguish between different color and texture patterns. 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 [Smith and Van Rosendale 1998].

More recently, we have initiated a new set of investigations that focus on the question: "Can we make our visualizations aesthetically pleasing?" The way an image initially attracts a viewer's attention is different from how it holds a viewer's interest over time. Cognitive scientists use the terms "orienting" and "engaging" to describe the distinct psychophysical processes involved in these two aspects of

attention [Coren et al. 2003]. Display techniques that invoke these responses could be used to direct attention to important properties in a visualization, and to then encourage the visual system to perform a more detailed analysis within these areas. The idea of building artistically motivated visualizations was also based on nonphotorealistic rendering algorithms in computer graphics [Curtis et al. 1997; Haberli 1990; Hertzmann 1998; Hsu and Lee 1994; Litwinowicz 1997; Meier 1996; Strassmann 1986], and by the efforts of researchers such as Interrante [Interrante 2000], Laidlaw [Kirby et al. 1999; Laidlaw 2001; Laidlaw et al. 1998], and Ebert and Rheingans [Ebert and Rheingans 2000; Rheingans and Ebert 2001] to extend this work to a visualization environment. Nonphotorealistic techniques represent a promising method to both orient and engage a viewer's attention within an image.

Certain movements in painting (e.g., Impressionism, Abstractionism, or watercolor) are characterized by a set of fundamental styles [Brown 1978; Schapiro 1997; Venturi 1978]. If the basic brush stroke properties embodied in these styles can be identified and simulated on a computer, we believe they can then be used to represent individual data attributes in a multidimensional dataset. Our goal is an image that looks like a painting, not of a real-world scene, but rather of the information contained in the dataset.

Such a technique might initially seem difficult to control and test. An important insight is that many brush stroke properties correspond closely to perceptual features that are detected by our visual system. In some sense this is not surprising. Artistic masters understood intuitively which properties of a painting would orient a viewer's gaze and engage their thoughts. We believe this overlap can act as a bridge between artistic styles and low-level vision, allowing us to apply our knowledge of perception to predict how nonphotorealistic techniques will perform in a visualization environment. In addition, psychophysical experiments offer a controlled method for studying the fundamental strengths and limitations of a given stroke property, both in isolation and in combination with other properties shown together in the same display. In order to use the correspondence between painting and perception during multidimensional visualization, we need to show that our perceptual guidelines extend to a nonphotorealistic domain. Perceptually salient displays will guarantee an *effective* presentation of information.

We begin this article with a brief overview of nonphotorealistic rendering, followed by a description of painting styles in Impressionism and their correspondence to perceptual features in vision. We continue with an explanation of the guidelines that are used to construct perceptually salient brush strokes. We next discuss a set of experiments that test the expressiveness of our nonphotorealistic properties to confirm that their abilities match the perceptual rules on which they are based. Finally, we describe a visualization system built from our experimental findings, and demonstrate its use for exploring a collection of multidimensional weather datasets.

2. NONPHOTOREALISTIC RENDERING

For many years, researchers in 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 such as 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 certain situations, these *nonphotorealistic renderings* are often considered more effective, more appropriate, or even more expressive than an equivalent photograph [Gooch and Gooch 2001; Strothotte and Schlechtweg 2002] (see Figure 1).

Different methods have been suggested to simulate different artistic styles. For example, researchers from the University of Washington presented a collection of techniques for generating pen-and-ink



Fig. 1. A simple nonphotorealistic rendering of a collection of water lilies; the original image is shown in the upper-left corner.

sketches. Their initial work focused on using multiresolution curves [Finkelstein and Salesin 1994] to build a *stroke texture*, a prioritized collection of simulated pen strokes that are drawn to create stroke patches with a specific texture and tone. The stroke textures are used to construct nonphotorealistic renderings of 3D polygonal models [Winkenbach and Salesin 1994], and to interactively convert greyscale reference images into pen-and-ink sketches [Salisbury et al. 1994]. Follow-on work extended the stroke textures to parametric surfaces [Winkenbach and Salesin 1996], allowed the definition of directional fields to control orientation during the sketching of 2D reference images [Salisbury et al. 1997], and discussed ways to guarantee a constant tone that is independent of scale and display resolution [Salisbury et al. 1996]. Related work by Takagi et al. used a voxel-based simulation of the physical properties of paper and colored lead pencils to construct color pencil drawings [Takagi and Fujishiro 1997; Takagi et al. 1999]. Finally, Sousa and Buchanan [1999a, 1999b, 2000] built a sophisticated simulation of graphite pencils, paper, blenders, and erasers to produce nonphotorealistic pencil drawings of 3D geometric models; their technique allows for the variation of numerous parameters such as pressure on the pencil, the shape of its tip, how the pencil is held by the artist, and how the pencil and paper interact.

Texture synthesis techniques have been proposed by a number of researchers to generate nonphotorealistic results. Initial work by Lewis [1984] allowed viewers to interactively paint spectral properties of a texture patch in the frequency domain. Convolution and an inverse Fourier transform were applied to generate a randomized spatial version of the texture. This result was “painted” onto a canvas in different ways using different pixel copy operations. Haberli and Segal [1993] showed how texture mapping can be used to produce a number of fundamental drawing primitives, including air brush and painted effects. More recently, Hertzmann et al. [2001] constructed feature analogies to automatically generate nonphotorealistic results. Texture synthesis techniques were applied to two images I and I' to learn how features in I map to corresponding features in I' (e.g., I could be a photograph and I' a painterly rendition of the photograph). Hertzmann used the resulting *image analogy* to automatically generate a nonphotorealistic image J' from a new source image J . The image analogy embodies J' with style properties similar to those seen in I' .

Our interests lie mainly in nonphotorealistic techniques that use simulated brush strokes to produce images that look like paintings. An early example of this idea was proposed by Strassmann [1986]; he constructed a “hairy brush”, a collection of bristles placed along a line segment. Japanese-style sumi brush strokes were produced by applying ink to the brush, then moving it along a path over a simulated paper surface. Later work by Haberli [1990] allowed users to paint by stroking a brush over an underlying target image. The size, shape, color, location, and direction of brush strokes were varied to produce different representations of the target. Hsu and Lee [1994] defined a reference backbone and reference thickness for a base texture, then warped these properties parametrically to produce line art images. This generated expressive strokes with complex paths of varying thickness. Litwinowicz [1997] clipped simple strokes to object boundaries in a reference image, then rendered the strokes as lines and texture maps with variable length, thickness, direction, and color. A stroke's properties were selected based on the image properties of the object it represented. Curtis et al. [1997] built a fluid-flow simulation to model the interactions of brush, pigment, and paper during the painting of watercolor images. Their system produced subtle watercolor effects such as dry-brush, backruns, and color glazing. Shiraishi and Yamaguchi [1999] computed image moments on a target image; these values controlled the color, location, orientation, and size of texture-mapped brush strokes in a painterly rendering of the target. Hertzmann [1998] decomposed a reference image into a level-of-detail hierarchy by applying Gaussian kernels with increasing support. This generated a collection of reference images, each with different amounts of blurring. The images were painted using strokes with a radius proportional to the kernel size. Each stroke was modeled as a spline that varied in its length, size, opacity, placement, and color jitter. The individual paintings were composited to produce a nonphotorealistic result. Meier [1996] addressed the goal of animating nonphotorealistic renderings by attaching particles to surfaces in a 3D geometric scene, then drawing a brush stroke with scene-controlled color, size, and direction at each particle position. Information stored within a particle ensured a consistent stroke appearance. This produced a smooth animation free of the visual artifacts that result from inconsistent variations in the appearance of a stroke across multiple frames. Gooch et al. [2002] segmented an image into closed regions representing image features; the medial axis of a region defined the locations and directions of brush strokes that were used to paint the region. Finally, Hertzmann [2002] proposed a method for simulating the appearance of lighting on the brush strokes in a painting. A height field was associated with each brush stroke, producing a global height map as the strokes were painted; the height map was then used to bump-map the painting with a Phong shading model.

More recently, researchers in scientific visualization have started to investigate how techniques from nonphotorealistic rendering might be used to improve the expressiveness of a data display. Laidlaw extended the layered approach of Meier to visualize multidimensional data in a painterly fashion [Kirby et al. 1999; Laidlaw 2001; Laidlaw et al. 1998]. He varied style properties such as underpainting

lightness and stroke size, transparency, direction, saturation, and frequency to display multiple layers of information in a single nonphotorealistic image. Interrante [2000] discussed constructing natural textures to visualize multidimensional data. Finally, Ebert and Rheingans used nonphotorealistic techniques such as silhouettes, sketch lines, and halos to highlight important features in a volumetric dataset [Ebert and Rheingans 2000; Rheingans and Ebert 2001]. More recent work applied stipple drawing techniques to interactively preview scientific and medical volumes [Lu et al. 2002].

Nonphotorealistic rendering produces images that are expressive by making use of a wide range of painting styles. Promising results from scientific visualization show that these ideas can be extended to the problem of representing information. This suggests that it may be possible to construct flexible brush stroke glyphs to visualize multidimensional data elements. To do this properly, however, we must ensure our brush strokes will form nonphotorealistic visualizations that are *effective* in their ability to represent multidimensional data. The use of nonphotorealistic techniques also holds promise for constructing visualizations that are seen as *aesthetic* or beautiful by our viewers.

Our investigations focus on understanding and controlling the expressive abilities of different nonphotorealistic brush stroke properties during visualization. These properties can then be used to produce nonphotorealistic images that are both effective and engaging.

3. PAINTING STYLES

The 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 nonphotorealistic properties 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 match brush stroke characteristics from the Impressionist painting style with corresponding visual features that have been shown to be effective in a perceptual visualization environment. There are no technical reasons for our choice of Impressionism over any other movement. In fact, we expect the basic theories behind our technique will extend 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 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 [Brown 1978; Schapiro 1997; Venturi 1978], 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) *Solicit a viewer’s optics.* Study the retinal system; divide tones as separate brush strokes to vitalize color rather than greying with overlapping strokes; harness simultaneous contrast; use models from color scientists such as Chevreul [1967] or Rood [1879].
- (4) *Minimize perspective.* Perspective is shortened and distance reduced to turn 3D space into a 2D image.
- (5) *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.

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 (e.g., Seurat's use of scientific models of color [Chevreul 1967; Hering 1964; Rood 1879]). Rather than building 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 features. Additional properties can be identified by studying the paintings themselves. These properties 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,
- Coverage*. The amount of canvas or underpainting that shows through the foreground strokes,
- 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 stroke boundaries, and produce ridges of paint that cause underhanging shadows when lit from the proper direction.

Figure 2 shows a close-up view of an oil painting that demonstrates different brush stroke properties such as color, path, size, and density. Although by no means exhaustive, this collection of features provides a good starting point for our work. All of the stroke properties 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.

4. PERCEPTUAL PROPERTIES

Recent research in visualization has explored ways to apply rules of perception to produce images that are visually salient [Ware 2000]. This work is based in large part on psychophysical studies of the low-level human visual system. One of the most important lessons of the past twenty-five years is that human vision does not resemble the relatively faithful and largely passive process of modern photography [Pomerantz and Pristach 1989; Triesman 1985; Triesman and Gormican 1988; Wolfe 1994; Wolfe et al. 2000]. The goal of human vision is not to create a replica or image of the seen world in our heads. A much better metaphor for vision is that of a dynamic and ongoing construction project, where the products being built are short-lived models of the external world that are specifically designed for the current visually guided tasks of the viewer [Egeth and Yantis 1997; Mack and Rock 1998; Rensink 2000; Simons 2000]. There does not appear to be any general purpose vision. What we "see" when confronted with a new scene depends as much on our goals and expectations as it does on the array of light that enters our eyes. Among the research findings responsible for this altered view of "seeing" is a greater appreciation of:

- (1) Detailed form and color vision is only possible for a tiny window of several degrees of arc surrounding the current gaze location. "Seeing" beyond the single glance therefore requires a time-consuming series of eye movements.



Fig. 2. A close-up view of a small section of an oil painting that demonstrates various stroke properties such as color, path, size, and density; the entire painting is shown in the upper-left corner.

- (2) The eye movements that are needed to process a “whole scene,” such as the 180° view we often assume we have, are discrete. Many of them must be made in order to see the detail in a large scene, and almost no visual information is gained during an eye movement itself.
- (3) Memory for information from one glance to the next is extremely limited. At most, the details from only three or four objects can be monitored between glances; perception is often limited to only one object at a time. What we see therefore depends critically on which objects in a scene we are looking for and attending to.
- (4) Human vision is designed to capitalize on the assumption that the world is generally a quiet place. Only differences need to be registered. Objects that are very different from their surroundings, or that change or move, draw attention to themselves because of the difference signals that emanate from these visual field locations.
- (5) The basic visual features that can be used to guide attention are not large in number. They include differences in the first order properties of luminance and hue, and the second-order properties of orientation, texture, and motion. Effective third-order properties are limited to some very simple characteristics of shape such as length, area, and convexity.

The reality of each of these findings can be illustrated through the so-called *change blindness* which affects us all [Mack and Rock 1998; Rensink 2000; Simons 2000]. It involves a task similar to a game that has amused children reading the comic strips for many years. Try to find the difference between the two pictures in Figures 3 and 4. Many viewers have a difficult time seeing any difference and often have to be coached to look carefully to find it. Once they have discovered it, they realize that the difference was



Fig. 3. An example of change blindness, the inability to quickly identify significant differences across separate views of a common scene; try to identify the difference between this photograph and the photograph shown in Figure 4 (the answer is included in footnote 1 on the next page).

not a subtle one. Change blindness is not a failure to see because of limited visual acuity; rather, it is a failure based on inappropriate attentional guidance. Some parts of the eye and the brain are clearly responding differently to the two pictures. Yet, this does not become part of our visual experience until attention is focused directly on the objects that vary.

Harnessing human vision for the purposes of data visualization therefore requires that the images themselves be constructed so as to draw attention to their important parts. Since the displays being shown are typically novel, we cannot rely on the expectations that might accompany the viewing of a familiar scene. Rather, we must build an effective mapping between data values and visual features, so that differences in the features draw the eyes, and more importantly the mind, on their own. Attracting the viewer's gaze to a particular object or location in a scene is the first step in having the viewer form a mental representation that may persist over subsequent scenes.

A data-feature mapping that builds on our knowledge of perception can support the exploration and analysis of large amounts of data in a relatively short period of time. The ability to take advantage of the low-level visual system is especially attractive, since:

- completion of high-level exploration and analysis tasks (e.g., target search, boundary and region identification, estimation, or spatial and temporal tracking) is rapid and accurate, usually requiring an exposure duration of 200 milliseconds or less,
- analysis is display size insensitive, so the time required to complete a task is independent of the number of elements in the display, and
- different features can interact with one another to mask information; psychophysical experiments allow us to identify and avoid these visual interference patterns.



Fig. 4. An example of change blindness, the inability to quickly identify significant differences across separate views of a common scene; try to identify the difference between this photograph and the photograph shown in Figure 3 (the answer is included in footnote 1 below).

Our most recent research has focused on the combined use of the fundamental properties of color (hue and luminance) and texture (size, contrast, orientation, and regularity) to encode multiple attributes in a single display [Healey 1996; Healey and Enns 1998, 1999]. A comparison of perceptual color and texture properties with painting 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 brush stroke properties (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 nonphotorealistic properties 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 knowledge to draw from to predict how we expect our brush stroke properties to react in a multidimensional visualization environment.

We applied three specific areas of research in perception and visualization to guide the use of properties of our nonphotorealistic brush strokes: color selection, texture selection, and feature hierarchies that cause visual interference and masking.

4.1 Color Selection

Color is a common feature used in many visualization designs. Examples of simple color scales include the rainbow spectrum, red-blue or red-green ramps, and the grey-red saturation scale [Ware 1988]. More sophisticated techniques attempt to control the difference viewers perceive between different

¹Hint: Look at the bushes immediately behind the back of the Sphinx.

colors, as opposed to the distance between their positions in RGB space. This improvement allows:

- perceptual balance*: a unit step anywhere along the color scale produces a perceptually uniform difference in color,
- distinguishability*: within a discrete collection of colors, every color is equally distinguishable from all the others (i.e., no specific color is “easier” or “harder” to identify), and
- flexibility*: colors can be selected from any part of color space (e.g., the selection technique is not restricted to only greens, or only reds and blues).

Color models such as CIE LUV, CIE Lab, or Munsell can be used to provide a rough measure of perceptual balance [Birren 1969; CIE 1978; Munsell 1905]. Within these models, Euclidean distance is used to estimate perceived color difference. More complex techniques refine this basic idea. Rheingans and Tebbs [1990] plotted a path through a perceptually balanced color model, then asked viewers to define how attribute values map to positions along the path. Non-linear mappings emphasize differences in specific parts of an attribute's domain (e.g., in the lower end with a logarithmic mapping, or in the higher end with an exponential mapping). Other researchers have constructed rules to automatically select a colormap for a target data attribute [Bergman et al. 1995; Rogowitz and Treinish 1993]. Properties of the attribute such as its spatial frequency, its continuous or discrete nature, and the type of analysis to be performed are used to choose an appropriate color representation. Ware constructed a color scale that spirals up around the luminance axis to maintain a uniform simultaneous contrast error along its length [Ware 1988]. His solution matched or outperformed traditional color scales for metric and form identification tasks. Healey and Enns showed that color distance, linear separation, and color category must all be controlled to select discrete collections of equally distinguishable colors [Healey 1996; Healey and Enns 1999].

Our color selection technique combines different aspects of each of these methods. A single loop spiraling up around the L -axis (the luminance pole) is plotted near the boundary of our monitor's gamut of displayable colors in CIE LUV space. The path is subdivided into r named color regions (i.e., a blue region, a green region, and so on). n colors can then be selected by choosing $\frac{n}{r}$ colors uniformly spaced along each of the r color regions. The result is a set of colors selected from a perceptually balanced color model, each with a roughly constant simultaneous contrast error, and chosen such that color distance and linear separation are constant within each named color region.

4.2 Texture Selection

Texture is often viewed as a single visual feature. Like color, however, it can be decomposed into a collection of fundamental perceptual dimensions. Researchers in computer vision have used properties such as regularity, directionality, contrast, size, and coarseness to perform automatic texture segmentation and classification [Haralick et al. 1973; Rao and Lohse 1993a, 1993b; Tamura et al. 1978]. These texture features were derived both from statistical analysis, and through experimental study. Results from psychophysics have shown that many of these properties are also detected by the low-level visual system, although not always in ways that are identical to computer-based algorithms [Aks and Enns 1996; Cutting and Millard 1984; Júlész 1975, 1984; Júlész et al. 1973, 1978; Snowden 1998; Triesman 1991; Wolfe 1994].

One promising approach in visualization has been to use perceptual texture dimensions to represent multiple data attributes. Individual values of an attribute control its corresponding texture dimension. The result is a texture pattern that changes its visual appearance based on data in the underlying dataset. Grinstein et al. [1989] visualized multidimensional data with “stick-man” icons whose limbs encode attribute values stored in a data element; when the stick-men are arrayed across a display, they form texture patterns whose spatial groupings and boundaries identify attribute correspondence. Ware




Fig. 5. Examples of target detection, color targets with constant orientation (top) and random orientation (bottom); (a) orange target in pink strokes, constant 45° background orientation; (b) green target in orange strokes, constant 60° background orientation; (c) green target in orange strokes, constant 45° background orientation; (d) orange target in pink strokes, random 45° and 60° background orientation; (e) green target in orange strokes, random 30° and 45° background orientation; (f) green target in orange strokes, random 45° and 60° background orientation.

and Knight [1995] designed Gabor filters that modified their orientation, size, and contrast based on the values of three independent data attributes. Healey and Enns [1998, 1999] constructed perceptual texture elements (or pexels) that varied in size, density, and regularity; results showed that size and density are perceptually salient, but variations in regularity are much more difficult to identify. More recent work found that orientation can also be used to encode information [Weigle et al. 2000]; a difference of 15° is sufficient to rapidly distinguish elements from one another.

We designed brush strokes that can vary in their area, orientation, spatial density, and regularity (in addition to color). These texture dimensions correspond closely to the nonphotorealistic properties size, direction, coverage, and placement. The displays in Figures 5 and 6, used to measure target detection performance, show examples of each of these properties.




Fig. 6. Examples of target detection, orientation targets with constant color (top) and random color (bottom); (a) 45° target in 30° strokes, constant pink background color; (b) 60° target in 45° strokes, constant green background color; (c) 45° target in 30° strokes, constant pink background color; (d) 45° target in 30° strokes, random pink and orange background color; (e) 60° target in 45° strokes, random pink and orange background color; (f) 45° target in 30° strokes, random orange and green background color.

4.3 Feature Hierarchy

A third issue that must be considered is visual interference. This occurs when the presence of one feature masks another. Although the need to measure a brush stroke's perceptual strength is not necessary in a painting, this information is critical for effective visualization design. The most important attributes (as defined by the viewer) should be displayed using the most salient features. Secondary data should never be visualized in a way that masks the information a viewer wants to see.

Certain perceptual features are ordered in a hierarchy by the low-level visual system. Results reported in both the psychophysical and visualization literature have confirmed a luminance–hue–texture interference pattern. Variations in luminance interfere with a viewer's ability to identify the presence

of individual hues and the spatial patterns they form [Callaghan 1990]. If luminance is held constant across the display, these same hue patterns are immediately visible. The interference is asymmetric: random variations in hue have no effect on a viewer's ability to see luminance patterns. A similar hue on texture interference has also been shown to exist [Healey and Enns 1998, 1999; Snowden 1998; Triesman 1985]; random variations in hue interfere with the identification of texture patterns, but not vice-versa.

Figure 5 shows examples of hue on orientation interference. The upper three displays use a constant background orientation (Figures 5(a)–(c)), while the lower three vary orientation randomly from stroke to stroke (Figures 5(d)–(f)). This has no effect on a viewer's ability to locate a target group by defined color; identification is rapid and accurate for both sets of displays. In Figure 6, the mapping is reversed: background color is held constant in the upper three displays (Figures 6(a)–(c)), and varied randomly in the lower three (Figures 6(d)–(f)). Locating a target group of strokes rotated counterclockwise from their neighbors is much harder when color varies randomly, compared to the displays where color is held constant. What the visual system “sees” initially is a random color pattern. Only with additional exposure time will differences in orientation be reported. Feature interference results suggest that luminance, then hue, then various texture properties should be used to display attributes in order of importance. Real-world evidence has confirmed that this technique works well in practice.

4.4 Orienting Versus Engaging Attention

We are interested in two properties of a nonphotorealistic visualization: its effectiveness and its aesthetic merit. These properties correspond to two basic aspects of human attention: *orienting* and *engaging* [Coren et al. 2003]. Orienting attention to a specific location in an image occurs when the location contains an abrupt transition in a visual feature that is processed by the low-level visual system (e.g., a high-contrast luminance edge, a brief flicker, or a motion discontinuity). This may include redirecting a viewer's gaze so that the foveal center of the eye is aimed at the region of interest, although this is not required. Visual processes can operate selectively on areas of high visual salience through a process called “covert orienting” [Posner and Raichle 1994]. Rapidly orienting a viewer's attention to novel or important areas in a visualization is the first step towards allowing the viewer to efficiently discover, explore, and analyze within their data.

The process of orienting is different from engaging attention in two important ways. First, while orienting is often a momentary event based largely on the nature of an image, engaging reflects the conscious intention of the viewer to search for specific information. For example, engaging is the process that allows the search for a difficult-to-find target to continue, even when no low-level visual evidence exists to orient the visual system to the target's location. Second, different neurological foundations are believed to control the two aspects of attention. Orienting is governed by the older, sub-cortical visual pathways. Engaging is determined by a network of cortical regions that are in close communication with the frontal lobes, the so-called “central executive” of the human brain [Posner and Raichle 1994].

Skilled visual artists are adept at exploiting these complimentary aspects of visual attention, even though they may do so intuitively, without understanding the underlying neural processes [Zeki 1999]. For example, masters of the human portrait such as Vermeer, Titian, and Rembrandt painted the faces of people such that the region of greatest detail and finest spatial resolution was the face itself. Properties of the background and the model's clothing are often presented in shadow or rendered with much less resolution and contrast. This has the effect of drawing the viewer's eye towards the face, which is the center of interest in the portrait. At the same time, these artists reserved another small region away from the face for the most extreme contrast. This was often the collar of the model, a piece of jewelry, or a background surface detail. This localized region of high contrast “pulls” at the viewer's

orienting system, even as the viewer tries to engage their attention on the portrait's face. It has recently been proposed that this interaction between orienting and engaging underlies our fascination with and artistic appreciation of these works [Ramachandran and Hirstein 1999; Zeki 1999]. Psychologists believe they may soon understand the neural substrate of this aspect of creative tension, an idea that is usually thought to be highly abstract.

We believe that orienting and engaging are both important to a successful visualization. Orienting allows us to highlight important regions in an image by capturing the viewer's focus of attention. Engaging encourages the visual system to continue to study the details of an image after orienting occurs. We are pursuing nonphotorealistic visualizations as a promising way to build images with exactly these characteristics. Orienting occurs through the careful use of visual features that are rapidly detected by the low-level visual system. Engaging is achieved by constructing visualizations that are perceived to be beautiful or artistic by the viewer. The studies described in this article represent our initial steps towards investigating different aspects of attention in the context of our nonphotorealistic visualization techniques.

5. EFFECTIVENESS STUDIES

The first question we wanted to answer is whether guidelines on the use of perception in glyph-based visualizations will extend to our nonphotorealistic domain. We conducted a set of psychophysical experiments to test this hypothesis. Our experiments were designed to investigate an observer's ability to rapidly and accurately identify target brush strokes defined by a particular color or orientation [Liu et al. 2003]. Observers were asked to determine whether a small, 3×3 group of strokes with a particular visual feature was present or absent in a display (e.g., a group of orange strokes, as in Figures 5(a), 5(c), 5(d), and 5(f), or a group strokes tilted 60° in Figures 6(b) and 6(e)). Background orientation, color, regularity, and density varied between displays. This allowed us to test for single-glance task performance, and for visual interference effects. Since observers need at least 200 milliseconds to initiate an eye movement, any task performed in 200 milliseconds or less is completed based on "a single glance" at the image. In all cases, observer accuracy and response times were recorded to measure performance. The experimental results were then used to identify similarities and differences between nonphotorealistic images and existing perceptual visualization techniques.

5.1 Design

Each experimental display contained a 22×22 array of simulated brush strokes (Figures 5 and 6). The color of the displays was calibrated to the monitor to ensure accurate reproduction. Observers were asked to determine whether a 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 observers to enter their answer: "target present" or "target absent." Observers were told to respond as quickly as possible, while still maintaining a high rate of accuracy. Feedback was provided after each display: a "+" sign if an observer's answer was correct, or a "-" sign if it was not.

The displays were equally divided into two groups: one studied an observer'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 single-glance performance and visual interference. For example, the following experimental conditions were used to investigate an observer's ability to identify colored strokes:

—*Two target-background color pairings.* An orange target in a pink background, or a green target in an orange background; this allowed us to test for generality in observer performance for different target-background color pairings,

- Two background orientations.* Constant (every stroke is oriented in the same direction, either 30° or 60°), or random (strokes are randomly oriented 30° and 45° , or 45° and 60°); any decrease in performance from a constant to a random background would indicate visual interference from orientation during the search for color targets,
- Three background densities.* The size of the strokes in the display were varied to produce sparse, dense, or very dense patterns; this allowed us to see how changes in density affected target identification, and
- Two background regularities.* Strokes were arrayed in a regular grid pattern, or jittered randomly across the display; this allowed us to test for visual interference caused by spatial irregularity in the global texture.

Our experimental conditions produced 24 different color display types (two target-background color pairings by two background orientations by three background densities by two background regularities). Observers were asked to view eight variations of each display type, for a total of 192 color trials. For each display type, half the trials were randomly chosen to contain a group of target strokes; the other half did not.

Examples of six color displays are shown in Figure 5. Each display contains either an orange target in a sea of pink strokes (Figures 5(a), 5(c), 5(d), and 5(f)), or a green target in a sea of orange strokes (Figures 5(b) and 5(e)). In the upper three displays, the background orientation of the strokes is constant (either 45° or 60°). The coverage is dense in Figure 5(a), sparse in Figure 5(b), and very dense in Figure 5(c). The strokes are arrayed in a regular pattern in Figure 5a, and randomly jittered in Figures 5(b) and 5(c). The lower three displays are identical, except for the background orientation. In Figure 5(d), half the strokes were randomly selected to be oriented 45° ; the other half are oriented 60° . In Figures 5(e) and 5(f), half the strokes are oriented 30° , and half are oriented 45° .

The displays that studied orientation were designed in an identical fashion. Two target-background orientation pairings were tested: target strokes oriented 45° in a sea of background strokes oriented 30° , or 60° targets in a 45° background. Two different color patterns were used to search for color on orientation interference: constant (every stroke has the same color, either green or pink), or random (strokes are randomly colored green and orange, or orange and pink). Background densities and regularities are identical to the color displays. As before, eight variations of each display type were shown for a total of 192 orientation trials.

Figures 6(a), 6(c), 6(d), and 6(f) show examples of 45° target strokes in a sea of 30° background elements. Figures 6(b) and 6(e) show a 60° target in a 45° background. The upper three displays have a constant background color (either pink or green). The strokes are densely packed and regularly positioned in Figure 6(a), sparsely packed and randomly jittered in Figure 6(b), and very densely packed and randomly jittered in Figure 6(c). The lower three displays are identical, except for the background color of the strokes. In Figures 6(d) and 6(e) half the strokes were randomly selected to be colored pink; the other half are colored orange. In Figure 6(f), half are colored orange, and half are colored green.

The colors, orientations, densities, and regularities we used were chosen based on results from previous experiments [Healey and Enns 1998; 1999; Weigle et al. 2000]. In particular, the colors and orientations we selected were shown to be rapidly distinguishable from one another when displayed in isolation (i.e., without variations in irrelevant background dimensions).

Eighteen observers (six males and twelve females ranging in age from 18 to 28) with normal or corrected acuity and normal color vision participated during our studies. The observers were undergraduate and graduate student volunteers, none of whom had any prior experience with scientific visualization. Every observer completed both the color and the orientation experiments within our minimum accuracy requirements of 60% or better for each target type. Observers were told before an

experiment that half the trials would contain a target, and half would not. Observers completed a practice session with 24 trials before each experiment (i.e., color practice trials before the color experiment, and orientation practice trials before the orientation experiment). Observers were counterbalanced: half started with the color experiment, while the other half started with the orientation experiment. We used a Macintosh computer with a 24-bit color display to run our studies. Answers (either “target present” or “target absent”) and response times for each trial an observer completed were recorded for later analysis.

5.2 Results

Each observer response collected during our experiments was classified by condition: target-background pairing, primary background type (either constant or random), density, regularity, and target present or absent. Trials with the same conditions were collapsed to produce an average accuracy a and an average response time t . We used these values to compute a measure of *search inefficiency* for each observer in each condition $e = \frac{t}{a}$; this is a common measurement for situations where the direction of change in accuracy and response time is the same in each experimental condition. If observer responses are perfect (i.e., $a = 1.0$), inefficiency e equals response time; as accuracy decreases, inefficiency e increases (i.e., search inefficiency increases both for longer response times and for increased error rates). Results were tested for significance with a multifactor analysis of variance (ANOVA). We used a standard 95% confidence interval to denote significant variation in mean inefficiency values.

We first conducted preliminary ANOVAs examining all possible factors, separately for accuracy data a and response time t . These analyses indicated that: (1) some factors were not significantly related to our measures of performance (specifically, target presence-absence and target-background pairing), and (2) a and t were highly correlated. Our primary analyses were therefore based on the search inefficiency measure e and the significant factors of target type (color or orientation), primary background (constant or random), density (sparse, dense, or very dense), and regularity (regular or irregular). In summary, our results showed:

- (1) Color targets were easy to detect at our 200 millisecond single-glance exposure duration (mean accuracy $a = 91.1\%$ and mean inefficiency $e = 811.9$ over all experimental conditions); a random orientation pattern had no interfering effect on performance.
- (2) Orientation targets were easy to detect when a constant color was displayed in the background ($a = 71.9\%$ and $e = 1327.7$ for constant color trials); a random background color pattern caused a significant reduction in performance ($a = 67.9\%$ and $e = 1437.8$ for random color trials).
- (3) Background density had a significant effect on both color and orientation targets; denser displays produced an improvement in performance.
- (4) Background regularity had a significant effect on both color and orientation targets; irregular displays caused a reduction in performance.

Color targets were easy to identify, moreover, a random variation in background orientation had no effect on performance ($F(1, 17) = 0.01$, $p < 0.94$ with $e = 813.7$, $a = 91.2\%$ for constant orientation, and $e = 810.2$, $a = 90.9\%$ for random orientation). Orientation targets were easy to identify in a constant color background, although performance was not as good as for color targets ($e = 1327.7$, $a = 71.9\%$). A random color pattern produced a significant reduction in performance ($F(1, 17) = 8.08$, $p < 0.05$, with $e = 1437.8$, $a = 67.9\%$).

Variation in background density had a significant effect on performance, both for color targets ($F(2, 34) = 30.84$, $p < 0.001$) and for orientation targets ($F(2, 34) = 7.85$, $p < 0.01$). In all cases accuracy and inefficiency were best for very dense packings ($e = 708.1$, $a = 96.9\%$ for very dense

color trials; $e = 1245.3$, $a = 75.2\%$ for very dense orientation trials), and worst for sparse packings ($e = 953.9$, $a = 83.0\%$ for sparse color trials; $e = 1511.2$, $a = 66.0\%$ for sparse orientation trials).

Variation in background regularity also had a significant effect on performance, both for color targets ($F(1, 17) = 5.10$, $p < 0.04$) and for orientation targets ($F(1, 17) = 24.89$, $p < 0.001$). In all cases accuracy and inefficiency were best for regular trials ($e = 787.5$, $a = 92.5\%$ for regular color trials; $e = 1235.7$, $a = 75.6\%$ for regular orientation trials), and worst for irregular trials ($e = 834.7$, $a = 89.7\%$ for irregular color trials; $e = 1523.7$, $a = 64.1\%$ for irregular orientation trials).

Finally, we observed a density $\times$ regularity interaction for color trials ($F(2, 34) = 5.34$, $p < 0.01$): variations in performance were larger for “harder” trials. For example, the effect of irregularity was larger for sparse color trials, compared to very dense color trials; the effect of density was larger for irregular trials, compared to regular trials. The same interaction pattern was seen for the orientation trials, but the effect was only marginally significant ($F(2, 34) = 2.93$, $p < 0.07$).

5.3 Interpretation

Our results match previous findings in both the psychophysical and the visualization literature, specifically: (1) color produces better performance than orientation during target identification ($F(1, 17) = 71.51$, $p < 0.001$ for our experiments), 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 reported in experimental [Callaghan 1990; Snowden 1998] and real-world visualization settings [Healey and Enns 1998, 1999]. Our results extend the work of Healey and Enns [1999], who found a general color on texture interference pattern, but no corresponding texture on color effect. This provides positive evidence to support the belief that perceptual findings will carry to a nonphotorealistic visualization environment.

The improvement in performance when density increased, both for color and orientation targets, was encouraging. An initial concern we discussed was that texture variations (e.g., orientation differences) would disappear when density increased and background color was held constant. Our results show that, for the types of strokes we displayed, different orientations are not “lost” in the background, even when a significant stroke overlap exists. This supports our goal of producing painterly images that contain dense stroke regions, yet at the same time allow viewers to rapidly identify variations in the underlying texture properties.

Finally, the reduction in performance when strokes were irregularly positioned was intriguing. We concluded that regularity acts as a reinforcing visual cue, helping observers identify targets based on some other feature (e.g., color or orientation). The presence of a target patch “breaks” the regularity pattern, providing an additional visual signal that identifies the presence of a target. Jittering the strokes removes this background support. In this sense, irregularity is not so much an “interfering” effect as it is the loss of a secondary feature that helps to highlight the presence of a group of target strokes.

6. NONPHOTOREALISTIC VISUALIZATION

Based on the results from our experiments, we built a nonphotorealistic visualization system that varied brush stroke color, orientation, coverage (i.e., spatial density), and size to encode up to four 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, coverage, size, and orientation in order of attribute importance (from most important to least important).

6.1 Painting Algorithm

To produce nonphotorealistic visualizations, we must convert a dataset D into a nonphotorealistic image using a data-feature mapping M . We wanted to design a technique that was based in part on the way that artists paint on a canvas. To this end, we implemented an algorithm that spatially subdivides D into common regions (objects) based on attribute value, then paints each region independently to produce a finished result. Our technique follows four basic steps:

- (1) Segment D into p spatially connected regions, where the attribute values $a_{i,j}$ for the elements in each region R_k are within a given tolerance ε_j of one another.
- (2) For each region R_k containing elements $e_1, \dots, e_t$, compute a region-global stroke coverage from the average value $\frac{1}{t} \sum_{i=1}^t a_{i,j}$, where A_j is the attribute represented by coverage.
- (3) “Paint” strokes at randomly selected positions within R_k . The color, orientation, and size of each stroke are controlled by the attribute values of the element closest to the stroke’s center. A stroke is accepted or rejected based on its overlap with existing strokes, and on its overlap with R_k . This process continues until R_k ’s required coverage is met.
- (4) After all p regions are painted, display the result to the viewer.

Step three represents an important difference between our nonphotorealistic technique and glyph-based visualizations. Most glyph algorithms use a one-to-one or one-to-many mapping to represent each data element with individual glyphs. We wanted a method that was more analogous to how paintings are constructed: “objects” in a scene are identified and painted in turn. This is done by segmenting a dataset into spatial regions, then painting strokes within each region until an appropriate stroke coverage is met. In our technique strokes do not correspond to specific data elements, rather, the strokes are bound to the elements indirectly through the segments they belong to.

Segmentation is performed using a modified region-growing algorithm. The first element e_1 of a new segment R_k is selected from a list of elements that do not belong to any segment. Average attribute values $\bar{a}_j = a_{1,j}$, $j = 1, \dots, m$ are initialized based on e_1 . R_k is then grown as follows:

- (1) Consider all elements in the eight-neighbor region around e_1 .
- (2) If a neighboring element e_i is not part of some other segment, and if $|\bar{a}_j - a_{i,j}| \leq \varepsilon_j \forall j$, add e_i to R_k .
- (3) Update $\bar{a}_j$ based on $a_{i,j}$, then recursively consider the neighbors of e_i .
- (4) Continue until no more elements can be added to R_k .

Some care must be used during the updating of $\bar{a}_j$. We do not use the initial $a_{1,j}$ as a fixed average, for example, since this produces segments that are too sensitive to the selection of e_1 . Consider the visual example shown in Figure 7, where we segment a dataset of pixels with $m = 3$ attributes: red, green, and blue. The segment generated with fixed averages and e_1 selected from the lower-left corner of the image is shown in grey in Figures 7(a) and 7(e). Because the choice of e_1 produced $\bar{a}_j$ that were relatively dark, the segment is smaller than expected. Since $\bar{a}_j$ do not change as the segment is constructed, we cannot correct for this initial decision. Updating $\bar{a}_j$ for each e_i forces the averages to follow the structure of the segment as it grows. New attribute values $a_{i,j}$ must be properly weighted when they are added to $\bar{a}_j$, however. Consider Figures 7(b) and 7(f), which use a simple running average $\bar{a}_j = \frac{1}{2}(\bar{a}_j + a_{i,j})$ for each new element e_i . This places too much importance on the attribute values of e_i , producing segments that are too large. Intuitively, a running average pushes $\bar{a}_j$ too far in the direction of e_i ; if neighboring elements have similar attribute values, this significantly increases the likelihood that these neighbors will be also accepted into the segment.




Fig. 7. Examples of different segmentation algorithms applied to an RGB image of a golden poppy; (a, b, c, d) segments with a fixed, running average, weighted average with $w = 1$, and weighted average with $w = \frac{7}{8}$, respectively; (e, f, g, h) segments overlaid on the original RGB image.

The technique we implemented uses weighted averages to build data segments. Given elements $e_1, e_2, \dots, e_t$, the average values at step t during segment construction are:

$$\bar{a}_j = \frac{1}{\sum_{i=1}^t w^{i-1}} (w^0 a_{1,j} + w^1 a_{2,j} + \dots + w^{t-1} a_{t,j}), \quad j = 1, \dots, m \quad (1)$$

w is used to weight the contribution of each new element. When $w = 1$, $\bar{a}_j$ is a simple average of the attribute values within R_k . When $w < 1$, each additional element has a monotonically smaller effect on $\bar{a}_j$, allowing the averages to converge to near-constant values. This is particularly useful when we visualize datasets with smooth gradients. Specifying $w < 1$ allows the construction of segments that do not expand to fill the entire gradient. The fraction $1/\sum_{i=1}^t w^{i-1}$ clamps $\bar{a}_j$ to lie in the range $0 \dots a_j^{max}$, where a_j^{max} is the largest possible value for attribute A_j .

Figure 7 shows examples of two weighted average segments. In Figures 7(c) and 7(g), the averages are updated using $w = 1$. In Figures 7(d) and 7(h), the segment is built with $w = 0.875$. This produces a smaller result, since elements past the first few contribute little to each $\bar{a}_j$ (e.g., the tenth element at $t = 10$ accounts for $0.875^9 / \sum_{i=1}^{10} 0.875^{i-1} = 0.051$, or approximately 5.1% of the segment average). By varying w , we can control the relative size of the segments we generate.

Each segment R_k is painted by randomly placing brush strokes inside it. The percentage of R_k to be covered by its strokes (coverage) is defined based on $\frac{1}{t} \sum_{i=1}^t a_{i,j}$, where A_j is the attribute that represents coverage. Because the elements e_i within R_k must have similar attribute values, a region-global coverage produces an acceptable representation of A_j within R_k . As each new stroke is placed, two values are computed: the overlap with existing strokes, and the overlap with R_k 's extent. If the stroke overlap is too high, or if the segment overlap is too low, the stroke is rejected. The allowable stroke overlap is slowly increased to ensure that R_k 's coverage can be met. The color, orientation, and size of each stroke are chosen using the attribute values of the element closest to the stroke's center.

The brush strokes used in our current prototype are identical to the ones shown during our experiments. They are constructed with a simple texture mapping scheme. This technique is common in



Fig. 8. An example of our segmentation and brush stroke model being used to produce a nonphotorealistic rendering from an RGB image of a golden poppy.

nonphotorealistic rendering (e.g., in Haberli [1990], Hertzmann [1998], Litwinowicz [1997], and Meier [1996]). Real painted strokes are digitally captured and converted into texture maps. The textures are applied to an underlying polygon to produce a collection of generic brush strokes. We use a small library of representative stroke textures. One of the textures is randomly selected and bound to a stroke when it is placed. This produces a more random, hand-generated feel to the resulting images. The nonphotorealistic rendering of the complete golden poppy image is shown in Figure 8. Additional examples of renderings and visualizations are shown in Figures 1, 9, and 10.

6.2 Practical Applications

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




Fig. 9. Nonphotorealistic visualization of weather conditions for February over the eastern United States: (top row) perceptual color ramps (dark blue for low to bright pink for high) of mean *temperature*, *pressure*, *wind speed*, and *precipitation* in isolation; (bottom row) combined visualization of *temperature* (dark blue to bright pink for cold to hot), *wind speed* (low to high coverage for weak to strong), *pressure* (small to large for low to high), and *precipitation* (upright to flat for light to heavy).



Fig. 10. Weather conditions over the continental United States: (a) mean *temperature, pressure, wind speed* and *precipitation* (represented by color, size, coverage, and orientation) for January; (b) mean conditions for August.

mean *temperature, wind speed, pressure, and precipitation*. Based on this order of importance, we built a data-feature mapping M that varies brush stroke color, coverage, size, and orientation. This mapping divides the concept of spatial density into two separate parts: *size*, the size of the 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 brush stroke features. Size describes the energy 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 data-feature mapping M :

- $A_1 = \text{temperature} \rightarrow V_1 = \text{color}$, $\Phi_1 = \text{dark blue for low temperature to bright pink for high temperature}$,
- $A_2 = \text{wind speed} \rightarrow V_2 = \text{coverage}$, $\Phi_2 = \text{low coverage for weak wind speed to full coverage for strong wind speed}$,
- $A_3 = \text{pressure} \rightarrow V_3 = \text{size}$, $\Phi_3 = \text{small strokes for low pressure to large strokes for high pressure}$, and
- $A_4 = \text{precipitation} \rightarrow V_4 = \text{orientation}$, $\Phi_4 = \text{upright (90° rotation) for light precipitation to flat (0° rotation) for heavy precipitation}$.

Figure 9 shows an example of applying M to data for February along the east coast of the continental United States. The top four 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 *temperature*, *pressure*, *wind speed*, and *precipitation*. The result of applying M to construct a nonphotorealistic visualization of all four 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 size boundaries, shown as neighboring regions with different sized strokes (e.g., larger strokes in Florida represent higher pressure readings). Increases in rainfall are shown as an increasing stroke tilt running from upright (light precipitation) to flat (heavy precipitation). Finally, the wind's magnitude modifies stroke coverage: weak *wind speed* values produce small numbers of strokes with a large amount of background showing through (e.g., north of the Great Lakes), while strong *wind speed* values produce larger numbers of strokes that completely fill their corresponding screen space (e.g., in central Texas and Kansas).

Figure 10 uses the same M to visualize weather conditions over the continental United States for January and August. These visualizations provide a number of interesting insights into historical weather conditions for this part of the world. In January (Figure 10(a)), weak *wind speed* and *pressure* values (shown as small, low coverage strokes) cover much of western, southeastern, and northeastern parts of the country. Regions of much higher *pressure* are shown as larger strokes in the center of the map. Typically heavy *precipitation* 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. Conditions in August (Figure 10(b)) are markedly different. Most of the United States is warm with areas of intense heat, shown as bright pink strokes, visible in southern California, the southwest, and most of the southern states. Little *precipitation* is evident apart from Florida, where tilted strokes are displayed. Finally, *wind speed* to the west of the Rocky Mountains is much weaker than to the east; the background is clearly visible through the strokes in the west, while almost no background can be seen in the east.

7. VALIDATION EXPERIMENT

In order to further explore the capabilities of our nonphotorealistic techniques, we conducted a basic validation experiment designed to:

- (1) Test the ability of our nonphotorealistic visualization to support common analysis tasks on real-world data.




Fig. 11. Example multidimensional visualizations: (a) standard visualization of *temperature* with color (dark green for cooler to yellow for warmer); (b) standard visualization of *precipitation* with Doppler radar (green for light rainfall to red for heavy rainfall); (c) standard visualization of *wind direction* with directed contours and *wind speed* with color (dark blue for low winds to bright green for high winds); (d) a combination of three individual visualizations to form a single, multidimensional image; (e) a nonphotorealistic visualization with simulated paint strokes that vary their color, coverage, orientation, and size to visualize the same data.

- (2) Compare our nonphotorealistic visualization with a more traditional display method.
- (3) Study whether the common method of combining displays that work well in isolation produces an effective multidimensional visualization.

Our experiment compared user performance in our nonphotorealistic weather visualizations with more traditional displays. The dataset we used for this experiment contained four data attributes $A = (\text{temperature}, \text{wind speed}, \text{wind direction}, \text{precipitation})$. Based on consultation with domain experts from the natural sciences, we decided to composite standard displays of the individual attributes to produce a multidimensional result. Anecdotal feedback from the scientists suggested that our nonphotorealistic visualizations were better than the collection of side-by-side displays they often employ (e.g., Figures 11(a)–(c), which were captured directly from online weather maps), particularly when searching for combinations of weather conditions. This is not surprising, since a search across multiple images will produce change blindness. The low-level visual system cannot remember image detail

beyond the local region containing the viewer's focus of attention (see the Perceptual Properties section for a more detailed discussion of change blindness). Because many people are already familiar with standard weather maps, the scientists wondered whether a combination of these displays would still be effective. Our experiment was designed to study this question, and to compare the performance of a combined display to our nonphotorealistic visualizations.

Figure 11(d) shows the result of applying $M_s(V_s, \Phi_s)$ with $V_s = (\text{color, luminance, directed contours, semi-transparent color})$, $\Phi_s = (\text{green} \dots \text{yellow, dark} \dots \text{bright, } 0^\circ \dots 360^\circ, \text{green} \dots \text{red})$. Certain modifications were needed to combine the data into a single image. For example, *temperature*, *wind speed*, and *precipitation* are all represented by color in the individual displays (Figures 11(a)–(c)); we continued to use color to represent *temperature*, but switched to luminance to represent *wind speed*. This variation of luminance makes areas of weaker winds appear darker (i.e., lower luminance), and areas of stronger winds appear lighter (i.e., higher luminance and therefore less saturated). We left the Doppler radar traces of *precipitation* intact, but made them semi-transparent to try to show the underlying *temperature*, *wind speed*, and *wind direction*.

Figure 11(e) displays the same data in Figure 11(d) as a nonphotorealistic visualization. Here, M_n is defined as $V_n = (\text{color, coverage, orientation, size})$, $\Phi_n = (\text{dark blue} \dots \text{bright pink, low} \dots \text{high, } 0^\circ \dots 360^\circ, \text{small} \dots \text{large})$. Since the nonphotorealistic technique was specifically designed to visualize multidimensional datasets, none of the tradeoffs used in Figure 11(d) were needed. Because of the coarseness of the available data, we imposed a more regular structure on the positions of our brush strokes. Apart from this modification, the painting algorithm used for Figure 11(e) was identical to the one used for the previous datasets.

Fifteen observers (computer science graduate students and staff ranging in age from 18 to 41) with normal or corrected acuity and normal color vision participated during this experiment. The data-feature mappings M_s and M_n were explained in detail to the observers using a pair of visualizations different from the ones shown during the experiment. Observers were encouraged to ask questions to ensure they understood how each attribute was being represented. Observers were then instructed to answer the following questions on a new pair of visualization images (Figures 11(d) and 11(e)):

- (1) In which visualization is it easiest to distinguish: *temperature*; *precipitation*; *wind speed*; *wind direction*?
- (2) Identify an area in each visualization that has: high *temperature*; high *precipitation*; low *wind speed*.
- (3) Identify an area in each visualization that has: high *precipitation* and low *temperature*; high *precipitation* and high *wind speed*.
- (4) Identify an area in each visualization where *temperature* changes rapidly.

The first question queried an observer's preferences about the representation techniques used for each attribute. The second question tested an observer's ability to identify values for three different attributes. The third question tested an observer's ability to identify combinations of attribute values. The final question tested an observer's ability to identify high spatial frequency changes in one attribute (*temperature*) in the presence of a second (*precipitation*). As with the construction of the standard visualization image, these questions were selected in part through suggestions from our natural science colleagues.

7.1 Results

Responses were recorded and tabulated for all fifteen observers. Chi-squared tests with a standard 95% confidence interval were used to denote significance. In summary, we found:

Table I. Combined Responses for the Question: “In Which Visualization is it Easiest to Distinguish the Given Data Attribute?”

Visualization	Temperature	Precipitation	Wind speed	Wind direction
Standard	1	13	1	5
Nonphotorealistic	14	2	14	10

Table II. Combined Responses for the Task: “Identify an Area in Each Visualization that has the Following Attribute Value”

Visualization	Response	High <i>temperature</i>	High <i>precipitation</i>	Low <i>wind speed</i>
Standard	correct	10	14	11
	incorrect	2	0	0
	hard to tell	3	1	4
Nonphotorealistic	correct	15	13	13
	incorrect	0	0	0
	hard to tell	0	2	2

- (1) Observers preferred the nonphotorealistic visualization’s method of representing *temperature* and *wind speed*.
- (2) Observers preferred the standard visualization’s method of representing *precipitation*.
- (3) Observers were better at identifying high *temperature* in the nonphotorealistic visualization.
- (4) Observers were better at identifying a combination of high *precipitation* and high *wind speed* in the nonphotorealistic visualization.
- (5) Observers were better at identifying areas of rapid *temperature* change in the nonphotorealistic visualization.

Table I details observer preferences for the visualization that they felt made each data attribute easiest to distinguish. A chi-squared test showed significant variation within the table as a whole ($\chi^2_3 = 28.8$, $p < 0.001$). Chi-squared tests on each attribute identified a significant preference for the nonphotorealistic visualization for *temperature* and *wind speed* ($\chi^2_1 = 11.267$, $p < 0.001$ in both cases), and a significant preference for the standard visualization for *precipitation* ($\chi^2_1 = 8.067$, $p < 0.01$). Observers indicated that it was easier to see *precipitation* in the standard visualization, since it sat “on top” of the other attributes. However, this made it difficult to distinguish *temperature* and *wind speed* in areas of high *precipitation* (and thus the preference for the nonphotorealistic visualization’s method of displaying these attributes). Although *wind direction* was also obscured by *precipitation* in the standard visualization, some observers felt they could infer its pattern from what they could see entering and exiting areas of high rainfall.

Table II details observer performance for the task of identifying the location of high or low attribute values in the visualization. For this task, “high” was considered to be any value in the top 10% of the range shown in the visualization, and “low” was any value in the bottom 10%. “Correct” means an observer correctly identified an area in the visualization that contained the target attribute value. “Incorrect” means an observer identified an area that did not contain the target attribute value. “Hard to tell” means an observer gave no answer, but instead reported it was “hard to tell” where the target value was located. Performance for identifying high *temperature* was significantly better in the nonphotorealistic visualization ($\chi^2_2 = 6.00$, $p < 0.05$). There was no statistical difference in performance for the other two attributes. Interestingly, although observers stated a preference for the way *precipitation* was displayed in the standard visualization (see Table I), this did not produce any improvement in identifying regions of high *precipitation* ($\chi^2_2 = 0.37$, $p < 0.90$).

Table III. Combined Responses for the Task: “Identify an Area in Each Visualization that has the Following Combinations of Attribute Values”

Visualization	Response	High <i>precipitation</i> and Low <i>temperature</i>	High <i>precipitation</i> and High <i>wind speed</i>
Standard	correct	9	7
	incorrect	5	8
	hard to tell	1	0
Nonphotorealistic	correct	13	14
	incorrect	2	1
	hard to tell	0	0

Table IV. Combined Responses for the Task: “Identify an Area in Each Visualization with Rapid *Temperature* Change”

Visualization	Response	Rapid <i>temperature</i> change
Standard	correct	4
	incorrect	11
Nonphotorealistic	correct	12
	incorrect	3

Table III details observer performance for the task of identifying the location of a combination of high and low attribute values in the visualization (with “high” and “low” defined as before). In both cases, absolute performance was better in the nonphotorealistic visualization, although it was statistically significant only for identifying combinations of high *precipitation* and high *wind speed* ($\chi^2_2 = 7.778$, $p < 0.05$). Observers reported that it was easier to see color differences (i.e., variations in *temperature*) through the semitransparent Doppler radar traces in the standard visualization, compared to luminance differences (i.e., variations in *wind speed*). This explained the slightly better absolute performance in the standard visualization for the first task: Identify areas of high *precipitation* and low *temperature* (versus the second task of identifying high *precipitation* and high *wind speed*).

Table IV details observer performance for the task of identifying rapid changes in *temperature*. These areas were known to be located within areas of high *precipitation*, so the question was designed to test an observer’s ability to identify sharp variations in one attribute (*temperature*) in the presence of a second (*precipitation*). Results showed a significant performance advantage in the nonphotorealistic visualization ($\chi^2_1 = 8.572$, $p < 0.01$).

7.2 Interpretation

Although the standard visualization appeals to our familiarity with the weather maps we often see in day-to-day life, it was not built with methods that support rapid and accurate multidimensional analysis. This fact was highlighted during our experiment. Results showed that performance with the nonphotorealistic visualization matched or exceeded the standard visualization in all cases. This suggests that a method specifically designed for multidimensional data produces better visualizations than a combination of displays that work well in isolation. It also demonstrates that the nonphotorealistic visualizations are effective at representing multidimensional data in a way that supports real-world analysis tasks.

Given the foundations used to build the visualizations (rules of perception versus effective visualizations in isolation), the fact that the nonphotorealistic visualization outperformed the standard visualization in certain situations is not surprising. What was unexpected was that the standard visualization was *never* better than the nonphotorealistic visualization for the tasks we tested. Choosing

representations in the standard visualization that favor some attributes (e.g., *precipitation*) at the expense of others should make these attributes highly salient. This was exactly what we observed, for example, in Tables III and IV where the presence of *precipitation* in the standard visualization interfered with the identification of *wind speed* and *temperature*, respectively. Our results therefore suggest that every attribute representation in the nonphotorealistic visualization is at least as good as the corresponding attribute representation in the standard visualization.

A number of issues were raised when we tried to combine the individual displays to produce the standard weather visualization. These included occlusion (e.g., semi-transparent Doppler radar traces obscured underlying *temperature*, *wind speed*, and *wind direction* values), and links between visual features that caused variations in one to affect another (e.g., luminance variations used to represent *wind speed* lightened or darkened the colors used to represent *temperature*).

A separate problem was the choice of features used in the individual displays. These choices were not always well-suited to the tasks the scientists said they wanted to perform. For example, the standard visualization uses a static colormap that assigns a fixed color to each range of temperatures. This is a common technique used to facilitate comparison across multiple weather maps. Unfortunately, it also results in a narrow range of colors when a user chooses to study a local region of interest. Our visualization scales the colormap to fit the range of attribute values being displayed.² The narrow color range made it difficult for users to identify specific *temperature* values in the standard visualization (both in isolation and in the presence of high *precipitation*). It may have been possible to replace the colormap to try to overcome some of these problems. This would not address the issues of variations in luminance to visualize *wind speed*, or the occlusion that occurs in areas of high *precipitation*, however. Our intuition is that the standard visualization would continue to produce poor representations for certain tasks, and would not outperform the nonphotorealistic visualization, even with a more expressive colormap.

Although these experiments visualized weather data, we are not restricted to this domain. We are currently applying our nonphotorealistic techniques to scientific simulation results in oceanography, and to the problem of tracking intelligent agents interacting in a simulated e-commerce auction environment. Building on the strengths of the low-level human visual system provides the flexibility needed to construct effective multidimensional visualizations for a wide range of problem environments.

8. CONCLUSIONS AND FUTURE WORK

This paper describes a method of visualization that uses painted brush strokes to represent multidimensional data elements. Our goal was to produce effective nonphotorealistic visualizations. We were motivated in part by nonphotorealistic rendering in computer graphics, and by the work of Laidlaw, Interrante, and Ebert and Rheingans to extend these techniques to a visualization environment. Our contributions to this work are the application of human perception during the selection of a data-feature mapping, and the use of controlled experiments to study the effectiveness of a nonphotorealistic visualization, both in a laboratory setting, and in a more practical, real-world context.

The brush strokes we used support the variation of visual features that were selected based on styles from the Impressionist school of painting. Each attribute in a dataset is mapped to a specific nonphotorealistic property; attribute values stored in a data element can then be used to vary the visual appearance of the brush strokes. The properties we chose correspond closely to perceptual features detected by the low-level human visual system. Experimental results show that existing guidelines on the use of perception during visualization extend to a nonphotorealistic environment. This allows us to

²In the case of visualizing more than one map, we first combine temperature ranges from each map, then scale our colormap to cover this combined range; in this way the same colors in different displays properly correspond to the same temperature values (e.g., see Figures 9 and 10).

optimize the selection and application of our brush stroke properties. The result is a “painted image” whose color and texture patterns can be used to explore, analyze, verify, and discover information stored in a multidimensional dataset. We are optimistic that future results from studies of perception in visualization will also apply to our nonphotorealistic domain.

In addition to being effective, our techniques try to produce visualizations that viewers perceive as engaging or aesthetic. Nonphotorealistic techniques that highlight important or unexpected properties can be used to orient a viewer’s attention to specific areas in the image. An engaging visualization will encourage a more in-depth examination of these details.

A number of areas for future work are now being considered. Experiments are currently underway to try to measure the level of artistic merit viewers attach to our visualizations, and to identify the basic emotional and visual composition properties of the images (e.g., pleasure, arousal, meaning, and complexity) that affect these judgments. One question of interest asks: “Can we use these results to vary a visualization’s composition in ways that improve its artistic merit?” For example, we could try to increase the meaning of a visualization image by explaining what it represents and how it is used. If meaning is a predictor of artistic beauty, we would expect to see an increase in observers’ artistic merit rankings of the visualization images. Another area for investigation asks: “How do knowledge and experience affect the rating scales?” Our observers are, for the most part, artistic novices. Conducting an experiment with participants who have some type of formal training in art theory and art history could offer important insights on how this knowledge affects appreciation of our different image types. Results from these two questions may show that our current emotional and visual composition properties need to be refined or extended to further differentiate the artistic merit attached to different images. We are evaluating new candidate properties to test during future studies.

Another interesting suggestion is to compare the artistic merit of our nonphotorealistic visualizations with traditional visualization techniques (e.g., multidimensional glyphs). We are now studying this possibility as a follow-on to our current experiments.

Our brush strokes support the variation of color, orientation, coverage, and size. We are working to identify new nonphotorealistic properties that could be integrated into our stroke model. Two promising candidates we have already discussed are coarseness and weight. Other properties are being sought using two complementary approaches. First, we are reviewing literature on technique and style in Impressionist art. Second, we are looking at perceptually salient visual features that may correspond to new nonphotorealistic properties. Increasing the number of features we can encode effectively in each brush stroke may allow us to represent datasets with higher dimensionality.

The need to display additional nonphotorealistic properties may exceed the abilities of our simple texture mapped stroke model. We are studying three techniques to overcome this limitation: (1) the creation of a larger library of texture mapped brush strokes that explicitly vary the properties that are not easy to modify within an individual texture map; (2) a model that uses spline surfaces to construct continuous representations of the multiple properties in a brush stroke, and (3) a model that uses a physical simulation to vary nonphotorealistic properties and construct visually realistic strokes.

Finally, we note one other important advantage we can derive from the correspondence between perceptual features and nonphotorealistic properties. We measure the perceptual salience of a visual feature using controlled psychophysical experiments. Exactly the same technique is used to investigate the strengths and limitations of new nonphotorealistic features, both in isolation and when displayed together with other stroke properties. Just as research in perception helps us to identify and control nonphotorealistic features during visualization, work on new features may offer insight into how the low-level visual system “sees” certain combinations of visual properties. These results could have an important impact on models of low-level human vision that are being constructed by researchers in the psychophysical community.

ACKNOWLEDGMENTS

The authors would like to thank Victoria Interrante, David Laidlaw, Penny Rheingans, and Theresa-Marie Rhyne for their suggestions and insights during discussions of this work. The authors would also like to thank Peter Rand for his suggestions for a standard visualization image and analysis tasks to test during our validation experiments.

REFERENCES

- AKS, D. J. AND ENNS, J. T. 1996. Visual search for size is influenced by a background texture gradient. *J. Experim. Psych.: Human Percept. Perf.* 22, 6, 1467–1481.
- BERGMAN, L. D., ROGOWITZ, B. E., AND TREINISH, L. A. 1995. A rule-based tool for assisting colormap selection. In *Proceedings of Visualization '95* (Atlanta, Ga.). 118–125.
- BIRREN, F. 1969. *Munsell: A Grammar of Color*. Van Nostrand Reinhold Company, New York, New York.
- BROWN, R. 1978. Impressionist technique: Pissarro's optical mixture. In *Impressionism in Perspective*, B. E. White, Ed. Prentice-Hall, Inc., Englewood Cliffs, N. J., 114–121.
- CALLAGHAN, T. C. 1990. Interference and dominance in texture segregation. In *Visual Search*, D. Brogan, Ed. Taylor & Francis, New York, 81–87.
- CHEVREUL, M. E. 1967. *The Principles of Harmony and Contrast of Colors and Their Applications to the Arts*. Reinhold Publishing Corporation, New York.
- CIE. 1978. *CIE Publication No. 15, Supplement Number 2 (E-1.3.1, 1971): Official Recommendations on Uniform Color Spaces, Color-Difference Equations, and Metric Color Terms*. Commission Internationale de L'Éclairage.
- COREN, S., WARD, L. M., AND ENNS, J. T. 2003. *Sensation and Perception (6th Edition)*. Wiley, New York, New York.
- CURTIS, C. J., ANDERSON, S. E., SEIMS, J. E., FLEISCHER, K. W., AND SALESIN, D. H. 1997. Computer-generated watercolor. In *SIGGRAPH 97 Conference Proceedings* (Los Angeles, Calif.) T. Whitted, Ed. ACM, New York, 421–430.
- CUTTING, J. E. AND MILLARD, R. T. 1984. Three gradients and the perception of flat and curved surfaces. *J. Experiment. Psych.: General* 113, 2, 198–216.
- EBERT, D. AND RHEINGANS, P. 2000. Volume illustration: Non-photorealistic rendering of volume models. In *Proceedings of Visualization 2000* (San Francisco, Calif.). 195–202.
- EGETH, H. E. AND YANTIS, S. 1997. Visual attention: Control, representation, and time course. *Ann. Rev. Psychol.* 48, 269–297.
- FINKELSTEIN, A. AND SALESIN, D. H. 1994. Multiresolution curves. In *SIGGRAPH 94 Conference Proceedings* (Orlando, Fla.), A. S. Glassner, Ed. ACM, New York, 261–268.
- GOOCH, B., COOMBE, G., AND SHIRLEY, P. 2002. Artistic vision: Painterly rendering using computer vision techniques. In *Proceedings of the NPAR 2002 Symposium on Non-Photorealistic Animation and Rendering* (Annecy, France). 83–90.
- GOOCH, B. AND GOOCH, A. 2001. *Non-Photorealistic Rendering*. A K Peters, Ltd., Natick, Mass.
- GRINSTEIN, G., PICKETT, R., AND WILLIAMS, M. 1989. EXVIS: An exploratory data visualization environment. In *Proceedings of Graphics Interface '89* (London, Ont., Canada). 254–261.
- HABERLI, P. 1990. Paint by numbers: Abstract image representations. *Comput. Graph. (SIGGRAPH 90 Conference Proceedings)* 24, 4, 207–214.
- HABERLI, P. AND SEGAL, M. 1993. Texture mapping as a fundamental drawing primitive. In *Proceedings of the 4th Eurographics Workshop on Rendering*, (Paris, France). M. Cohen, C. Puech, and F. Sillion, Eds. 259–266.
- HARALICK, R. M., SHANMUGAM, K., AND DINSTEN, I. 1973. Textural features for image classification. *IEEE Trans. Syst., Man, and Cybernet.* SMC-3, 6, 610–621.
- HEALEY, C. G. 1996. Choosing effective colors for data visualization. In *Proceedings of Visualization '96* (San Francisco, Calif.). 263–270.
- HEALEY, C. G., BOOTH, K. S., AND ENNS, J. T. 1996. High-speed visual estimation using preattentive processing. *ACM Trans. Computer-Hum. Interact.* 3, 2, 107–135.
- HEALEY, C. G. AND ENNS, J. T. 1998. Building perceptual textures to visualize multidimensional datasets. In *Proceedings of Visualization '98* (Research Triangle Park, N. C.). 111–118.
- HEALEY, C. G. AND ENNS, J. T. 1999. Large datasets at a glance: Combining textures and colors in scientific visualization. *IEEE Trans. Visual. Comput. Graph.* 5, 2, 145–167.
- HERING, E. 1964. *Outlines of a Theory of Light Sense*. Harvard University Press, Cambridge, Mass.
- HERTZMANN, A. 1998. Painterly rendering with curved brush strokes of multiple sizes. In *SIGGRAPH 98 Conference Proceedings* (Orlando, Fla.). M. Cohen, Ed. ACM, New York, 453–460.

- HERTZMANN, A. 2002. Fast texture maps. In *Proceedings of NPAR 2002 Symposium on Non-Photorealistic Animation and Rendering* (Annecy, France). 91–96.
- HERTZMANN, A., JACOBS, C. E., OLIVER, N., CURLESS, B., AND SALESIN, D. H. 2001. Image analogies. In *SIGGRAPH 2001 Conference Proceedings* (Los Angeles, Calif.). E. Fiume, Ed. ACM, New York, 327–340.
- HSU, S. C. AND LEE, I. H. H. 1994. Drawing and animation using skeletal strokes. In *SIGGRAPH 94 Conference Proceedings* (Orlando, Fla.). A. Glassner, Ed. ACM, New York, 109–118.
- INTERRANTE, V. 2000. Harnessing natural textures for multivariate visualization. *IEEE Comput. Graph. Applic.* 20, 6, 6–11.
- JULÉSZ, B. 1975. Experiments in the visual perception of texture. *Scient. Amer.* 232, 34–43.
- JULÉSZ, B. 1984. A brief outline of the texton theory of human vision. *Trends Neurosci.* 7, 2, 41–45.
- JULÉSZ, B., GILBERT, E. N., AND SHEPP, L. A. 1973. Inability of humans to discriminate between visual textures that agree in second-order statistics—revisited. *Perception* 2, 391–405.
- JULÉSZ, B., GILBERT, E. N., AND VICTOR, J. D. 1978. Visual discrimination of textures with identical third-order statistics. *Biologic. Cybernet.* 31, 137–140.
- KIRBY, R. M., MARMANIS, H., AND LAIDLAW, D. H. 1999. Visualizing multivalued data from 2D incompressible flows using concepts from painting. In *Proceedings of Visualization '99* (San Francisco, Calif.). 333–340.
- LAIDLAW, D. H. 2001. Loose, artistic “textures” for visualization. *IEEE Comput. Graph. Applic.* 21, 2, 6–9.
- LAIDLAW, D. H., AHRENS, E. T., KREMERS, D., AVALOS, M. J., JACOBS, R. E., AND READHEAD, C. 1998. Visualizing diffusion tensor images of the mouse spinal cord. In *Proceedings of Visualization '98* (Research Triangle Park, N. C.). 127–134.
- LEWIS, J.-P. 1984. Texture synthesis for digital painting. *Computer Graphics (SIGGRAPH 84 Proceedings)* 18, 3, 245–252.
- LITWINOWICZ, P. 1997. Processing images and video for an impressionist effect. In *SIGGRAPH 97 Conference Proceedings*, (Los Angeles, Calif.). T. Whitted, Ed. ACM, New York, 407–414.
- LIU, G., HEALEY, C. G., AND ENNS, J. T. 2003. Target detection and localization in visual search: A dual systems perspective. *Percept. Psychophys.* 65, 5, 678–694.
- LU, A., MORRIS, C. J., EBERT, D. S., RHEINGANS, P., AND HANSEN, C. 2002. Non-photorealistic volume rendering using stippling techniques. In *Proceedings of Visualization 2002* (Boston, Mass.), 211–218.
- MACEachREN, A. M. 1995. *How Maps Work*. Guilford Publications, Inc., New York.
- MACK, A. AND ROCK, I. 1998. *Inattentional Blindness*. MIT Press, Menlo Park, Calif.
- MCCORMICK, B. H., DEFANTI, T. A., AND BROWN, M. D. 1987. Visualization in scientific computing. *Comput. Graph.* 21, 6, 1–14.
- MEIER, B. J. 1996. Painterly rendering for animation. In *SIGGRAPH 96 Conference Proceedings* (New Orleans, La.) H. Rushmeier, Ed. ACM, New York, 477–484.
- MUNSELL, A. H. 1905. *A Color Notation*. Geo. H. Ellis Co., Boston, Mass.
- POMERANTZ, J. AND PRISTACH, E. A. 1989. Emergent features, attention, and perceptual glue in visual form perception. *J. Experim. Psych.: Human Percept. Perf.* 15, 4, 635–649.
- POSNER, M. I. AND RAICHLE, M. E. 1994. *Images of mind*. Scientific American Library.
- RAMACHANDRAN, V. S. AND HIRSTEIN, W. 1999. The science of art: A neurological theory of aesthetic experience. *J. Conscious. Stud.* 6, 6-7, 15–51.
- RAO, A. R. AND LOHSE, G. L. 1993a. Identifying high level features of texture perception. *CVGIP: Graph. Models Image Process.* 55, 3, 218–233.
- RAO, A. R. AND LOHSE, G. L. 1993b. Towards a texture naming system: Identifying relevant dimensions of texture. In *Proceedings of Visualization '93* (San Jose, Calif.), 220–227.
- RENSINK, R. A. 2000. Seeing, sensing, and scrutinizing. *Vision Res.* 40, 10–12, 1469–1487.
- RHEINGANS, P. AND EBERT, D. 2001. Volume illustration: Nonphotorealistic rendering of volume models. *IEEE Trans. Vis. Comput. Graph.* 7, 3, 253–264.
- RHEINGANS, P. AND TEBBS, B. 1990. A tool for dynamic explorations of color mappings. *Comput. Graph.* 24, 2, 145–146.
- ROGOWITZ, B. E. AND TREINISH, L. A. 1993. An architecture for rule-based visualization. In *Proceedings of Visualization '93* (San Jose, Calif.), 236–243.
- ROOD, O. N. 1879. *Modern Chromatics, with Applications to Art and Industry*. Appleton, New York.
- ROSENBLUM, L. J. 1994. Research issues in scientific visualization. *IEEE Comput. Graph. Applic.* 14, 2, 61–85.
- SALISBURY, M., ANDERSON, C., LISCHINSKI, D., AND SALESIN, D. H. 1996. Scale-dependent reproduction of pen-and-ink illustrations. In *SIGGRAPH 96 Conference Proceedings* (New Orleans, La.) H. Rushmeier, Ed. ACM, New York, 461–468.
- SALISBURY, M., ANDERSON, S. E., BARZEL, R., AND SALESIN, D. H. 1994. Interactive pen-and-ink illustrations. In *SIGGRAPH 94 Conference Proceedings*, (Orlando, Fla.). A. S. Glassner, Ed. ACM, New York, 101–108.

- SALISBURY, M., WONG, M. T., HUGHES, J. F., AND SALESIN, D. H. 1997. Orientable textures for image-based pen-and-ink illustration. In *SIGGRAPH 97 Conference Proceedings* (Los Angeles, Calif.). T. Whitted, Ed. 401–406.
- SCHAPIRO, M. 1997. *Impressionism: Reflections and Perceptions*. George Brazillier, Inc., New York.
- SHIRAISHI, M. AND YAMAGUCHI, Y. 1999. Image moment-based stroke placement. In *SIGGRAPH 99: Sketches & Applications* (Los Angeles, Calif.). R. Kidd, Ed. ACM, New York, 247.
- SIMONS, D. J. 2000. Current approaches to change blindness. *Vis. Cognit.* 7, 1/2/3, 1–15.
- SLOCUM, T. A. 1998. *Thematic Cartography and Visualization*. Prentice-Hall, Inc., Upper Saddle River, N. J.
- SMITH, P. H. AND VAN ROSENDALE, J. 1998. Data and visualization corridors report on the 1998 CVD workshop series (sponsored by DOE and NSF). Tech. Rep. CACR-164, Center for Advanced Computing Research, California Institute of Technology.
- SNOWDEN, R. J. 1998. Texture segregation and visual search: A comparison of the effects of random variations along irrelevant dimensions. *J. Experiment. Psych.: Human Percept. Perform.* 24, 5, 1354–1367.
- SOUSA, M. C. AND BUCHANAN, J. W. 1999a. Computer-generated graphite pencil rendering of 3d polygon models. *Comput. Graph. Forum (Proceedings Eurographics '99)* 18, 3, 195–208.
- SOUSA, M. C. AND BUCHANAN, J. W. 1999b. Computer-generated pencil drawings. In *Proceedings SKIGRAPH '99* (Banff, Canada).
- SOUSA, M. C. AND BUCHANAN, J. W. 2000. Observational models of graphite pencil materials. *Comput. Graph. Forum* 19, 1, 27–49.
- STRASSMANN, S. 1986. Hairy brushes. *Comput. Graph. (SIGGRAPH 86 Proceedings)* 20, 4, 185–194.
- STROTHOTTE, T. AND SCHLECHTWEIG, S. 2002. *Non-Photorealistic Computer Graphics: Modeling, Rendering and Animation*. Morgan-Kaufmann, Inc., San Francisco, Calif.
- TAKAGI, S. AND FUJISHIRO, I. 1997. Microscopic structural modeling of colored pencil drawings. In *SIGGRAPH 97 Sketches & Applications* (Los Angeles, Calif.). D. S. Ebert, Ed. ACM, New York, 187.
- TAKAGI, S., FUJISHIRO, I., AND NAKAJIMA, M. 1999. Volumetric modeling of artistic techniques in colored pencil drawing. In *SIGGRAPH 99 Sketches & Applications*, (Los Angeles, Calif.). R. Kidd, Ed. ACM, New York, 283.
- TAMURA, H., MORI, S., AND YAMAWAKI, T. 1978. Textural features corresponding to visual perception. *IEEE Trans. Syst. Man, and Cybernet. SMC-8*, 6, 460–473.
- TRIESMAN, A. 1985. Preattentive processing in vision. *Comput. Vis. Graph. Image Process.* 31, 156–177.
- TRIESMAN, A. 1991. Search, similarity, and integration of features between and within dimensions. *J. Experiment. Psych.: Human Percept. Perf.* 17, 3, 652–676.
- TRIESMAN, A. AND GORMICAN, S. 1988. Feature analysis in early vision: Evidence from search asymmetries. *Psychol. Rev.* 95, 1, 15–48.
- TUFTE, E. R. 1983. *The Visual Display of Quantitative Information*. Graphics Press, Cheshire, Conn.
- TUFTE, E. R. 1990. *Envisioning Information*. Graphics Press, Cheshire, Conn.
- TUFTE, E. R. 1997. *Visual Explanations: Images and Quantities, Evidence and Narrative*. Graphics Press, Cheshire, Conn.
- VENTURI, L. 1978. Impressionist style. In *Impressionism in Perspective*, B. E. White, Ed. Prentice-Hall, Inc., Englewood Cliffs, N. J., 105–113.
- WARE, C. 1988. Color sequences for univariate maps: Theory, experiments, and principles. *IEEE Comput. Graph. Applic.* 8, 5, 41–49.
- WARE, C. 2000. *Information Visualization: Perception for Design*. Morgan-Kaufmann, San Francisco, Calif.
- WARE, C. AND KNIGHT, W. 1995. Using visual texture for information display. *ACM Trans. Graph.* 14, 1, 3–20.
- WEIGLE, C., EMIGH, W. G., LIU, G., TAYLOR, R. M., ENNS, J. T., AND HEALEY, C. G. 2000. Oriented texture slivers: A technique for local value estimation of multiple scalar fields. In *Proceedings of Graphics Interface 2000* (Montréal, Quebec, Canada), 163–170.
- WINKENBACH, G. AND SALESIN, D. H. 1994. Computer-generated pen-and-ink illustration. In *SIGGRAPH 94 Conference Proceedings*. (Orlando, Fla.). A. Glassner, Ed. ACM, New York, 91–100.
- WINKENBACH, G. AND SALESIN, D. H. 1996. Rendering free-form surfaces in pen-and-ink. In *SIGGRAPH 96 Conference Proceedings*, (New Orleans, La.). H. Rushmeier, Ed. ACM, New York, 469–476.
- WOLFE, J. M. 1994. Guided Search 2.0: A revised model of visual search. *Psycho. Bull. Rev.* 1, 2, 202–238.
- WOLFE, J. M., KLEMPEN, N., AND DAHLEN, K. 2000. Post attentive vision. *J. Experiment. Psych.: Human Percept. Perform.* 26, 2, 693–716.
- ZEKI, S. 1999. *Inner Vision*. Oxford University Press, Oxford, U.K.

Received January 2003; revised May 2003; accepted August 2003