

# Perceptually-Motivated Graphics, Visualization and 3D Displays

**Ann McNamara\***

Department of Visualization  
Texas A&M University

**Katerina Mania<sup>†</sup>**

Department of Electronic & Computer Engineering  
Technical University of Crete

**Marty Banks<sup>‡</sup>**

Visual Space Perception Laboratory  
University of California, Berkeley

**Christopher Healey<sup>§</sup>**

Department of Computer Science  
North Carolina State University

May 18, 2010

---

\*ann@viz.tamu.edu

<sup>†</sup>mania@ced.tuc.gr

<sup>‡</sup>martybanks@berkeley.edu

<sup>§</sup>healey@csc.ncsu.edu

## **Abstract**

This course presents timely, relevant examples on how researchers have leveraged perceptual information for optimization of rendering algorithms, to better guide design and presentation in (3D stereoscopic) display media, and for improved visualization of complex or large data sets. Each presentation will provide references and short overviews of cutting-edge current research pertaining to that area. We will ensure that the most up-to-date research examples are presented by sourcing information from recent perception and graphics conferences and journals such as ACM Transactions on Perception, paying particular attention work presented at the 2010 Symposium on Applied Perception in Graphics and Visualization.

## About the Lecturers

### **Ann McNamara**

Department of Visualization  
Texas A&M University  
3137 TAMU  
College Station, TX 77843-3137

+1-979-845-4715  
ann@viz.tamu.edu  
<http://www.viz.tamu.edu/people/ann>

Ann McNamara received her undergraduate and graduate degrees from the University of Bristol, UK. Anns research focuses on the advancement of computer graphics and scientific visualization through novel approaches for optimizing an individuals experience when creating, viewing and interacting with virtual spaces. She investigates new ways to exploit knowledge of human visual perception to produce high quality computer graphics and animations more efficiently. She joined the faculty of the newly formed Department of Visualization at Texas A&M University in 2008, where she is currently an assistant professor. Ann serves on several IPCs including APGV.

### **Katerina Mania**

Department of Electronic and Computer Engineering  
Technical University of Crete  
University Campus Kounoupidiana Chania Crete Greece

+30 28210 37222  
k.mania@ced.tuc.gr  
<http://www.music.tuc.gr/kmania>

Katerina Mania completed a B.Sc. in Mathematics, University of Crete, Greece, an M.Sc./Ph.D in Computer Science, University of Bristol, UK, funded by HP Labs. She worked at HP Labs as a researcher before serving on Faculty in the Department of Informatics, University of Sussex. Katerina spent her sabbatical at NASA Ames Research Centre (Advanced Displays and Spatial Perception Laboratory) in 2003. She is currently an Assistant Professor with tenure at the Technical University of Crete, Greece. Katerina is the program co-chair for APGV 2010 and Associate Editor of ACM Transactions on Applied Perception and Presence Teleoperators and Virtual Environments.

## **Marty Banks**

Marty Banks' Lab  
University of California, Berkeley  
Vision Science, 360 Minor Hall  
Berkeley, CA 94720-2020

+1-510-642-7679  
martybanks@berkeley.edu  
<http://bankslab.berkeley.edu/>

Martin S. Banks received his Bachelors degree at Occidental College (1970). After one year in Germany teaching, he entered the graduate program in Psychology at UC San Diego. He received a Masters degree in Experimental Psychology (1973). Banks then transferred to the graduate program at the University of Minnesota where he received his PhD. in Developmental Psychology (1976). He was Assistant & Associate Professor of Psychology at the University of Texas at Austin (1976-1985). He moved to UC Berkeley School of Optometry in 1985 where his is now a Full Professor of Optometry and Vision Science.

## **Christopher Healey**

Department of Computer Science  
North Carolina State University  
890 Oval Drive #8206

+1 919.513.8112  
healey@csc.ncsu.edu  
<http://www.csc.ncsu.edu/faculty/healey>

Christopher G. Healey received a B.Math from the University of Waterloo in Waterloo, Canada, and an M.Sc. and Ph.D. from the University of British Columbia in Vancouver, Canada. Following a postdoctoral fellowship at the University of California at Berkeley, he joined the Department of Computer Science at North Carolina State University, where he is currently an Associate Professor. His research interests include visualization, graphics, visual perception, and areas of applied mathematics, databases, artificial intelligence, and aesthetics related to visual analysis and data management.

## Course Overview

### 5 minutes: Welcome and Introductions

*Ann McNamara*

Welcome, overview of course & motivation for attending. Speaker Introductions

### 40 minutes: Perceptually Motivated 3D Displays & Depth Perception

*Martin Banks*

An overview of Depth Perception and important phenomenon when presenting information on 3D Displays

### 40 minutes: Perceptually Motivated Visualization

*Chris Healey*

A look at Visual Attention, Visual Memory, and its Role in Visualization.

### 15 minutes: Break

### 30 minutes: Perceptually Motivated Rendering

*Ann McNamara*

Overview of how knowledge from perceptual research feeds into optimized rendering algorithms.

### 30 minutes: Perceptually Motivated Simulation and Virtual Environments

*Katerina Mania*

Perceptually-based Optimizations & Fidelity Metrics for Simulation Technology

### 30 minutes: Leading-edge research and APGV 2010

*Katerina Mania & Martin Banks*

A summary of cutting edge perceptual research selected from APGV 2010.

### 10 minutes: A look to the future

*Ann McNamara & Katerina Mania*

Discussion of trends for APGV 2010

### 10 minutes: Conclusion, Questions & Answers

*All*

Wrap up, review, questions and discussion.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Motivation . . . . .	1
1.2	Course Overview . . . . .	1
1.3	Focus Areas . . . . .	1
1.3.1	Perceptually Motivated 3D Displays & Depth Perception	1
1.3.2	Perceptually Motivated Visualization . . . . .	2
1.3.3	Exploitation of the limitations of the HVS to reduce rendering times . . . . .	2
1.3.4	Exploration of incorporating perceptual and cognitive aspects to Virtual Environments (VEs). . . . .	2
1.4	Summary . . . . .	2
<b>2</b>	<b>Perceptually Motivated 3D Displays &amp; Depth Perception</b>	<b>3</b>
2.1	Introduction . . . . .	3
<b>3</b>	<b>Perceptually Motivated Visualization</b>	<b>37</b>
3.1	Introduction . . . . .	37
3.2	Visual Attention and Preattentive Processing . . . . .	37
3.3	Theories of Preattentive Processing . . . . .	39
3.3.1	Feature Integration . . . . .	39
3.3.2	Textons . . . . .	39
3.3.3	Guided Search . . . . .	40
3.3.4	Boolean Maps . . . . .	41
3.3.5	Ensemble Coding . . . . .	42
3.3.6	Feature Hierarchies . . . . .	42
3.4	Visual Memory . . . . .	43
3.4.1	Change Blindness . . . . .	43
3.4.2	Inattentional Blindness . . . . .	43
3.4.3	Attentional Blink . . . . .	45
3.5	Conclusions . . . . .	45
<b>4</b>	<b>Perceptually Motivated Rendering</b>	<b>46</b>
4.1	Visual Perception in Realistic Image Synthesis . . . . .	46
4.1.1	Visual Perception . . . . .	46
4.1.2	The Human Visual System . . . . .	46
4.1.3	Contrast . . . . .	48
4.1.4	Constancy . . . . .	49
4.1.5	Human Visual Perception . . . . .	49
4.1.6	Lightness Perception . . . . .	51
4.2	Perceptually driven rendering . . . . .	52
4.3	Conclusion . . . . .	57

<b>5</b>	<b>Perceptually Motivated Simulation and Virtual Environments</b>	<b>93</b>
5.1	Introduction . . . . .	93
5.2	Perceptually-based Selective Rendering . . . . .	93
5.3	Behavioral Fidelity of Simulations based on Space Memory . . . . .	98
5.4	Investigating Perceptual Sensitivity to Head Tracking Latency . . . . .	102
<b>6</b>	<b>Trends from the 7th annual ACM/Eurographics Symposium on Applied Perception in Graphics and Visualization 2010</b>	<b>107</b>
6.1	APGV Proceedings and Resources . . . . .	107

# 1 Introduction

## 1.1 Motivation

The (re) introduction of 3D cinema, advent of affordable stereoscopic display technology, and seamless integration of real-world scenes with computer graphics fuels our continuing ability to create and display stunning realistic imagery. With the arrival of new technology, algorithms and display methods comes the realization that gains can be made by tailoring output to the intended audience; humans. Human beings have an amazingly complex perceptual systems, which have the ability to quickly capture and process vast amounts of complex data. With all its capability however, the Human Visual System (HVS) has some surprising nuances and limitations that can be exploited to the benefit of numerous graphics applications. This new tutorial will provide insight into those aspects of the HVS and other perceptual systems that can serve as both a guide and yard-stick to further the development and evaluation of computer graphics imagery and presentations. The literature on perception provides a rich source of knowledge that can be applied to the realm of computer graphics for immediate and direct benefit, generating images that not only exhibit higher quality, but use less time and resources to process. In addition, knowledge of the HVS serves as a guide on how best to present the images to fulfill the application at hand.

## 1.2 Course Overview

We will present timely, relevant examples on how researchers have leveraged perceptual information for optimization of rendering algorithms, to better guide design and presentation in (3D stereoscopic) display media, and for improved visualization of complex or large data sets. Each section will provide references and short overviews of cutting-edge current research pertaining to that area. We will ensure that the most up-to-date research examples are presented by sourcing information from recent perception and graphics conferences and journals such as ACM Transactions on Perception, paying particular attention work presented at the 2010 Symposium on Applied Perception in Graphics and Visualization.

## 1.3 Focus Areas

We will focus on four key areas in which perceptual knowledge has been successfully interleaved with computer graphics.

### 1.3.1 Perceptually Motivated 3D Displays & Depth Perception

3D stereoscopic displays are being used in a wide range of fields. To understand how better to present information on such displays, a comprehensive understanding of depth perception is necessary. This area will focus on depth perception and applications of such to image presentation.

### **1.3.2 Perceptually Motivated Visualization**

Discussion of recent research pertaining to psychophysics and application to scientific and information visualization. A closer look at visual attention and visual memory will provide the framework for steering perceptually informed visualizations.

### **1.3.3 Exploitation of the limitations of the HVS to reduce rendering times**

while improving resulting image quality. This includes real-time and non-real time graphics, image quality metrics and high dynamic range imagery.

### **1.3.4 Exploration of incorporating perceptual and cognitive aspects to Virtual Environments (VEs).**

Such principles could be applied to selective real-time rendering algorithms, positive transfer of training as well as to optimizations for latency degradations and predictive tracking.

## **1.4 Summary**

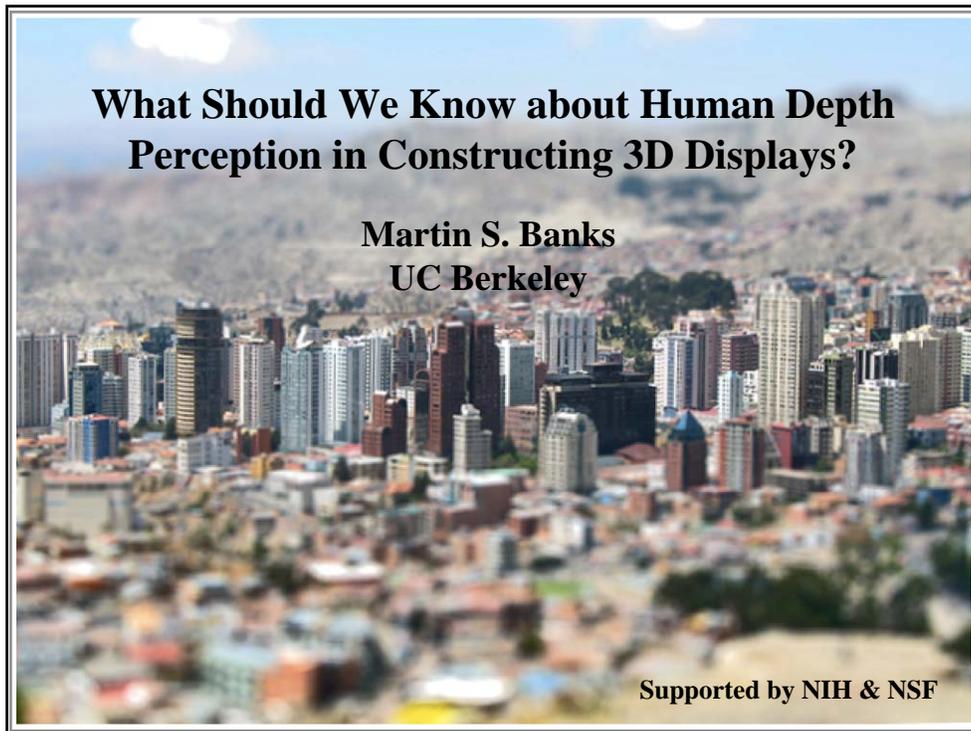
In summary, this course represents a whirlwind tour of insights into how the eye and brain capture and process visual information through our perceptual systems, and how we can use those insights to further advance many areas in computer graphics.

**Martin S. Banks**

## **2 Perceptually Motivated 3D Displays & Depth Perception**

### **2.1 Introduction**

The human visual system has evolved in an environment with constrained relationships between objects and retinal images. That relationship is often altered in stereoscopic displays, so it is important to understand the situations in which the alteration is insignificant and the situations in which it causes undesirable perceptual or ergonomic effects. This section will review the current literature on visual perception and human ergonomics in the context of the viewing of stereo displays. The literature shows that stereo displays can be associated with viewer fatigue/discomfort, reduced visual performance, and distorted 3D perception. This section will also discuss ways to minimize these adverse viewer effects.



## Excitement about Stereo Display Applications

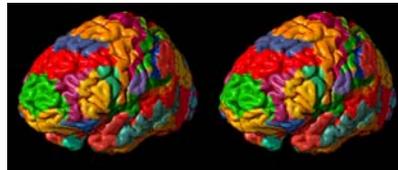
### Entertainment

stereo cinema  
video games



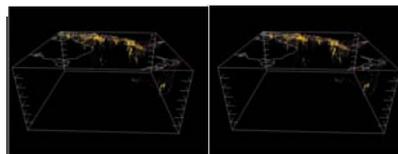
### Medical Imaging

diagnosis  
surgery  
instruction



### Scientific Visualization

geology  
molecular biology



## Problems with Using Stereo Displays

### Technical Issues

- Developing content
- Sufficient resolution over time: “temporal aliasing
- Sufficient separation between two eyes’ images: “ghosting”

### User Issues

- Perceptual distortions due to incorrect viewing position
- Vergence-accommodation conflict: distortion, fatigue
- Treatment of blur

## Problems with Using Stereo Displays

### Technical Issues

- Developing content
- Sufficient resolution over time: “temporal aliasing
- Sufficient separation between two eyes’ images: “ghosting”

### User Issues

- Perceptual distortions due to incorrect viewing position
- Vergence-accommodation conflict: distortion, fatigue
- Treatment of blur

## Viewing Pictures

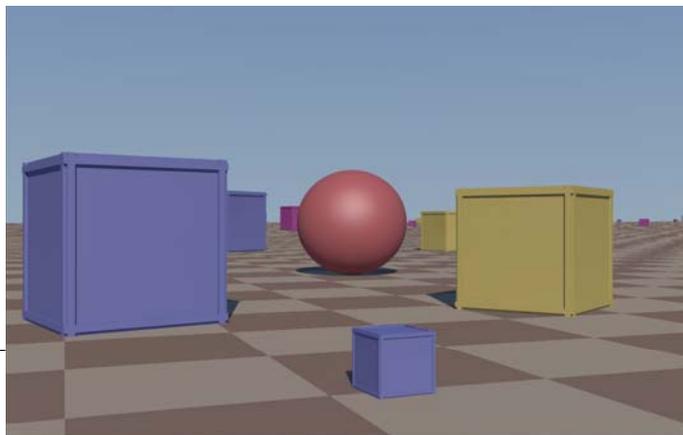
Almost never view pictures from correct position.

Retinal image thus specifies different scene than depicted.

Do people compensate, and if so, how?



## Stimuli



Vishwanath, Girshick, & Banks (2005), *Nature Neuroscience*.

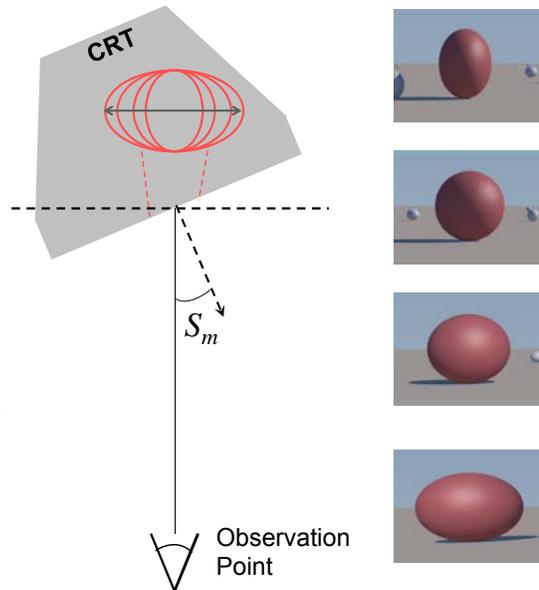
### Experimental Task

Stimulus: simulated 3D ovoid with variable aspect ratio.

Task: adjust ovoid until appears spherical.

Vary monitor slant  $S_m$  to assess compensation for oblique viewing positions.

If compensate, will set ovoid to sphere *on screen* (ellipse on retina).

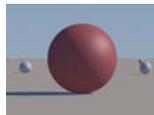


Vishwanath, Girshick, & Banks (2005), *Nature Neuroscience*.

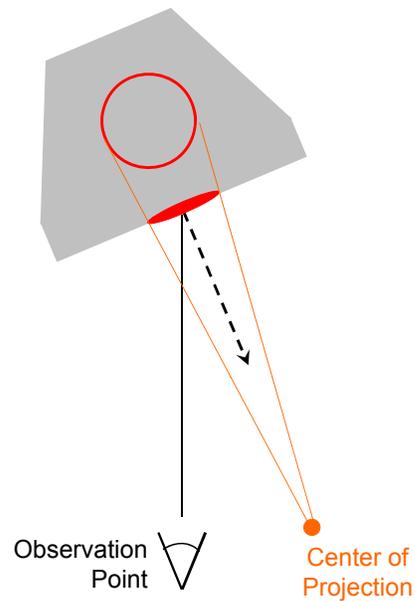
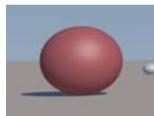
### Predictions

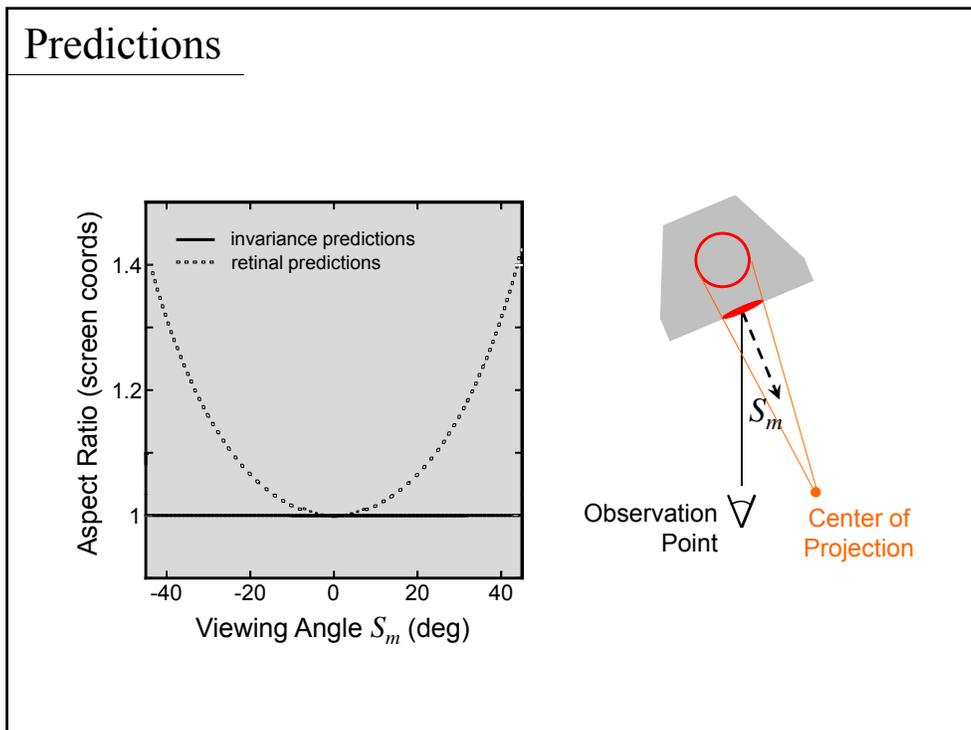
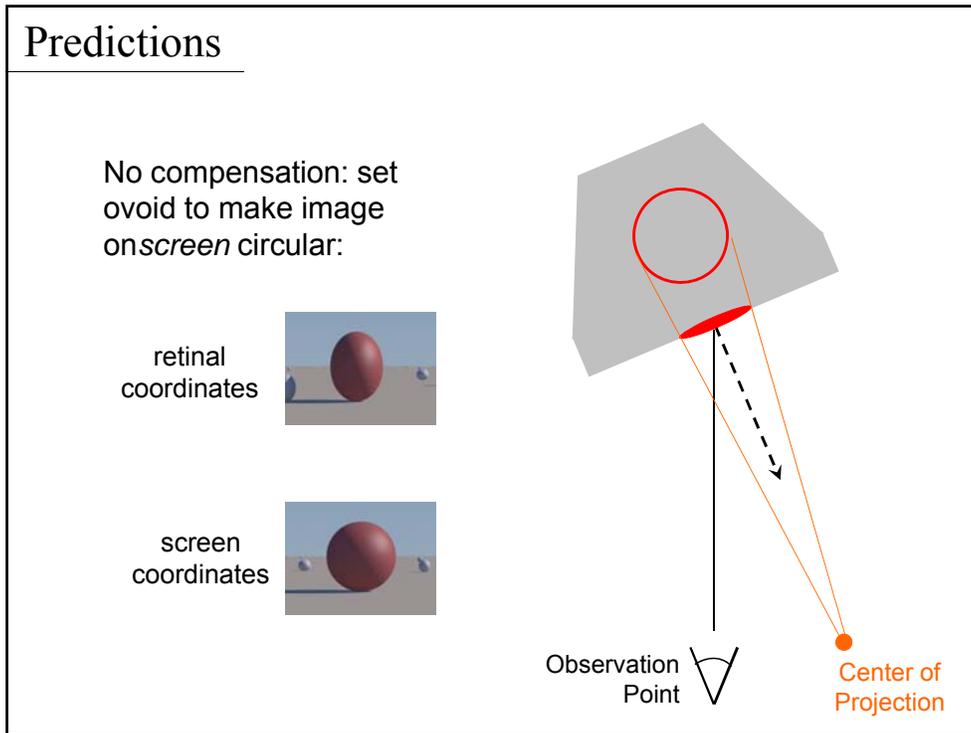
No compensation: set ovoid to make image on *retina* circular:

retinal coordinates

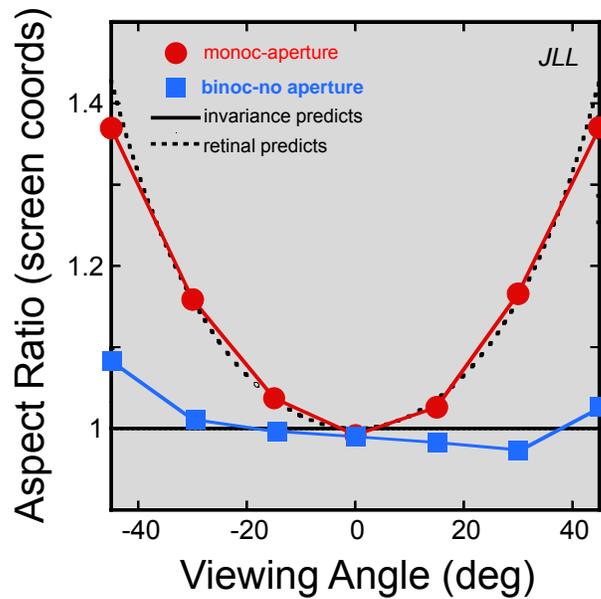


screen coordinates





## Results



Vishwanath, Girshick, & Banks (2005), *Nature Neuroscience*.

## Compensation for Incorrect Viewing Position

- Pictures not useful unless percepts are robust to changes in viewing position.
- People compensate for oblique viewing position when viewing 2d pictures.
- Two theories of compensation: pictorial & surface. Data clearly favor surface compensation.
- Two versions of surface method: global & local. Data clearly favor local slant.

## 2D Pictures vs Stereo Pictures



- Two eyes presented same image
- Binocular disparities specify orientation & distance of picture surface; hence useful for compensation

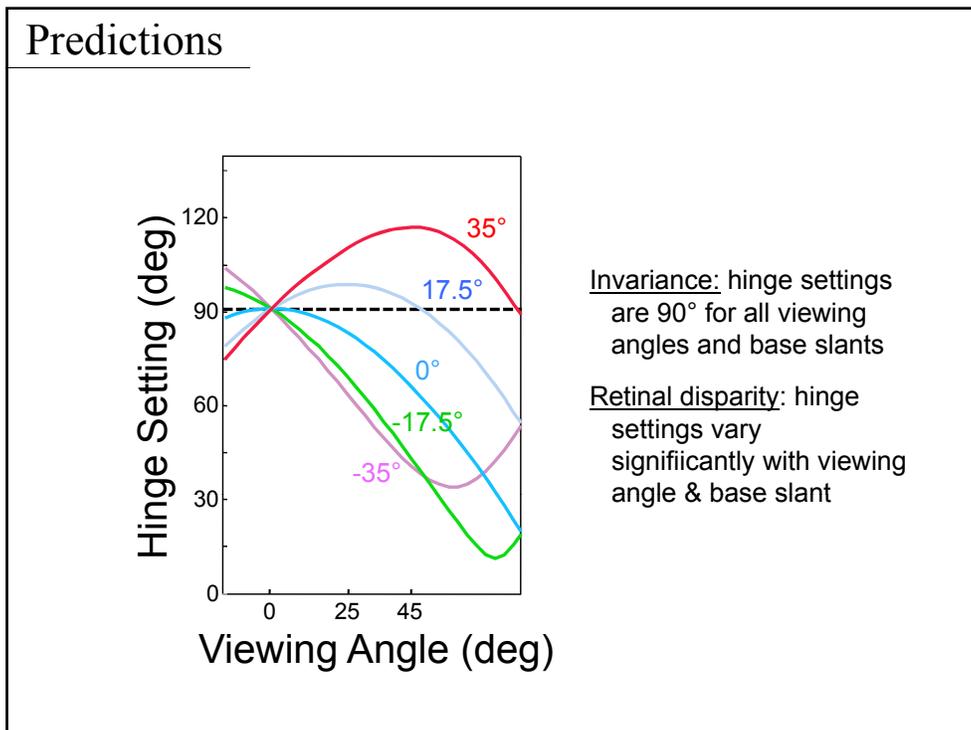
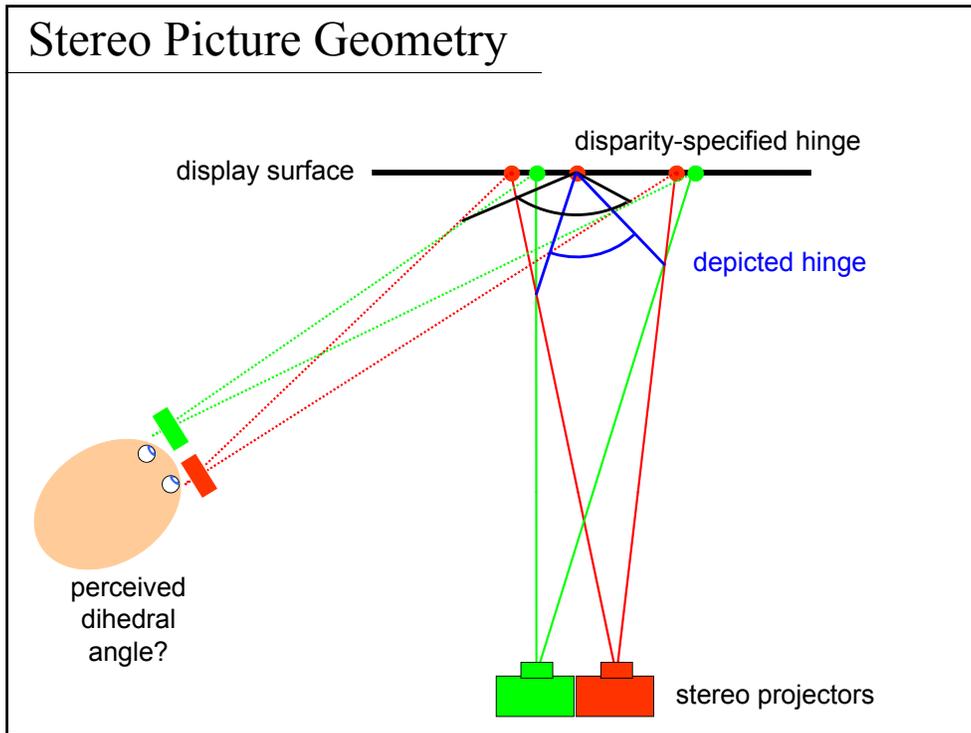


- Two eyes presented different images
- Binocular disparities specify orientation & distance of picture surface *and* layout of picture contents; hence not useful for compensation

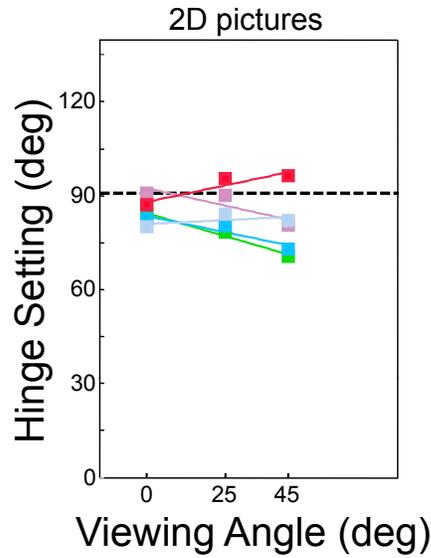
## Stereo Pictures

- For most applications, viewers will not be at correct position.
- Retinal disparities thus specify a different layout than depicted.
- Do people compensate?
- Is correct seating position for a 3D movie more important than for 2D movie?



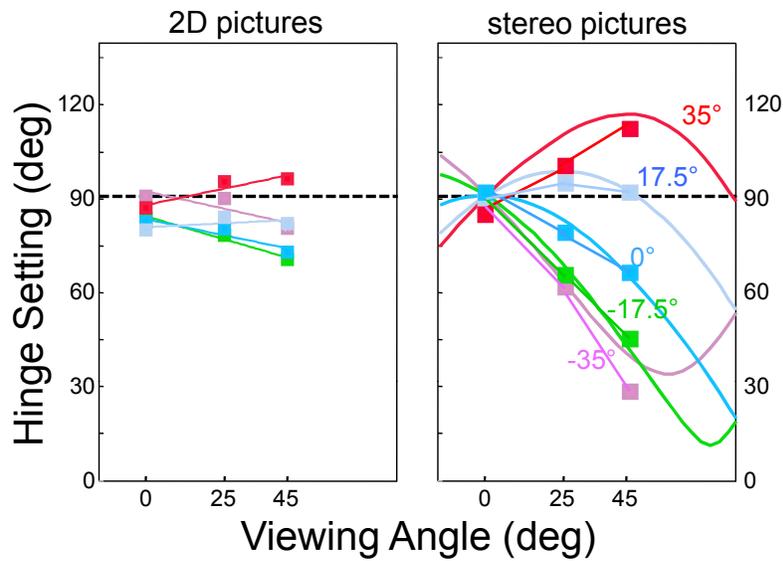


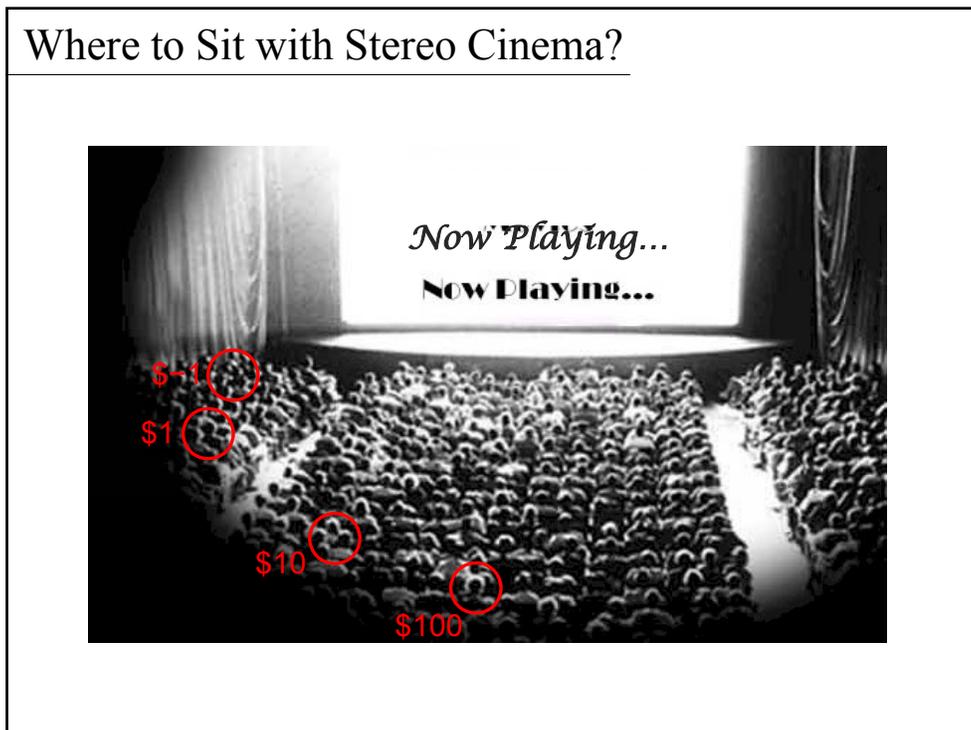
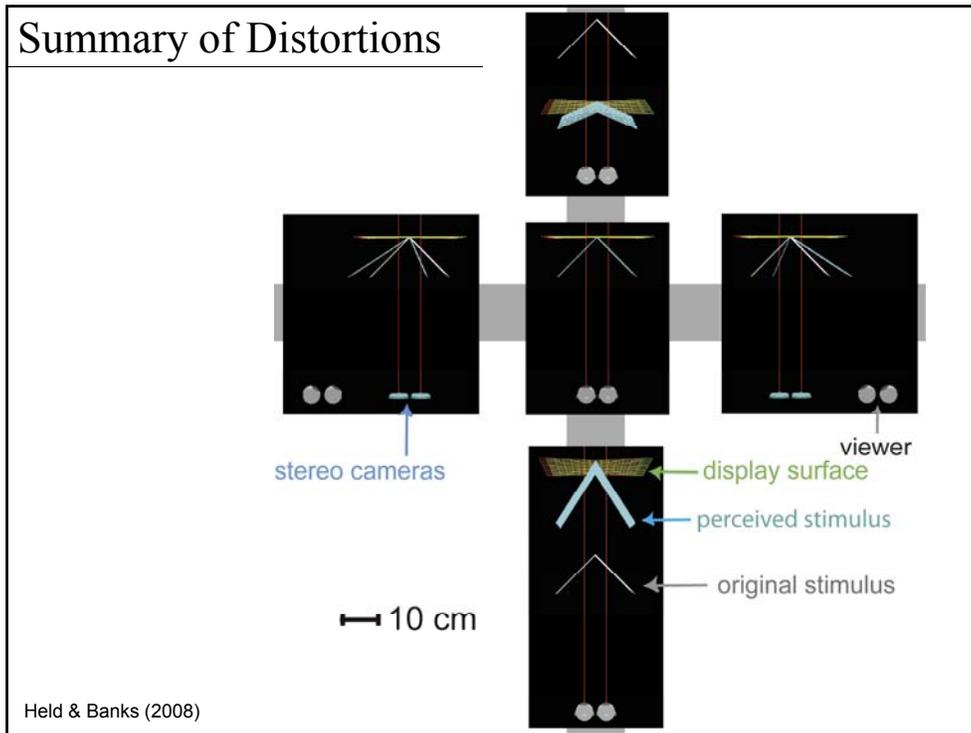
Results: 2D Pictures



Invariance: hinge settings are 90° for all viewing angles and base slants

Results: 2D & Stereo Pictures





## Problems with Using Stereo Displays

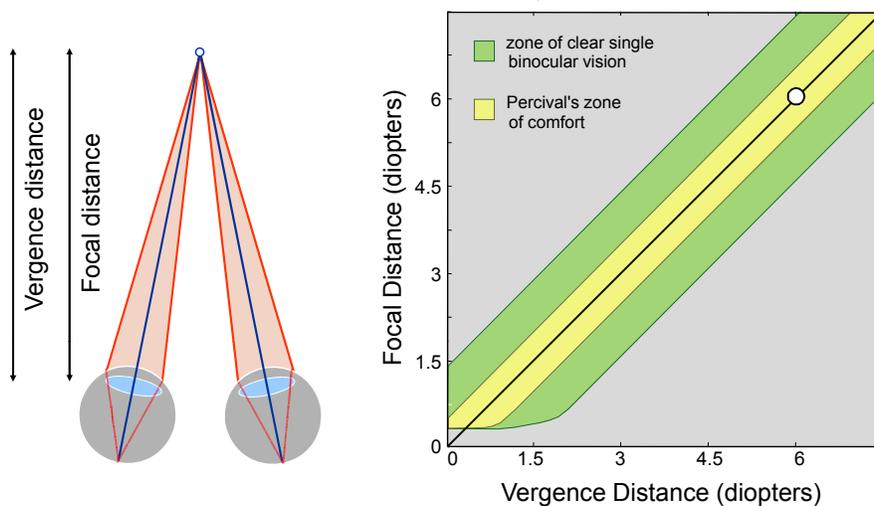
### Technical Issues

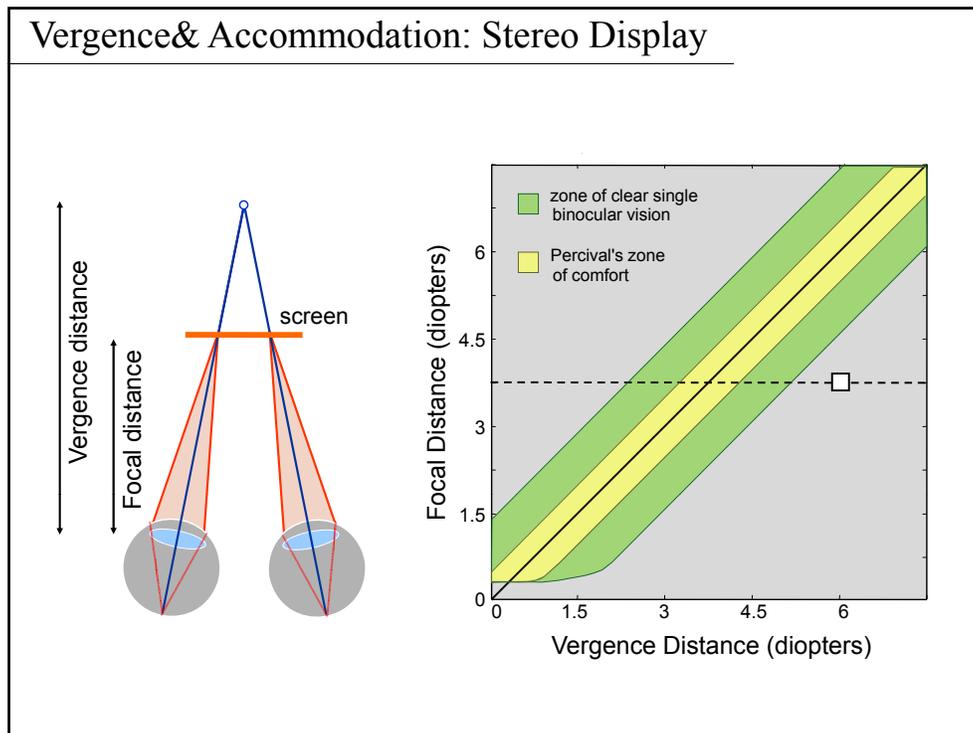
- Developing content
- Sufficient resolution over time: “temporal aliasing”
- Sufficient separation between two eyes’ images: “ghosting”

### User Issues

- Perceptual distortions due to incorrect viewing position
- Vergence-accommodation conflict: distortion, fatigue
- Treatment of blur

## Vergence & Accommodation: Natural Viewing





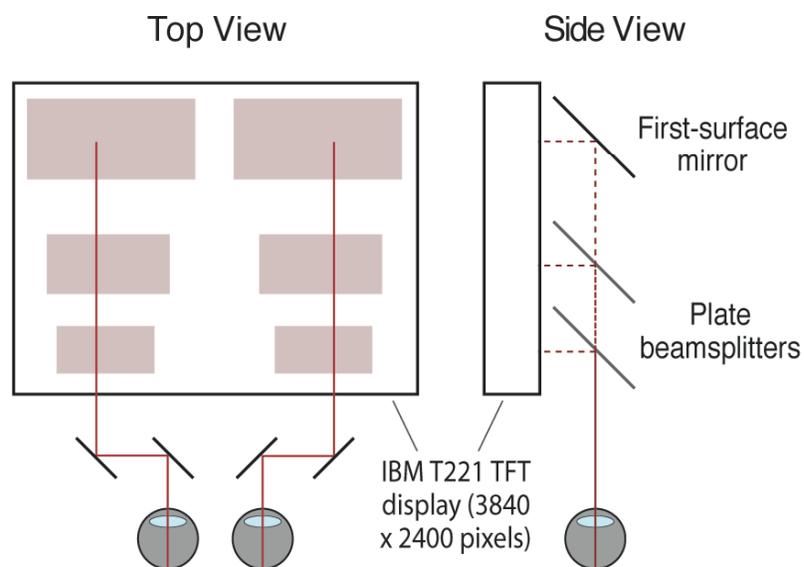
- ### Consequences of Vergence-accommodation Conflicts
- Does accommodation affect 3d shape perception?
  - Is vergence-accommodation conflict the cause of discomfort & fatigue in viewing stereo displays?

## Displays with Nearly Correct Focus Cues

Two multi-focal displays we've developed:

1. Fixed-viewpoint, volumetric display with mirror system & 3 focal planes (Akeley, Watt, Girshick, & Banks, *SIGGRAPH2004*).
2. Fixed-viewpoint, volumetric display with switchable lens & 4 focal planes (Love, Hoffman, Kirby, Hands, Gao, & Banks, *Optics Express*, 2009)

## Multi-focal Display



Akeley, Watt, Girshick & Banks (2004), *SIGGRAPH*.

Multi-focal Display



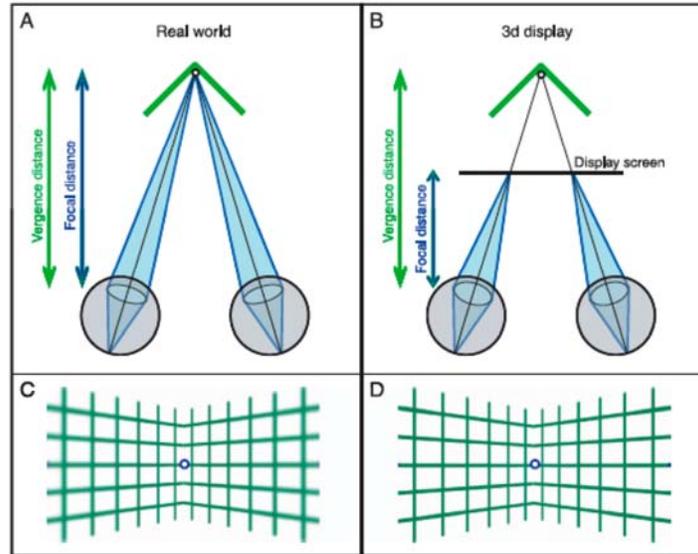
Akeley, Watt, Girshick & Banks (2004), SIGGRAPH.

Multi-focal Display



Akeley, Watt, Girshick & Banks (2004), SIGGRAPH.

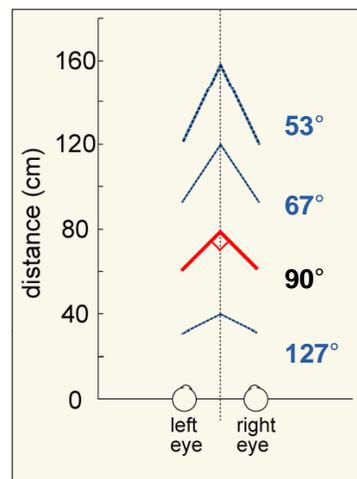
Vergence-accommodation Conflict & Perceived Shape



Watt, Ernst, Akeley, & Banks (2005), *Journal of Vision*

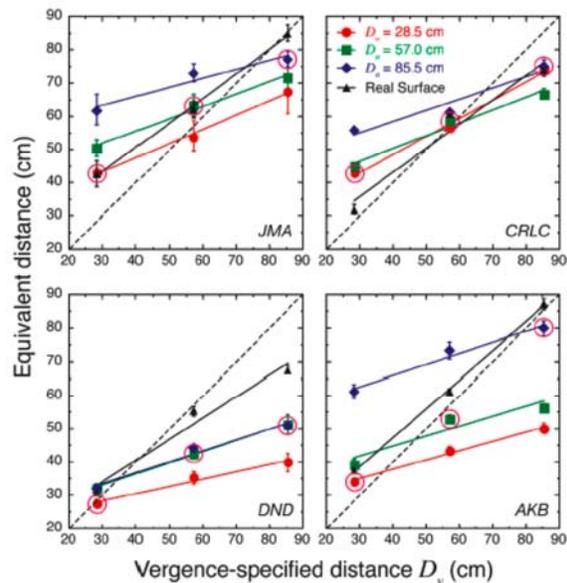
Vergence-accommodation Conflict & Perceived Shape

- For estimating 3d shape, horizontal disparities must be scaled by an estimate of fixation distance
- Information about fixation distance is available from vergence AND accommodation



Watt, Ernst, Akeley, & Banks (2005), *Journal of Vision*

## Vergence-accommodation Conflict &amp; Perceived Shape



- Vergence and focal distance affect estimated distance used to scale disparity
- Thus focal distance affects perceived 3d shape

Watt, Ernst, Akeley, & Banks (2005), *Journal of Vision*

## Problems with Using Stereo Displays

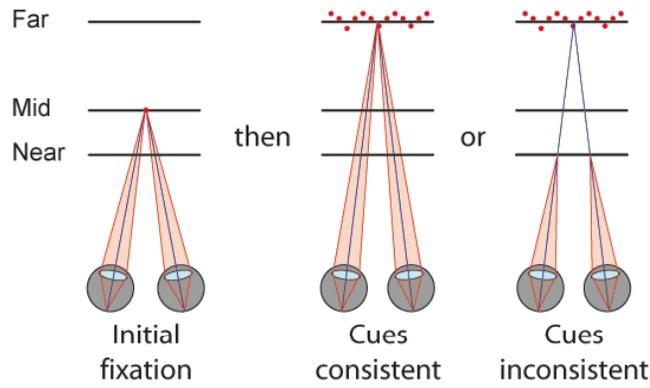
Technical Issues

- Developing content
- Sufficient resolution over time: "temporal aliasing"
- Sufficient separation between two eyes' images: "ghosting"

User Issues

- Perceptual distortions due to incorrect viewing position
- Vergence-accommodation conflict: distortion, fatigue
- Treatment of blur

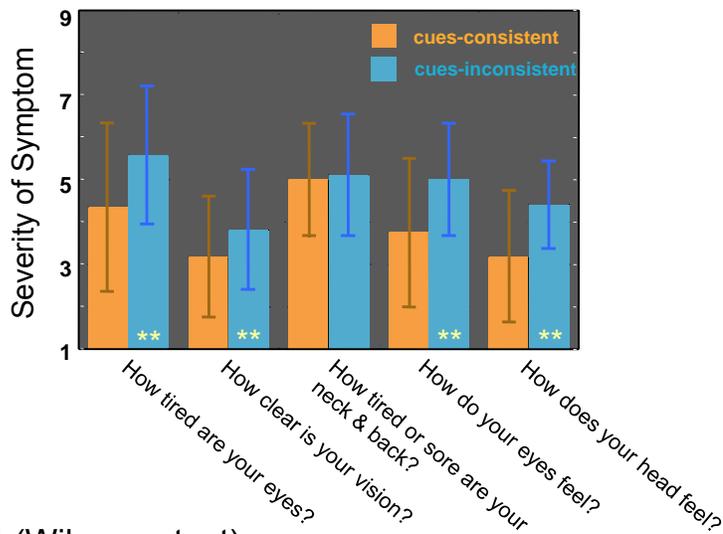
### Fatigue Experiment: Stimuli



- 600-ms stimulus at near or far vergence-specified distance
- Appeared at each focal distance

Hoffman, Girshick, Akeley, & Banks (2008), *Journal of Vision*

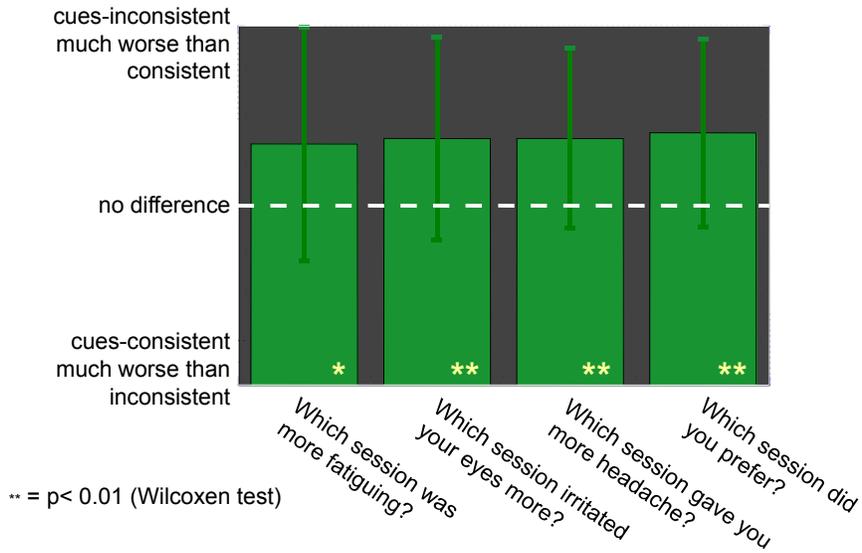
### Fatigue Experiment: Results



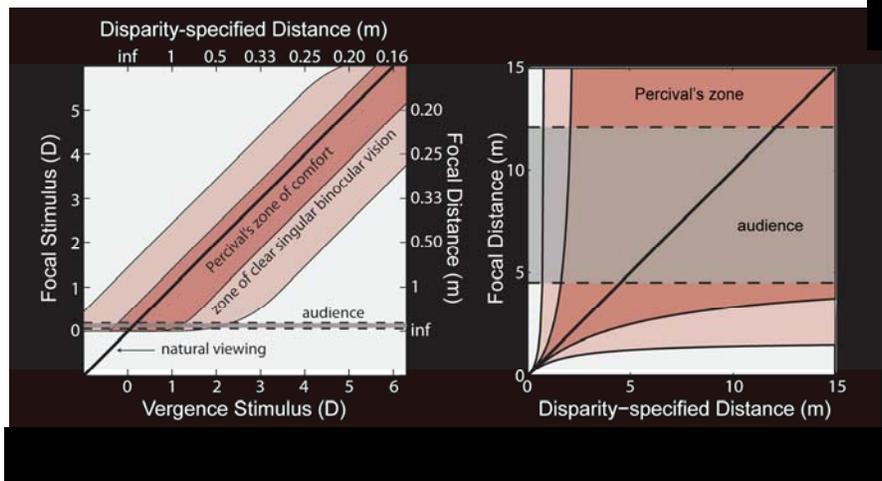
\*\*= p < 0.01 (Wilcoxon test)

Hoffman, Girshick, Akeley, & Banks (2008), *Journal of Vision*

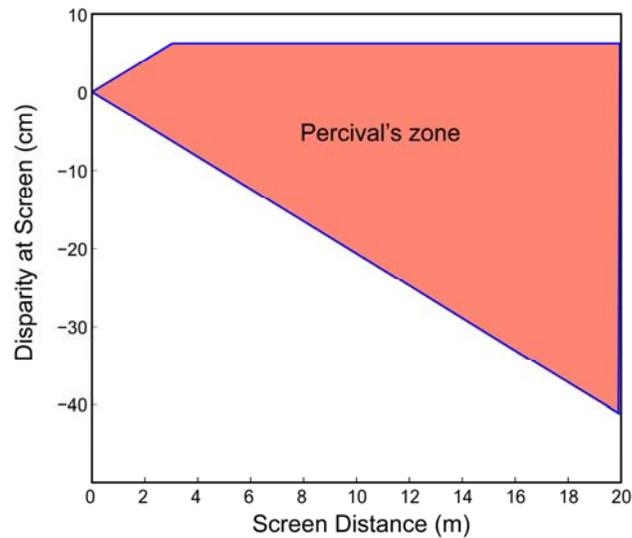
### Fatigue Experiment: Results



Hoffman, Girshick, Akeley, & Banks (2008), *Journal of Vision*



### Discomfort & Stereo Cinema



### Problems with Using Stereo Displays

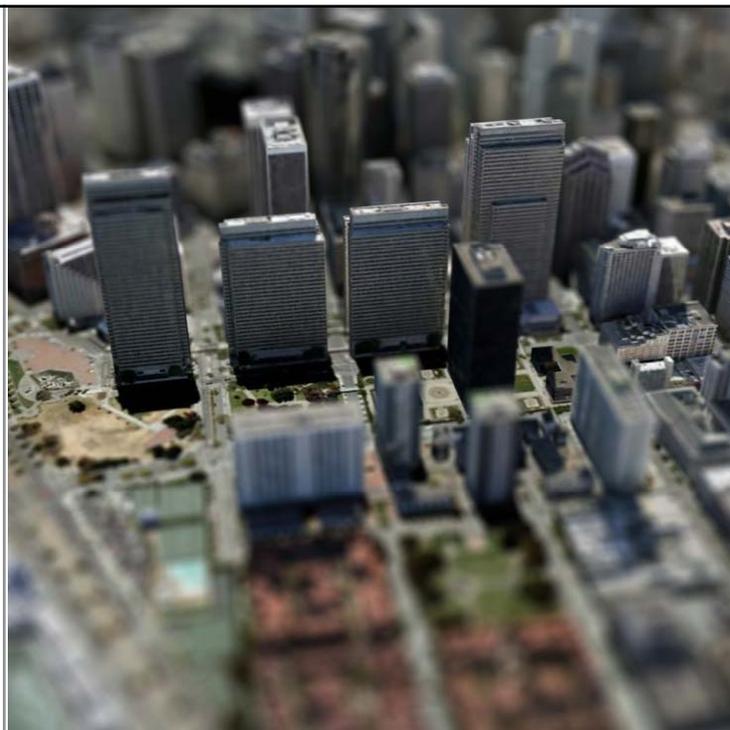
#### Technical Issues

- Developing content
  - Sufficient resolution over time: “temporal aliasing”
- Sufficient separation between two eyes’ images: “ghosting”

#### User Issues

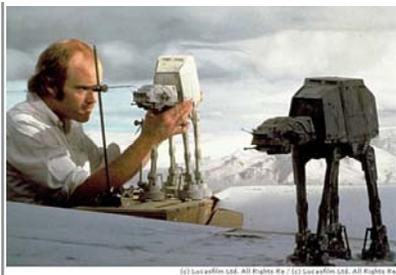
- Perceptual distortions due to incorrect viewing position
- Vergence-accommodation conflict: distortion, fatigue
  - Treatment of blur

Blur as Cue to Absolute Distance





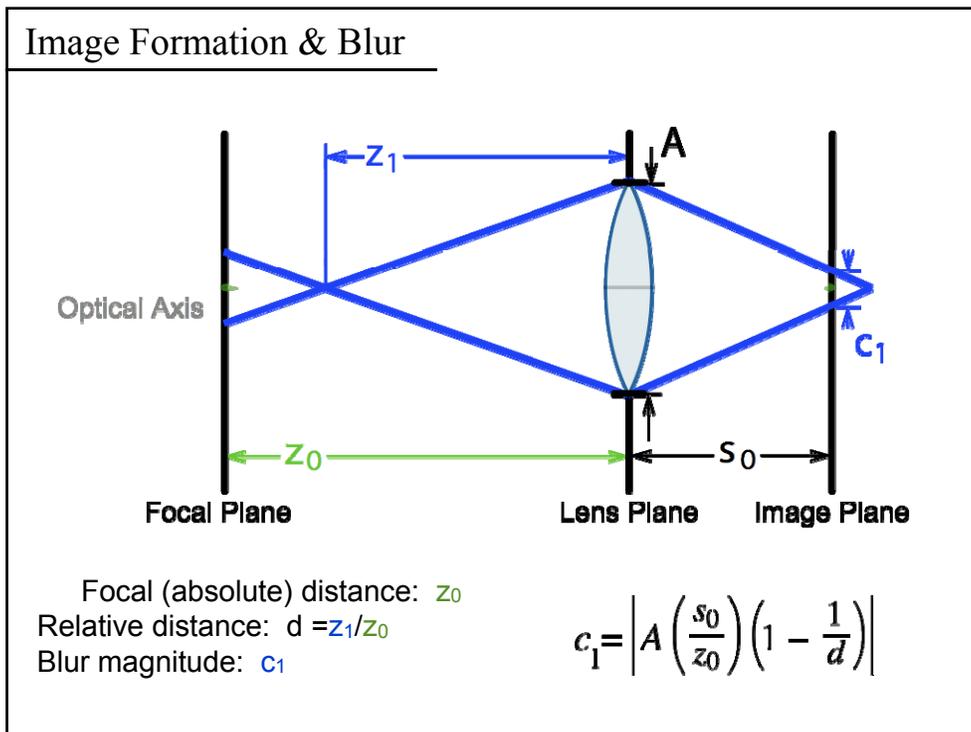
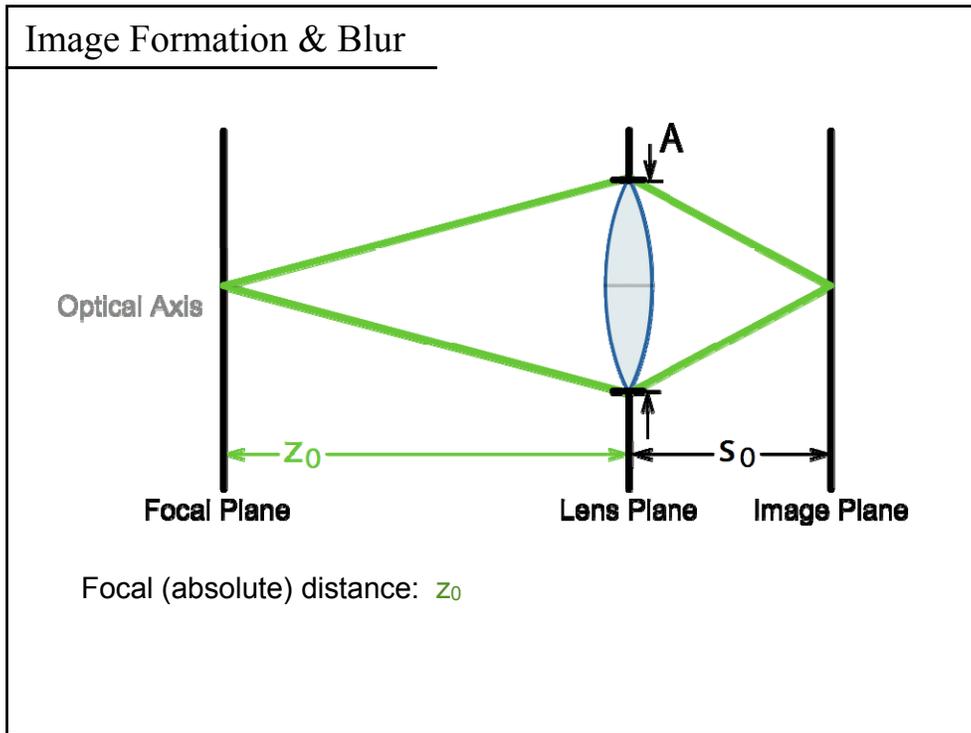
### Blur in Cinematography



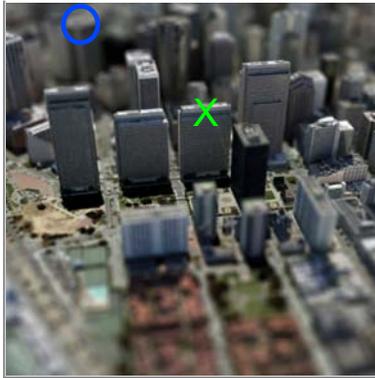
Small camera aperture to increase depth of field & minimize blur

Scale models appear much larger





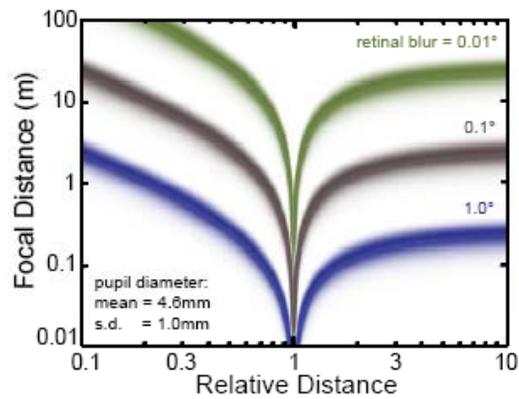
Distance Information from Blur



$$z_0 = \left| \frac{A}{c_l} s_0 \left( 1 - \frac{1}{d} \right) \right|$$

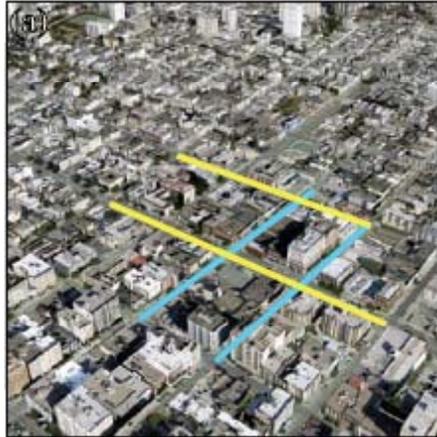
Solve for absolute distance ( $z_0$ ) given blur, aperture, & relative distance ( $d$ )

Distance Information from Blur



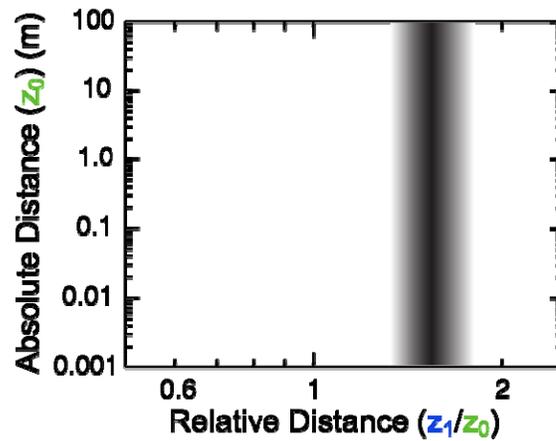
Can only place rough bounds on absolute distance from measurement of blur

### Estimating Relative Distance from Perspective



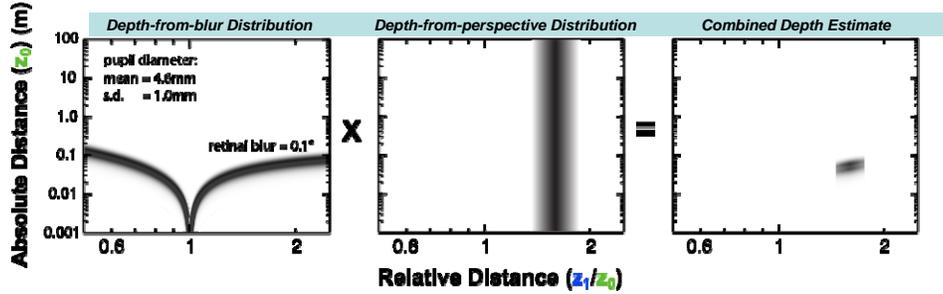
- Grid lines placed on image to determine vanishing points
- Estimate local slant from linear perspective
- Calculate relative distances

### Distance Information from Perspective



Can't estimate absolute distance from perspective

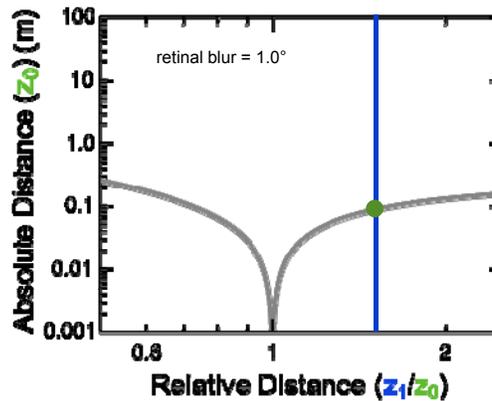
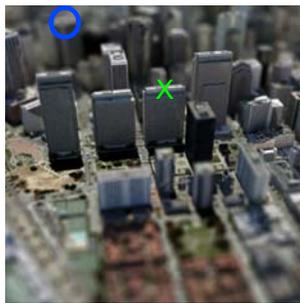
### Probabilistic Model



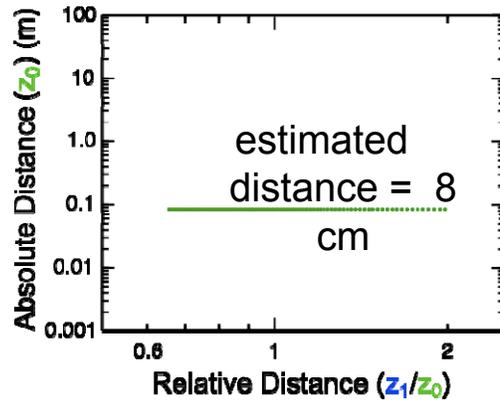
By combining information from blur & perspective, can estimate absolute distance & therefore absolute size

Held, Cooper, O'Brien, & Banks, TOG, 2009

### Estimating Absolute Distance



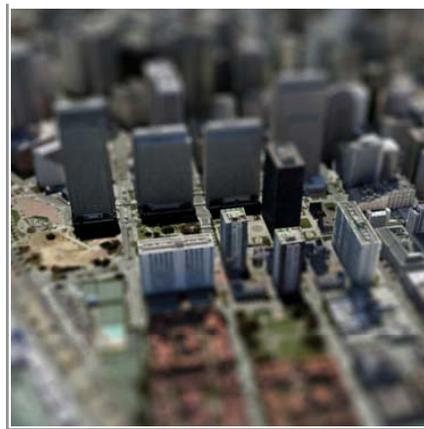
### Estimating Absolute Distance



### Accuracy of Blur-distance Signals

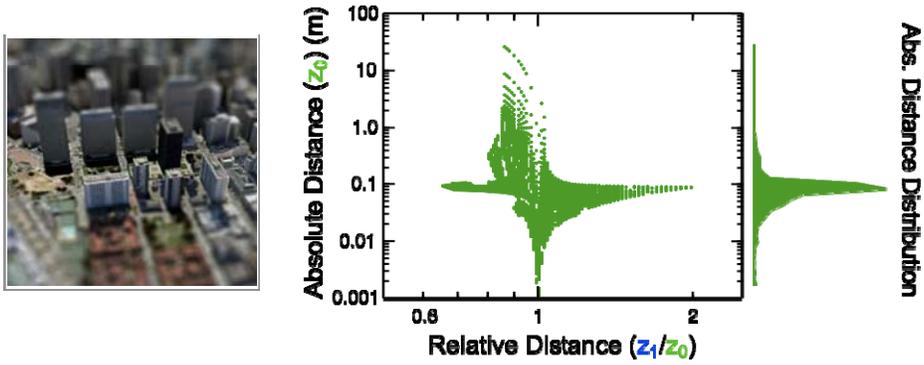


Blur consistent with distance



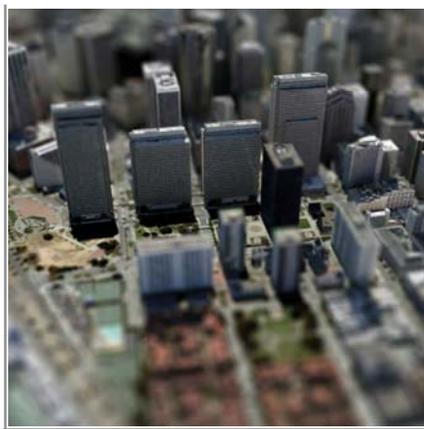
Blur & distance gradients aligned

### Distance Estimate with Aligned Gradients



Estimated distance = ~10 cm

### Accuracy of Blur-distance Signals

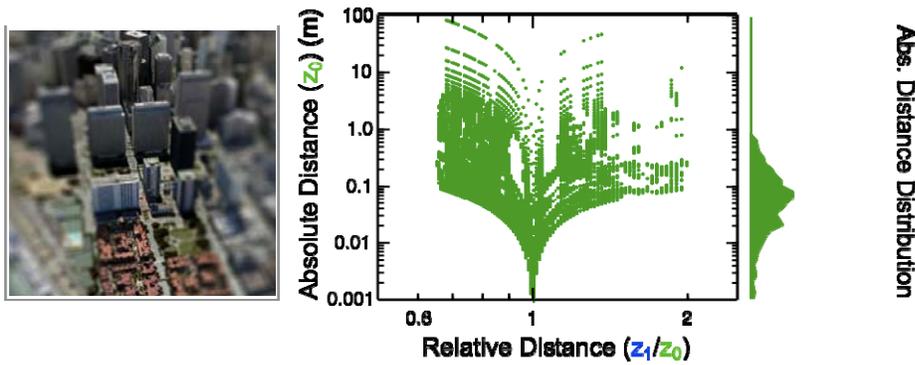


Blur consistent with distance



Blur & distance gradients not aligned

### Distance Estimates with Unaligned Gradients

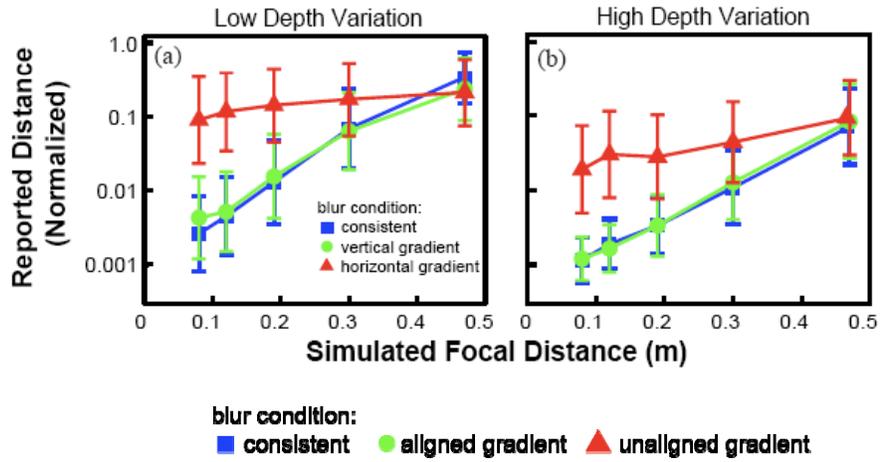


Uncertain distance estimate

### Psychophysical Experiment

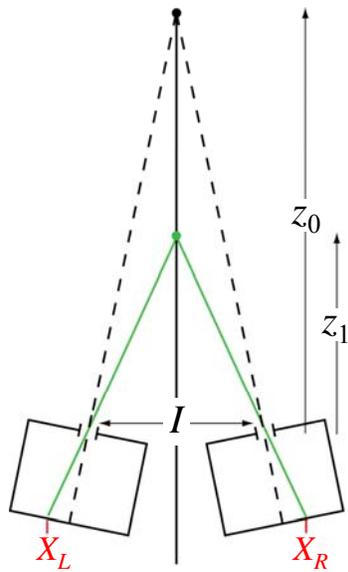
- 7 scenes from GoogleEarth
- Each scene rendered 4 ways: no blur, blur consistent with distance, blur & distance gradients aligned, blur & distance gradients orthogonal
- 5 blur magnitudes
- Naïve subjects viewed each image monocularly for 3 sec
- Reported distance from marked building in image center to the camera that produced the image
- 7 repetitions, random order

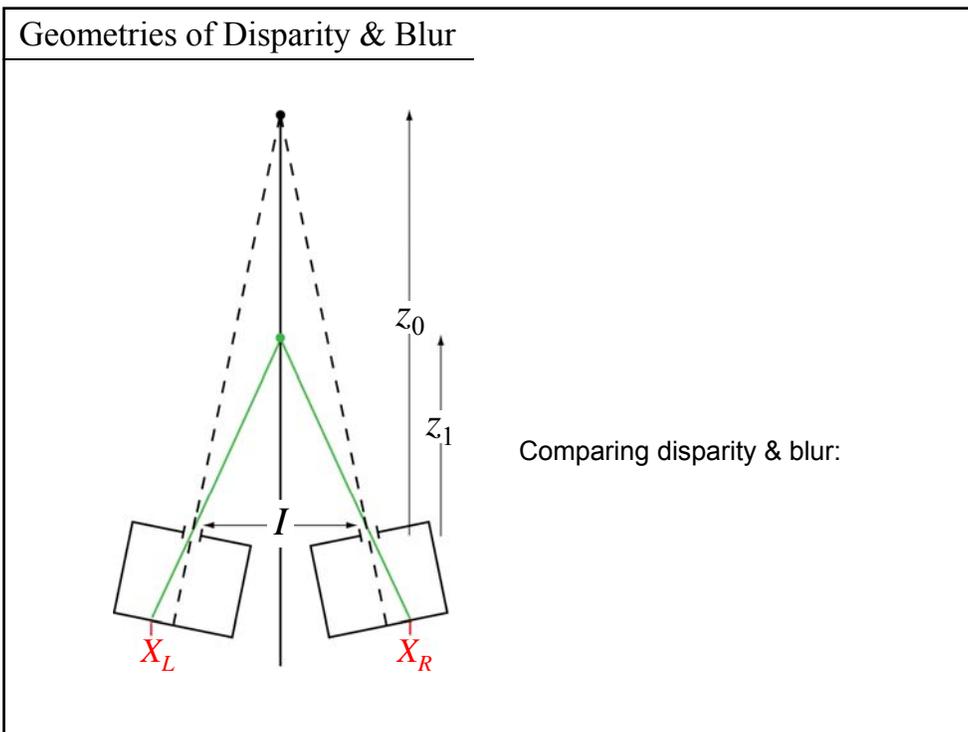
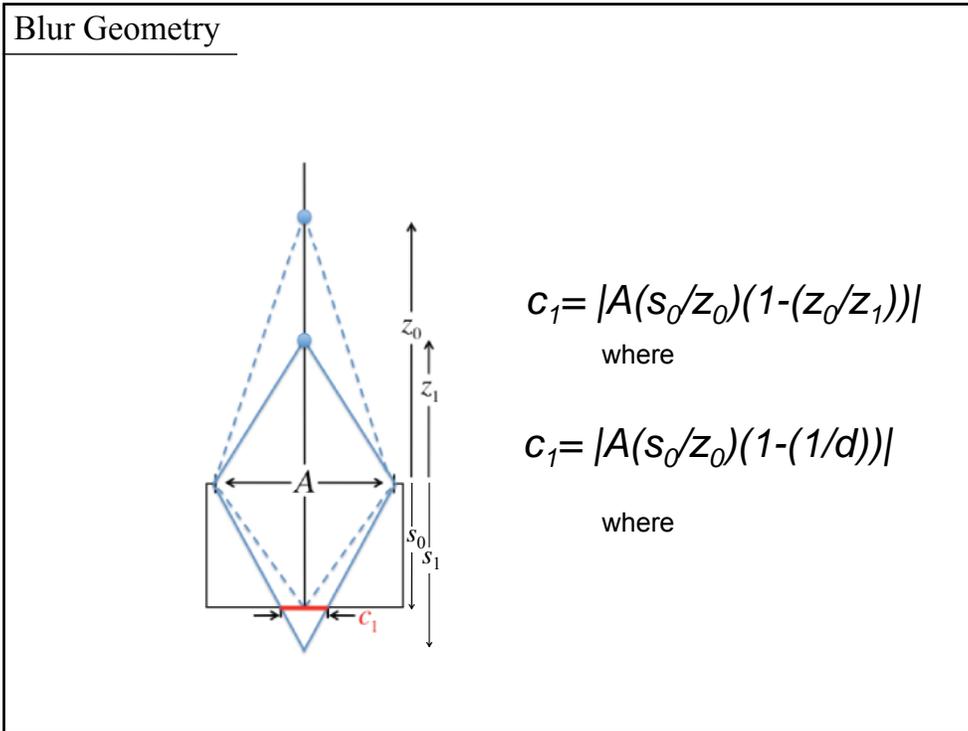
Experimental Results



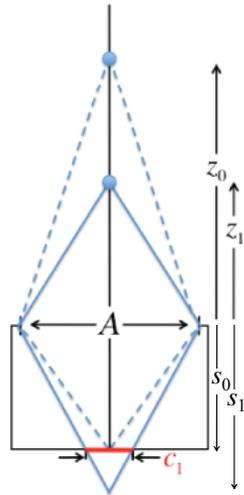
Held, Cooper, O'Brien, & Banks, TOG, 2009

Disparity Geometry





### Natural Depth of Field



### Summary

- Various user issues associated with stereo displays
- Compensation for incorrect viewing position is much less complete with stereo pictures than with conventional pictures
  - Blur & accommodation matter to depth perception
    - Accommodation affects disparity scaling and hence perceived 3d shape
    - Vergence-accommodation conflicts cause visual discomfort/fatigue
    - Blur in combination with other pictorial cues can be effective cue to absolute distance
- Blur and disparity have similar geometries and may provide complimentary depth information in natural viewing

### Participants



Kurt Akeley  
Microsoft



Marty Banks  
UC Berkeley



Emily Cooper  
UC Berkeley



Gordon Love  
Durham Univ.



Ahna Girshick  
NYU



Robin Held  
UC Berkeley



David Hoffman  
UC Berkeley



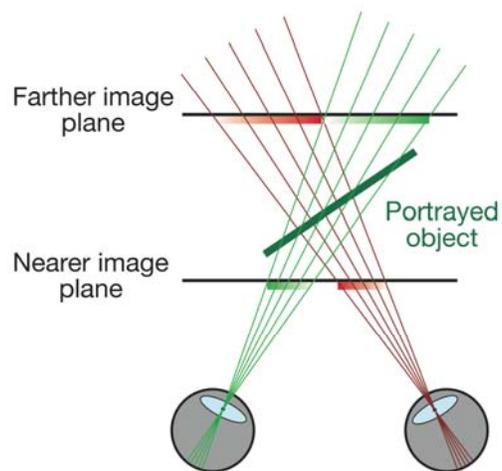
James O'Brien  
UC Berkeley



Simon Watt  
Univ of Wales

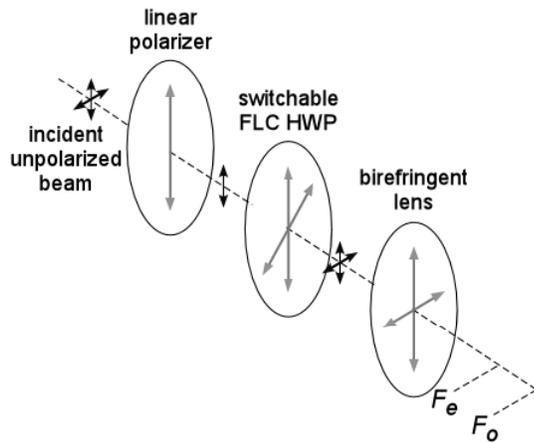
### Depth-weighted Blending

- Depth-weighted blending along lines of sight
- Weights dependent on dioptric distances to planes



Akeley, Watt, Girshick, & Banks (2004), *SIGGRAPH*.

Switchable Lens & Multi-plane Display



Birefringent material has ordinary & extra-ordinary refractive indices depending on polarization  
 Calcite lens has two focal lengths  $F_o$  &  $F_e$   
 Variation in power effected by polarization modulator (FLC)  
 Can stack lenses;  $2^N$  states

Christopher Healey

## 3 Perceptually Motivated Visualization

### 3.1 Introduction

A fundamental goal of visualization is to produce images of data that support visual analysis, exploration and discovery, and identifying novel insights. An important consideration during visualization design is the role of human visual perception [79, 111, 114, 125]. How we “see” details in an image can directly impact a user’s efficiency and effectiveness. This article surveys research on attention and visual perception, with a specific focus on results that have direct relevance to visualization and visual analytics. We discuss theories of low-level visual perception, then show how these findings form a foundation for more recent work on visual memory and visual attention.

### 3.2 Visual Attention and Preattentive Processing

For many years vision researchers have been investigating how the human visual system analyzes images. An important initial result was the discovery of a limited set of visual properties that are detected very rapidly by low-level and fast-acting visual processes. These properties were initially called preattentive, since their detection seemed to precede focused attention. We now know that attention plays a critical role in what we see, even at this early stage of vision. The term preattentive continues to be used, however, since it conveys an intuitive notion of the speed and ease with which these properties are identified.

Typically, tasks that can be performed on large multi-element displays in less than 200–250 milliseconds (msec) are considered preattentive. Eye movements take at least 200 msec to initiate, and random locations of the elements in the display ensure that attention cannot be prefocused on any particular location, yet viewers report that these tasks can be completed with very little effort. This suggests that certain information in the display is “seen” in parallel by low-level visual processes.

A simple example of a preattentive task is the detection of a red circle in a group of blue circles (Fig. 1). The target object has a visual property red that the blue distractor objects do not. A viewer can tell at a glance whether the target is present or absent. Here the visual system identifies the target through a difference in hue, specifically, a red target in a sea of blue distractors. Hue is not the only visual feature that is preattentive. For example, viewers can just as easily find a red circle in a background of red squares. Here, the visual system identifies the target through a difference in curvature (or form).

A unique visual property in the target—a red hue or a curved form—allows it to “pop out” of a display. A conjunction target made up of a combination of non-unique features normally cannot be detected preattentively. For example, consider combining the two backgrounds and searching for a red circle in a sea of blue circles and red squares. The red circle target is made up of two features: red

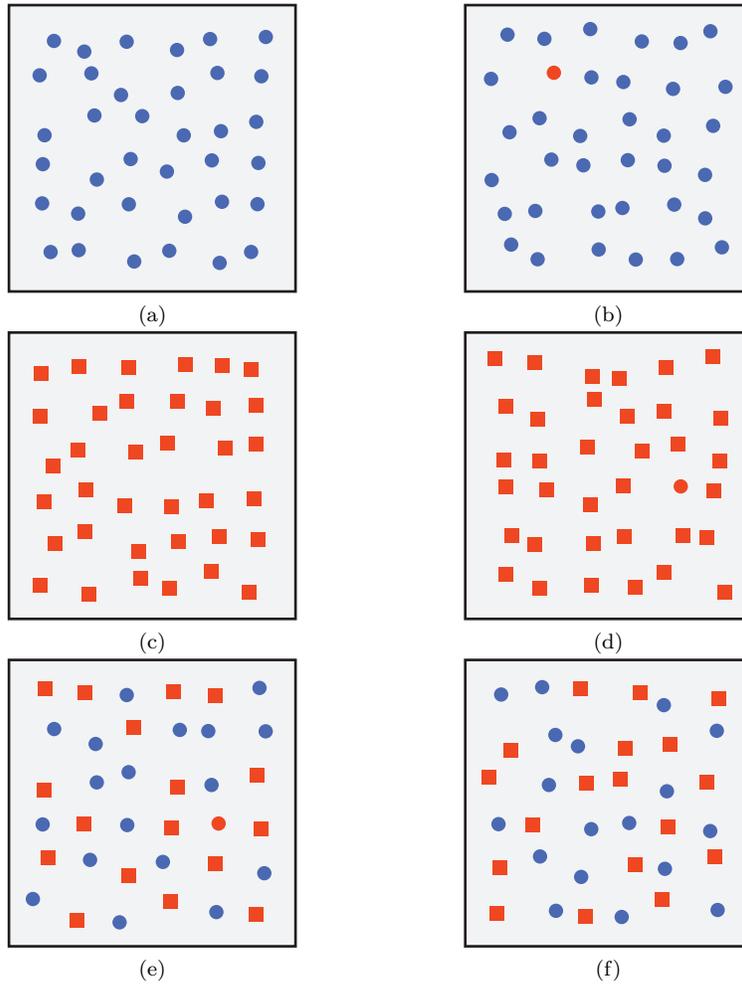


Figure 1: Target detection: (a) hue target red circle absent; (b) target present; (c) shape target red circle absent; (d) target present; (e) conjunction target red circle present; (f) target absent

and circular. One of these features is present in each of the distractor objects—red squares and blue circles. The visual system has no unique visual property to search for when trying to locate the target. A search for red items always returns true because there are red squares in each display. Similarly, a search for circular items always sees blue circles. Numerous studies have shown that a conjunction target cannot be detected preattentively. Viewers must perform a time-consuming serial search through the display to confirm its presence or absence.

If low-level visual processes can be harnessed during visualization, it can

draw attention to areas of potential interest in a display. This cannot be accomplished in an ad-hoc fashion, however. The visual features assigned to different data attributes—the data-feature mapping—must take advantage of the strengths of our visual system, must be well-suited to the analysis needs of the viewer, and must not produce visual interference effects (e.g., conjunction search) that could mask information.

### 3.3 Theories of Preattentive Processing

A number of theories have been proposed to explain how preattentive processing occurs within the visual system: feature integration, textons, guided search, and boolean maps. We provide an overview of these theories, then discuss briefly feature hierarchies, which describes situations where the visual system favors certain visual features over others, and ensemble coding, which shows that viewers can generate summaries of the distribution of visual features in a scene, even when they are unable to locate individual elements based those same features.

#### 3.3.1 Feature Integration

Anne Treisman was one of the original researchers to document the area of preattentive processing [117, 115, 116]. In order to explain the phenomena, Treisman proposed a model low-level human vision made up of a set of feature maps and a master map of locations. Each feature map registers activity for a specific visual feature. Treisman suggested a manageable number of feature maps, including one for each of the opponent colors, as well as separate maps for orientation, shape, and texture. When the visual system first sees an image, all the features are encoded in parallel into their respective maps. A viewer can access a particular map to check for activity, and perhaps to determine the amount of activity. The individual feature maps give no information about location, spatial arrangement, or relationships to activity in other maps, however.

#### 3.3.2 Textons

Bela Julesz was also instrumental in expanding our understanding of what we “see” in an image. Julesz initially focused on statistical analysis of texture patterns [55, 56, 57, 58, 59]. His goal was to determine whether variations in a particular order statistic were detected by the low-level visual system, for example contrast—a first-order statistic—orientation and regularity—a second-order statistic—and curvature—a third-order statistic. Based on these findings, Julesz suggested that the early visual system detects a group of features called textons, which fall into three general categories:

1. Elongated blobs—line segments, rectangles, or ellipses—with specific properties of hue, orientation, width, and so on.
2. Terminators—ends of line segments.

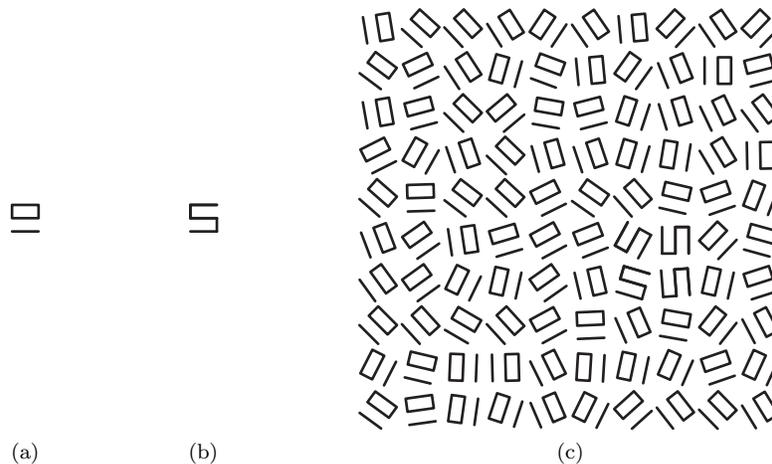


Figure 2: Textons: (a,b) two textons  $A$  and  $B$  that appear different in isolation, but have the same size, number of terminators, and join points; (c) a target group of  $B$ -textons is difficult to detect in a background of  $A$ -textons when random rotation is applied

### 3. Crossings of line segments

Julész believed that only a difference in textons or in their density could be detected preattentively (Fig. 2). No positional information about neighboring textons is available without focused attention. Like Treisman, Julész suggested that preattentive processing occurs in parallel and focused attention occurs in serial.

#### 3.3.3 Guided Search

More recently, Jeremy Wolfe has proposed a theory that he calls “guided search.” He hypothesized that an activation map based on both bottom-up and top-down information is constructed during visual search. Attention is drawn to peaks in the activation map that represent areas in the image with the largest combination of bottom-up and top-down influence [131, 132, 130].

As with Treisman, Wolfe believes early vision divides an image into individual feature maps. In his theory, there is one map for each feature type—a color map, an orientation map, and so on. Within each map a feature is filtered into multiple categories. Bottom-up activation follows feature categorization. It measures how different an element is from its neighbors. Top-down activation is a user-driven attempt to find items with a specific property or set of properties. The activation map is a combination of bottom-up and top-down activity. Hills in the activation map mark regions that generate relatively large amount of bottom-up or top-down influence, but without providing information about the source of a hill. A subject’s attention is drawn from hill to hill in order of decreasing activation.

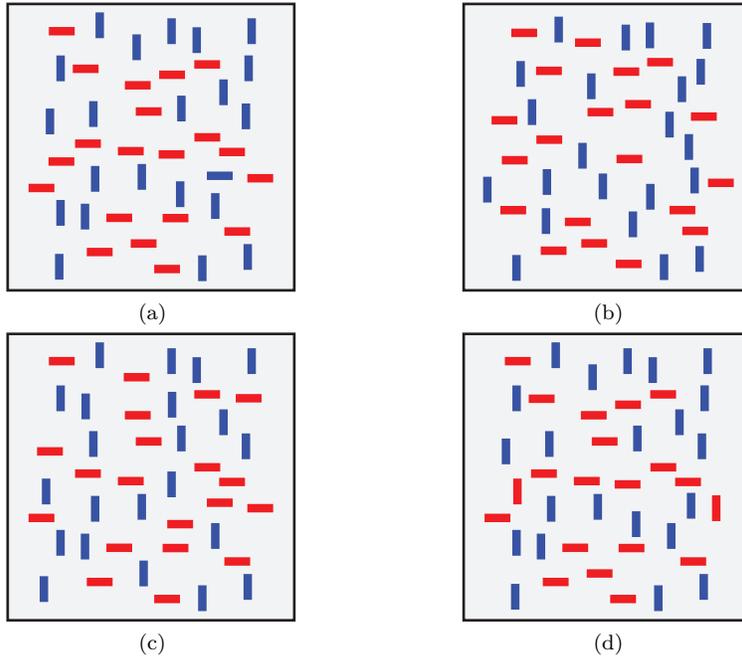


Figure 3: Conjunction search with boolean maps: (a–b) blue horizontal target, select “blue” objects, then search within for a horizontal target, present in (a), absent in (b); (c–d) red vertical target, select “red” objects, then search within for a vertical target, absent in (c), present in (d)

### 3.3.4 Boolean Maps

A more recent model of low-level vision has been presented by Huang et al. [52, 53]. This theory carefully divides visual search into two parts: selection and access. Selection involves choosing a set of objects from a scene. Access determines what properties of the selected objects a viewer can apprehend. Although both operations are implicitly present in previous theories, they are often described as a whole and not as separate steps.

Huang et al. suggest that the visual system can divide a scene into exactly two parts: selected elements and excluded elements. This is the “boolean map” that underlies their theory. The visual system can then access certain properties of the selected elements in the map. Once a boolean map is created, two properties are available to a viewer: the label for any feature in the map, and the spatial location of the selected elements (Fig. 3). Boolean maps can be created in two ways. First, a viewer can specify a single value of an individual feature to select all objects that contain that feature. Second, union or intersection can be applied to two existing maps. In either case, only the result is retained, since evidence suggests that a viewer can only hold and access one boolean map at a time. Viewers can chain these operations together to search for targets in a

fairly complex scene.

### 3.3.5 Ensemble Coding

Existing characterizations of preattentive vision have focused on how low level-visual processes can be used to guide attention to specific location or object in a larger scene. An equally important characteristic of low-level visual processes is their ability to generate a quick summary of how simple visual features are distributed across the field of view. The ability of humans to register a rapid and in-parallel summary of a scene in terms of its simple features was first reported by Ariely [4]. He demonstrated that observers could extract the average size of a large number of dots from only a single glimpse at a display. Yet, when observers were tested on the same displays and asked to indicate whether a single dot of a given size was present, they were unable to do so. This suggests that there is a preattentive mechanism that records summary statistics of visual features without retaining information about the constituent elements that generated the summary.

This ability to rapidly identify scene-based averages may offer important advantages in certain visualization environments. For example, given a stream of real-time data, ensemble coding would allow viewers to observe the stream at a high frame rate, yet still identify individual frames with interesting distributions of visual features (i.e. attribute values). Ensemble coding would also be critical for any situation where viewers want to estimate the amount of a particular data attribute in a display. These capabilities were hinted at in a paper by Healey et al., but without the benefit of ensemble coding as a possible explanation.

### 3.3.6 Feature Hierarchies

One promising strategy for multidimensional visualization is to assign different visual features to different data attributes. This allows multiple data values to be shown simultaneously in a single image. A key requirement of this method is a data-feature mapping that does not produce visual interference. One example of interference is a conjunction target. Another example is the presence of feature hierarchies that appears to exist in the visual system. For certain tasks one visual feature may be “more salient” than another. Researches in psychophysics and visualization have demonstrated a hue-shape hierarchy: the visual system favors color over shape [15, 16, 17, 48, 49]. Background variations in hue interfere with a viewer’s ability to identify the presence of individual shapes and the spatial patterns they form. If hue is held constant across the display, these same shape patterns are immediately visible. The interference is asymmetric: random variations in shape have no effect on a viewer’s ability to see color patterns. Similar luminance-hue and hue-texture hierarchies have also been identified.

## 3.4 Visual Memory

Preattentive processing asks in part: “What visual properties draw our eyes, and therefore our focus of attention to a particular object in a scene?” An equally interesting question is: “What do we remember about an object or a scene when we stop attending to it and look at something else?” Many viewers assume that as we look around us we are constructing a high-resolution, fully detailed description of what we see. Researchers in psychophysics have known for some time that this is not true. In fact, in many cases our memory for detail between glances at a scene is very limited. Evidence suggests that a viewer’s current state of mind can play a critical role in determining what is seen and what is not.

We present three theories that demonstrate and attempt to explain this phenomena: change blindness, inattention blindness, and attentional blink. Understanding what we remember as we focus on different parts of a visualization is critical to designing visualizations that encourage locating and retaining the information that is most important to the viewer.

### 3.4.1 Change Blindness

New research in psychophysics has shown that an interruption in what is being seen—a blink, an eye saccade, or a blank screen—renders us “blind” to significant changes that occur in the scene during the interruption [104, 69, 105, 110]. This change blindness phenomena can be illustrated using a task similar to one shown in comic strips for many years. A viewer is shown two pairs of images. A number of significant differences exists between the images. Many viewers have a difficult time seeing any difference and often have to be coached to look carefully to find it. Once they discover it, they realize that the difference was 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.

The presence of change blindness has important implications for visualization. The images we produce are normally novel for our viewers, so prior expectations cannot be used to guide their analyses. Instead, we strive to direct the eye, and therefore the mind, to areas of interest or importance within a visualization. This ability forms the first step towards enabling a viewer to abstract details that will persist over subsequent images.

### 3.4.2 Inattention Blindness

A related phenomena called inattention blindness suggests that viewers fail to perceive objects or activities that occur outside of the focus of attention [69]. This phenomena is illustrated through an experiment conducted by Neisser [90, 109]. His experiment superimposed video streams of two basketball games. Players wore white shirts in one stream and black shirts in the other. Subjects



(a)



(b)

Figure 4: Change blindness, a major difference exists between the two images

attended to one team—either white or black—and ignored the other. Whenever the subject’s team made a pass, they were told to press a key. After about 30 seconds of video, a third stream was superimposed showing a woman walking through the scene with an open umbrella. The stream was visible for about 4 seconds, after which another 25 seconds of basketball video was shown. Following the trial, only six of twenty-eight naive observers reported seeing the woman. When subjects only watched the screen and did not count passes, 100% noticed the woman.

Additional issues with relevance to visualization are also being investigated. Most et al. are studying the relationship between inattentive blindness and attentional capture, the ability of an object to draw the focus of attention without a viewer’s active participation. Researchers are also studying how perceptual load affects inattentive blindness. Finally, results suggest meaningful objects (e.g., a person’s name or a happy face icon) may be easier to notice.

### 3.4.3 Attentional Blink

In each of the previous methods for studying visual attention, the primary emphasis is on how human attention is limited in its ability to represent the details of a scene (change blindness) and in its ability to represent multiple objects at the same time (inattention blindness). But attention is also severely limited in its ability to process information that arrives in quick succession, even when that information is presented at a single location in space. The attentional blink paradigm is currently the most widely used method to study the availability of attention across time. Its name—“blink”—derives from the finding that when two targets are presented in rapid succession, the second of the two targets cannot be detected or identified when it appears within approximately 100–500 msec following the first target [14, 101]. This suggests that that attention operates over time like a window or gate, opening in response to finding a visual item that matches its current criterion or template and then closing shortly thereafter to consolidate that item as a distinct object or event from others. The attentional blink is an index of the “dwell-time” needed to consolidate a rapidly presented visual item into visual short term memory.

## 3.5 Conclusions

This presentation surveys past and current theories of low-level visual perception and visual attention. Initial work in preattentive processing identified basic visual features that can implicitly or explicitly capture a viewer’s focus of attention. More recent work has extended this to study limited visual memory for change—change blindness and attentional blink—and being “blind” to objects that are outside the focus of attention—inattention blindness. Each of these phenomena have significant consequences for visualization. We strive to produce images that are salient and memorable, and that guide attention to locations of importance within the data. Understanding what the visual seems to see and does not see is critical to designing effective visual displays.

**Ann McNamara**

## **4 Perceptually Motivated Rendering**

### **4.1 Visual Perception in Realistic Image Synthesis**

Realism is often a primary goal in computer graphics imagery. We strive to create images that are perceptually indistinguishable from an actual scene. Rendering systems can now closely approximate the physical distribution of light in an environment. However, physical accuracy does not guarantee that the displayed images will have an authentic visual appearance. In recent years the emphasis in realistic image synthesis has begun to shift from the simulation of light in an environment to images that look as real as the physical environment they portray. In other words the computer image should be not only physically correct but also perceptually equivalent to the scene it represents. This implies aspects of the Human Visual System (HVS) must be considered if realism is required. Visual perception is employed in many different guises in graphics to achieve authenticity [92, 7]. Certain aspects of the HVS must be considered to identify the perceptual effects that a realistic rendering system must achieve in order to effectively reproduce a similar visual response to a real scene. This section outlines the main characteristics of the HVS and the manner in which knowledge about visual perception is increasingly appearing in state-of-the-art realistic image synthesis. Perception driven rendering algorithms are described, which focus on embedding models of the HVS directly into global illumination computations in order to improve their efficiency.

#### **4.1.1 Visual Perception**

Perception is the process by which humans, and other organisms, interpret and organize sensation in order to understand their surrounding environment. Sensation refers to the immediate, relatively unprocessed result of stimulation of sensory receptors. Perception, on the other hand, is used to describe the ultimate experience and interpretation of the world and usually involves further processing of sensory input. Sensory organs translate physical energy from the environment into electrical impulses processed by the brain. In the case of vision light, in the form of electromagnetic radiation, activates receptor cells in the eye triggering signals to the brain. These signals are not understood as pure energy, rather, perception allows them to be interpreted as objects, events, people and situations.

#### **4.1.2 The Human Visual System**

Vision is a complex process that requires numerous components of the human eye and brain to work together. Vision is defined as the ability to see the features of objects we look at, such as color, shape, size, details, depth, and contrast. Vision begins with light rays bouncing off the surface of objects. These reflected

light rays enter the eye and are transformed into electrical signals. Millions of signals per second leave the eye via the optic nerve and travel to the visual area of the brain. Brain cells then decode the signals providing us with sight. The response of the human eye to light is a complex, still not well understood process. It is difficult to quantify due to the high level of interaction between the visual system and complex brain functions. A sketch of the anatomical components of the human eye is shown in Figure 5. The main structures are the iris, lens, pupil, cornea, retina, vitreous humor, optic disk and optic nerve.

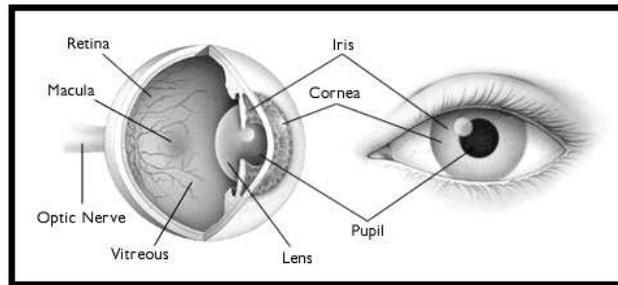


Figure 5: Cross section of the human eye

The path of light through the visual system begins at the pupil, is focused by the lens, then passes onto the retina, which covers the back surface of the eye. The retina is a mesh of photoreceptors, which receive light and pass the stimulus on to the brain. The internal structure of the human eye, a sphere, typically 12mm in radius, is enclosed by a protective membrane, the sclera. At the front of the sclera lies the cornea, a protruding opening, and an optical system comprising the lens and ciliary muscles which change the shape of the lens providing variable focus. Light enters the eye through the lens and proceeds through the vitreous humor, a transparent substance, to the rear wall of the eye, the retina. The retina has photoreceptors coupled to nerve cells, which intercept incoming photons and output neural signals. These signals are transmitted to the brain through the optic nerve, connected to the retina at the optic disk or papilla, more commonly known as the blind spot. The retina is composed of two major classes of receptor cells known as rods and cones. The rods are extremely sensitive to light and provide achromatic vision at low (scotopic) levels of illumination. The cones are less sensitive than the rods but provide color vision at high (photopic) levels of illumination. A schematic drawing of rod and cone cells is shown in Figure 5. Cones are nerve cells that are sensitive to light, detail, and color. Millions of cone cells are packed into the macula, aiding it in providing the visual detail needed to scan the letters on an eye chart, see a street sign, or read the words in a newspaper. Rods are designed for night vision. They also provide peripheral vision, but they do not see as acutely as cones. Rods are insensitive to color. When a person passes from a brightly lit place to one that is dimly illuminated, such as entering a movie theatre during the day, the interior seems very dark. After some minutes this impression passes

and vision becomes more distinct. In this period of adaptation to the dark, the eye becomes almost entirely dependent on the rods for vision, which operate best at very low light levels. Since the rods do not distinguish color, vision in dim light is almost colorless.

Cones provide both luminance and color vision in daylight. They contain three different pigments, which respond either to blue, red, or green wavelengths of light. A person who is missing one or more of the pigments is said to be color-blind and has difficulty distinguishing between certain colors, such as red from green. These photoreceptor cells are connected to each other and the ganglion cells which transmit signals to and from the optic nerve. Connections are achieved via two layers, the first and second synaptic layers. The interconnections between the rods and cones are mainly horizontal links, indicating a preferential processing of signals in the horizontal plane.

Normal daytime vision, where the cones predominate visual processing, is termed photopic, whereas low light levels where the rods are principally responsible for perception is termed scotopic vision. When both rods and cones are equally involved then vision is termed mesopic. Visual acuity is the ability of the Human Visual System (HVS) to resolve detail in an image. The human eye is less sensitive to gradual and sudden changes in brightness in the image plane but has higher sensitivity to intermediate changes. Acuity decreases with increase in distance. Visual acuity can be measured using a Snellen Chart, a standardized chart of symbols and letters. Visual field indicates the ability of each eye to perceive objects to the side of the central area of vision. A normal field of vision is 180 degrees .

### 4.1.3 Contrast

Contrast is defined as:

$$lmax = \frac{lmax - lmin}{lmax + lmin} \quad (1)$$

where  $lmax$  and  $lmin$  are the maximum and minimum luminance. Human brightness sensitivity is logarithmic, so it follows that for the same perception, higher brightness requires higher contrast. Apparent brightness is dependent on background brightness. This phenomenon, termed simultaneous contrast, is illustrated in Figure 6. Despite the fact that all centre squares are the same brightness, they are perceived as different due to the different background brightness.

Depth Perception is the ability to see the world in three dimensions and to perceive distance. Images projected onto the retina are two-dimensional, and from these flat images vivid three dimensional worlds are constructed. Binocular Disparity and monocular cues provide information for depth perception. Binocular disparity is the difference between the images projected onto the left and right eye. The brain integrates these two images into a single three dimensional image to allow depth and distance perception. Monocular cues are cues to depth that are effective when viewed with only one eye, including interposition,

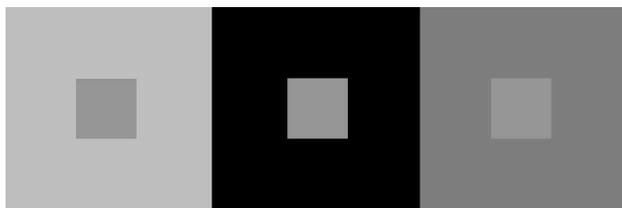


Figure 6: Simultaneous Contrast. Despite the fact that all centre squares are the same brightness, they are perceived as different due to the different background brightness.

atmospheric perspective, texture gradient, linear perspective, size cues, height cues and motion parallax.

#### 4.1.4 Constancy

Perceptual Constancy is a phenomenon which enables the same perception of an object despite changes in the actual pattern of light falling on the retina. Psychologists have identified a number of perceptual constancies including lightness constancy, color constancy, size constancy and shape constancy. Lightness Constancy: The term lightness constancy describes the ability of the visual system to perceive surface lightness correctly despite changes in the level of illumination.

Color Constancy: Closely related to lightness constancy, this is the ability of the HVS to perceive the correct color of an object despite changes in illumination.

Shape Constancy: Objects are perceived as having the same shape regardless of changes in their orientation. -example with cube, from front and side

Size Constancy: This is the tendency to perceive objects as staying the same size despite changes in viewing distance.

#### 4.1.5 Human Visual Perception

A number of psychophysical experimental studies have demonstrated many features of how the HVS works. However, problems arise when trying to generalize these results for use in computer graphics. This is because, often, experiments are conducted under limited laboratory conditions and are typically designed to explore a single dimension of the HVS. As described earlier, the HVS comprises complex mechanisms, which rather than working independently, often features work together, and therefore it makes sense to examine the HVS as a whole. Instead of reusing information from previous psychophysical experiments, new experiments are needed. Some examples will support this. Figure 2.6: When a black and white patterned top shown on the left is rotated at 5-10 revolutions per second, colored rings can be seen. The light intensity distribution of the rotating pattern as a function of time is shown on the right. Spatiotem-

poral interactions between antagonistic, spectrally opponent color mechanisms account for this phenomenon. A Benhams disk is a flat disc, half of which is black and the other half has three sets of lines like the grooves on a record but more spaced out, Figure 7. When the disk is spun a human observer sees red, yellow and green rings, despite the fact that there are no colors in the pattern. The curves on the right of the pattern begin to explain what happens. Each curve plots the temporal light intensity distribution at the different radii from the centre, created when the top is spun. These changing light patterns produce spatiotemporal interaction in the HVS that unbalance antagonistic, spectrally-opponent mechanisms to create the appearance of colored rings. This illusion demonstrates that, although it may be convenient to model the HVS in terms of unidimensional responses to motion, pattern and color, human percepts are in fact the product of complex multidimensional response.

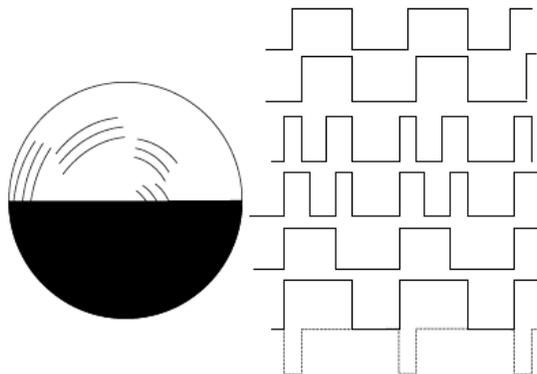


Figure 7: When a black and white patterned top shown on the left is rotated at 5-10 revolutions per second, colored rings can be seen. The light intensity distribution of the rotating pattern as a function of time is shown on the right. Spatiotemporal interactions between antagonistic, spectrally opponent color mechanisms account for this phenomenon.

A second example, Figure 8, shows the panels in checkerboard block on the left and a flat pattern on the right, which have the same reflectance, but differences in their three-dimensional organization means they are perceived differently. The two panels marked with Xs have the same reflectance, but on the block they appear to have different reflectance under different levels of illumination. Conversely, the two panels marked with Os have different reflectance values but on the block appear to be the same color due to the different illumination conditions. This demonstrates the complexity of interactions between apparent reflectance, apparent illumination and apparent shape that can dramatically affect human perception.

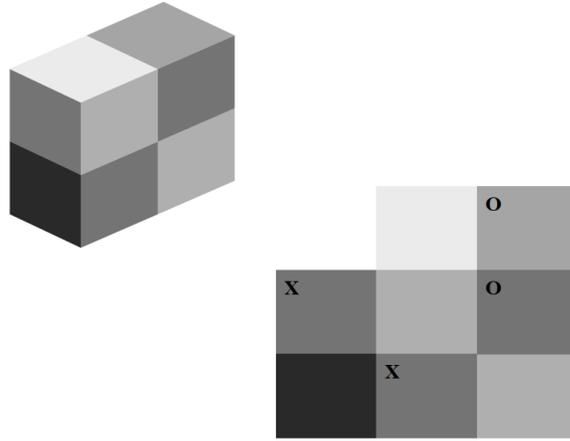


Figure 8: Interaction between apparent reflection, apparent illumination and apparent three-dimensional shape. Corresponding panels in the two patterns have the same physical reflectance. Differences in the perceived spatial organization of the patterns produces differing interpretations in terms of lightness (apparent reflectance) and brightness (apparent illumination).

#### 4.1.6 Lightness Perception

Gilchrist [43, 44, 18] justified the systematic study of lightness error as an understanding of the HVS. He found that there are always errors when judging lightness, and these errors are not random, but systematic. The pattern of these systematic errors therefore provide a signature of the visual system. He defines a lightness error as any difference between the actual reflectance of a target surface and the reflectance of the matching chip selected from a Munsell chart. The task defined for the psychophysical experiments described later in this thesis involves asking human observers to match the reflectance of real world objects to a Munsell chart, which gives a measure of errors in lightness matching. The observer is then asked to match the reflectance of simulated objects (in a computer generated rendition of the real world) to the same Munsell chart. This gives a measure of lightness errors with respect to the computer image. There are limitations on the HVS, so there will be errors (systematic errors) in both cases. For the rendered image to be deemed a faithful representation, both sets of lightness errors should be close to each other.

Gilchrist (1977) [45] showed that the perception of the degree of lightness of a surface patch (i.e. whether it is white, gray or black) is greatly affected by the perceived distance and orientation of the surface in question, as well as the perceived illumination falling on the surface -where the latter was experimentally manipulated through a variety of cues such as occlusion, or perspective. Perception of the lightness of patches varying in reflectance may thus be a suitable candidate for the choice of visual task. It is simple to perform, and it is

known that lightness constancy depends on the successful perception of lighting and the 3D structure of a scene, for example Figure 9. When viewed in isolation, the patches on the top left hand corner appear to be of different luminance. However, when examined in the context of the entire scene, it can be seen that the patches have been cut from the edge of the stairwell, and are perceived as an edge where the entire stairwell has the same luminance. Lightness has been applied when developing Tone Mapping techniques for High Dynamic Range Imagery [61, 62].

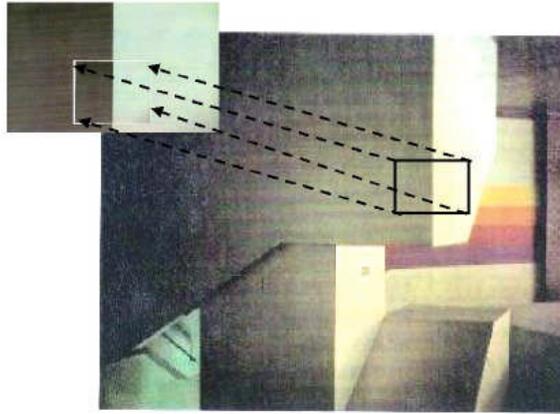


Figure 9: Importance of depth perception for lightness constancy.

## 4.2 Perceptually driven rendering

Recent years have seen an increase in the application of visual perception to computer graphics. As mentioned earlier, in certain applications it is important that computer images should not only be physically correct but also perceptually equivalent to the scene it is intended to represent. Realism implies computational expense, and research is beginning to emerge to investigate how knowledge of the human visual system can be used to cut corners and minimize rendering times by guiding algorithms to compute only what is necessary to satisfy the observer. Perception based image quality metrics, which can be used to evaluate, validate and compare imagery have been presented [99, 97, 80, 76, 86].

Even for realistic image synthesis there may be little point spending time or resources to compute detail in an image that would not be detected by a human observer. By eliminating any computation spent on calculating image features which lie below the threshold of visibility, rendering times can be shortened leading to more efficient processing. Because the chief objective of physically based rendering is realism, incorporating models of HVS behavior into rendering algorithms can improve performance, as well as improving the quality of the

imagery produced. So by taking advantage of the limitations of the human eye, just enough detail to satisfy the observer can be computed without sacrificing image quality. Several attempts have been made to develop image synthesis algorithms that detect threshold visual differences and direct the algorithm to work on those parts of an image that are in most need of refinement.

Raytracing produces an image by computing samples of radiance, one for each pixel in the image plane. Producing an anti-aliased image is difficult unless very high sampling densities are used. Mitchell [85] realized that deciding where to do extra sampling can be guided by knowledge of how the eye perceives noise as a function of contrast and color. Studies have shown that the eye is most sensitive to noise in intermediate frequencies [123]. While frequencies of up to 60 cycles per degree (cpd) can be visible, the maximum response to noise is at approximately 4.5 cpd, so sampling in regions with frequency above this threshold can be minimized, without affecting the visual quality of the image. Mitchell begins by sampling the entire image at low frequency then uses an adaptive sample strategy on the image according to the frequency content. This results in a non uniform sampling of the image, which enables aliasing noise to be channelled into high frequencies where artifacts are less conspicuous. However, non-uniform sampling alone doesn't eliminate aliasing, just changes its characteristics to make it less noticeable. Mitchell applies two levels of sampling. To decide whether the high sampling density should be invoked the variance of samples could be used [89], but this is a poor measure of visual perception of local variation. Instead Mitchell chooses to use contrast to model the non-linear response of the eye to rapid variations in light intensity:

As each sample consists of three separate intensities for red, green and blue, three separate contrasts can be computed for each of them. These three contrasts are tested against separate thresholds, 0.4, 0.3 and 0.6 for red, green and blue respectively, and super-sampling is done if any one exceeds the threshold. The contrast metric is then used to determine when the high sampling density should be invoked. This test is most sensitive to green in accordance with the human eyes response to noise as a function of color. Multi stage filters are then used to reconstruct the non-uniform samples into a digital image. Although this idea has the beginnings of a perceptual approach, it is at most a crude approximation to the HVS. Only two levels of sampling are used and it doesn't account for visual masking <sup>1</sup>.

The HVS exhibits different spatial acuities in response to different colors. Evidence exists that color spatial acuity is less than monochrome spatial acuity. Exploiting this poor color spatial acuity of the HVS, Meyer and Liu [84] developed an adaptive image synthesis algorithm which uses an opponent's processing model of color vision [61] comprising chromatic and achromatic color channels. Using a Painter and Sloan [121] adaptive subdivision, a k-D <sup>2</sup> tree representation of the image is generated. Areas of the image containing high frequency

---

<sup>1</sup>The presence of high spatial frequency in an image can mask the presence of other high frequency information

<sup>2</sup>A KD Tree is a data structure that is used in computer science during orthogonal range searching

information are stored at the lower levels of the tree. They then modified a screen subdivision raytracer to limit the depth to which the k-D tree must be descended to compute the chromatic color channels. The limit is determined by psychophysical results describing the color spatial frequency. They achieved a modest saving in computational effort and showed, using a psychophysical experiment, that decreasing the number of rays used to produce the chromatic channels had less of an effect on image quality than reducing the number of rays used to create the achromatic channels. This was the first work to attempt to minimize the computation of color calculations, as opposed to just decreasing costly object intersection calculations.

Bolin and Meyer [9] took a frequency based approach to raytracing, which uses a simple vision model, making it possible for them to control how rays are cast in a scene. Their algorithm accounts for the contrast sensitivity, spatial frequency and masking properties of the HVS. The contrast sensitivity response of the eye is non-linear. So, when deciding where rays should be cast, the algorithm deems a luminance difference at low intensity to be of greater importance than the same luminance difference at high intensity. The spatial response of the HVS is known to be less for patterns of pure color than for patterns that include luminance differences. This means that it is possible to cast fewer rays into regions with color spatial variations than are cast in regions with spatial frequency variations in luminance. Finally, it is known that the presence of high spatial frequency can mask the presence of other high frequency information (masking). When used in conjunction with a Monte Carlo raytracer, more rays are spawned when low frequency terms are being determined than when high frequency terms are being found. Using this strategy, the artifacts that are most visible in the scene can be eliminated from the image first, then noise can be channelled into areas of the image where artifacts are less conspicuous. This technique is an improvement on Mitchells method because the vision model employed accounts for contrast sensitivity, spatial frequency and masking.

Despite the simplicity of the vision models used in these approaches, the results are promising, especially as they demonstrate the feasibility of embedding HVS models into the rendering systems to produce more economical systems without forfeiting image quality. Fueled by the notion that more sophisticated models of the HVS would yield even greater speedup, several researchers began to introduce more complex models of the HVS into their global illumination computations.

Myszkowski [88] applied a more sophisticated vision model to steer computation of a Monte Carlo based raytracer. Aiming to take maximum advantage of the limitations of the HVS, his model included threshold sensitivity, spatial frequency sensitivity and contrast masking. A perceptual error metric is built into the rendering engine allowing adaptive allocation of computation effort into areas where errors remain above perceivable thresholds and allowing computation to be halted in all other areas (i.e. those areas where errors are below the perceivable threshold and thus not visible to a human observer). This perceptual error metric takes the form of Dalys [25] Visible Difference Predictor (VDP).

Bolin and Meyer [10] devised a similar scheme, also using a sophisticated vision model, in an attempt to make use of all HVS limitations. They integrated a simplified version of the Sarnoff Visible Discrimination Model (VDM) into an image synthesis algorithm to detect threshold visible differences and, based on those differences direct subsequent computational effort to regions of the image in most need of refinement. The VDM takes two images, specified in CIE XYZ color space, as input. Output of the model is a Just Noticeable Difference (JND) map. One JND corresponds to a 75% probability that an observer viewing the two images would detect a difference [81]. They use the upper and lower bound images from the computation results at intermediate stages and used the predictor to get an error estimate for that stage. Drettakis et al introduced a perceptual rendering pipeline which takes into account visual masking due to contrast and spatial frequency[29]. Scenes are split into layers to account for inter-object masking. Using a perceptually driven level of detail algorithm the layers are then used to choose an appropriate level of detail for each object based on predicted contrast and and spatial masking. A subsequent user study showed that their algorithmic choices corresponded well with perceived differences in the images. Masking has also been used in geometric modeling, Lavoué et. al. [65] introduced the notion of *roughness* for a 3D mesh. Roughness gives a measure of geometric noise on the surface, based on this noise masking can be invoked to hide geometric distortions.

Applying a complex vision model at each consecutive time step of image generation requires repeated evaluation of the embedded vision model. The VDP can be expensive to process due to the multi-scale spatial processing involved in some of its components. This means that in some cases the cost of recomputing the vision model may cancel the savings gained by employing the perceptual error metric to speed up the rendering algorithm. To combat this, Ramasubramanian [100] introduced a metric that handles luminance-dependent processing and spatially-dependent processing independently, allowing the expensive spatially-dependent component to be precomputed. Ramasubramanian developed a physical error metric that predicts the perceptual threshold for detecting artifacts in the image. This metric is then used to predict the sensitivity of the HVS to noise in the indirect lighting component. This enables a reduction in the number of samples needed in areas of an image with high frequency texture patterns, geometric details, and direct lighting variations, giving a significant speedup in computation.

Using validated image models that predict image fidelity, programmers can work towards achieving greater efficiencies in the knowledge that resulting images will still be faithful visual representations. Also in situations where time or resources are limited and fidelity must be traded off against performance, perceptually based error metrics could be used to provide insights into where computation could be economized with least visual impact.

In addition to Tone Mapping Operators (TMOs) being useful for rendering calculated luminance to the screen [119, 63, 26, 103], they are also useful for giving a measure of the perceptible difference between two luminances at a given level of adaptation. This function can then be used to guide algorithms, such as

discontinuity meshing, where there is a need to determine whether some process would be noticeable or not to the end user.

Gibson and Hubbold [42] have used features of the threshold sensitivity displayed by the HVS to accelerate the computation of radiosity solutions. A perceptually based measure controls the generation of view independent radiosity solutions. This is achieved with an a priori estimate of real-world adaptation luminance, and uses a TMO to transform luminance values to display colors and is then used as a numerical measure of their perceived difference. The model stops patch refinement once the difference between successive levels of elements becomes perceptually unnoticeable. The perceived importance of any potential shadow falling across a surface can be determined, this can be used to control the number of rays cast during visibility computations. Finally, they use perceptual knowledge to optimize the element mesh for faster interactive display and save memory during computations. This technique was used on the adaptive element refinement, shadow detection, and mesh optimization portions of the radiosity algorithm.

Discontinuity meshing is an established technique used to model shadows in radiosity meshes. It is computationally expensive, but produces meshes which are far more accurate and which also contain fewer elements. Hedley et al. [50] used a perceptually informed error metric to optimize adaptive mesh subdivision for radiosity solutions, the goal being to develop scalable discontinuity meshing methods by considering visual perception. Meshes were minimized by discarding discontinuities which had a negligible perceptible effect on a mesh. They demonstrated that a perception-based approach results in a greater reduction in mesh complexity, without introducing more visual artifacts than a purely radiometrically-based approach.

Farrugia and Peroche [36] used a perceptual metric for discontinuity refinement to develop a progressive radiance evaluation based on the work of Guo et al. [46]. Guo used an iterative process to construct an irregular subdivision of the image in blocks which refer to smooth regions, or discontinuous regions to build a Directional Coherence Map (DCM). As the algorithm proceeds the current DCM dictates where new samples are taken. A contrast based perceptual heuristic based on contrast over samples corresponding to the corners of each block is used. Farrugia and Peroch extend this by applying a visual differences predictor based on Pattanaik et al.s Multiscale Model of Adaptation and Spatial Vision to classify their subdivision cells [94]. To speed up computation they apply the metric over each cell pair using a statistical approach based on Albin et. al. [2].

Recognizing that the illumination on a surface can be split into separable components, which can be individually computed, Stokes et. al [112] introduced a new approach which applied a perceptual metric on each component. After a suite of psychophysical experiments to probe various global illumination scenarios they determined limited contribution of light path interaction and fitted a mathematical model used to guide rendering based on the metric predicted relative importance of each component as a function of visible surface materials.

Ramanarayanan et. al. noticed that when viewing an aggregate, observers

attend less to individual objects and focus more on overall properties such as numerosity, variety, and arrangement. They also noted that rendering and modeling costs increase with aggregate complexity, exactly when observers are attending less to individual objects. They presented new aggregate perception metrics to simplify scenes by substituting geometrically simpler aggregates for more complex ones without changing appearance [98].

Ramanarayanan [99] worked toward developing a perceptual metrics based on higher order aspects of visual coding and introduced the term "Visual equivalence". Images are visually equivalent if they convey the same impressions of scene appearance, even if they are visibly different. They conducted a series of psychophysical experiments to investigate how object geometry, material, and illumination interact influence appearance. In their paper they characterized conditions under which two classes of transformations on illumination maps (blurring and warping) yield images that are visually equivalent to reference solutions, and from this developed a metric to predict visual equivalence.

### 4.3 Conclusion

Using validated image models that predict image fidelity, programmers can work toward achieving greater efficiencies in the knowledge that resulting images will still be faithful visual representations. Also in situations where time or resources are limited and fidelity must be traded off against performance, perceptually based error metrics could be used to provide insights into where computation could be economized with least visual impact.

Some of the applications of visual perception in computer graphics were explored. For many applications computer imagery should not only be physically correct but also perceptually equivalent to the scene it represents. Knowledge of the HVS can be employed to greatly benefit the synthesis of realistic images at various stages of production. Global illumination computations are costly in terms of computation. There is a great deal of potential to improve the efficiency of such algorithms by focusing computation on the features of a scene which are more conspicuous to the human observer. Those features that are below perceptual visibility thresholds have no impact on the final solution, and therefore can be omitted from the computation, increasing efficiency without causing any perceivable difference to the final image. Perceptual metrics involving advanced HVS models can be used to determine the visible differences between a pair of images. These metrics can then be used to compare and evaluate image quality. They can also be used within the rendering framework to steer computation into regions of an image which are in most need of refinement, and to halt computation when differences in successive iterations of the solution become imperceptible.

**Katerina Mania**

## **5 Perceptually Motivated Simulation and Virtual Environments**

### **5.1 Introduction**

Computer graphics algorithms have for long dealt with simulation of physics: simulation of the geometry of a real-world space, simulation of the light propagation in a real environment and simulation of motor actions with appropriate tracking. Perception principles have subsequently been incorporated into rendering algorithms [82], in order to save rendering computation, mainly following the generic idea of “do not render what we cannot see” [77, 54, 68]. However, with Virtual Environment (VE) simulator technologies aiming at simulating real-world task situations, the research community is challenged to produce a much more complex system which is perceptually optimized. We do not necessarily require accurate simulation of physics to induce reality. Much less detail is often adequate [126, 37, 72, 73].

### **5.2 Perceptually-based Selective Rendering**

Perception principles have been incorporated into rendering algorithms in order to optimize rendering computation and produce photorealistic images from a human rather than a machine point of view. In order to economize on rendering computation, selective rendering guides high level of detail to specific regions of a synthetic scene and lower quality to the remaining scene, without compromising the level of information transmitted. Scene regions that have been rendered in low and high quality can be combined to form one complete scene. Such decisions are guided by predictive attention modeling, gaze or task-based information.

In order to economize on rendering computation, previous research dealing with interactive synthetic scenes has been focused on either rendering in high quality the 2-3 degrees foveal region of vision and with less detail the periphery of vision based on gaze information [70], or rendering in high quality the foveal area based on a-priory knowledge of the viewers task focus [19, 113]. Gaze-dependent rendering encounters difficulties of maintaining display updates free of visual artifacts after a fast (4ms) eye saccade. Such processes are quite computationally demanding, however, if the speed gaze-to-rendering issue is resolved, task performance results are indistinguishable to a fully fledged, high resolution real-time environment. It has also been proposed to assign selective high quality rendering in the visual angle of the fovea (2o) centered on the users task focus [19, 113]. This approach, however, cannot be applied when there is no overt task to be conducted. Moreover, there is no acceptable model of comparing or predicting task-relevant saccades. Following a different approach, Haber et al. [47] suggested rendering the informative areas of a scene in varying quality

based on saliency models. Such models aim to predict the visual features that involuntarily attract visual attention such as object edges, sudden color changes or movements. It was proposed that the most noticeable areas as derived from saliency modeling should be rendered in higher quality. Bottom-up visual attention models are not shown to predict attention regions successfully [77]. Correlation between actual human and computationally-derived scan-paths was found to be much lower than predicted when carrying out a real-world task such as making a cup of tea [35]. Moreover, we have no generally accepted model of comparing scan paths.

A comprehensive approach should be task and gaze-independent, simulating cognitive processes rather than predicting attention employing bottom-up processes such as saliency models. A recent selective rendering approach exploits existing research on memory schemata which could ultimately guide selective rendering based on spatial cognition processes. Schemata are knowledge structures based on the notion that an individual's prior experience will influence how he or she perceives, comprehends and remembers new information. When participants are exposed to a large amount of information in a scene, cognitive psychologists have suggested that schemata are used to guide the search for information in memory [13]. A general premise derived from this research is that information which is not related to the schema being used in retrieval will be harder to recall than information which is schema related. In terms of real world scenes, schemata represent the general meaning of a scene such as office, theatre etc. Schemata influence memory of the objects in a given context according to their association with the schema in place. When being exposed to a synthetic environment, similar information should be transmitted between the simulated scene and the real-world scene, both depicting a specific schema. This would, in due course, indicate which objects or areas in a synthetic scene could be rendered in lower quality without affecting information uptake but at the same time reducing computational complexity [87].

Flannery and Walles [39] investigated how schema theories apply to real versus virtual memories. Participants were instructed to explore either a virtual or a similar real environment for 20 seconds, without prior knowledge that their memory of the space would be subsequently assessed. Participants then completed a recognition task. Recognition scores revealed that participants had better recognition for consistent objects, but were more confident for the recognition of the inconsistent objects.

Previous work [71, 72], included a preliminary investigation of the effect of object type (consistent vs. inconsistent) and shadows (flat-shaded scene vs. radiosity scene) on object memory recognition in a VE. The computer graphics simulation was displayed on a Head Mounted Display (HMD) utilizing stereo imagery and head tracking. Thirty-six participants across three conditions of varied rendering quality of the same space were exposed to the computer graphics environment and completed a memory recognition task. The high-quality and mid-quality conditions included a pre-computed radiosity simulation of an academics office (with 80% and 40% radiosity iterations computed respectively). The low-quality condition consisted of a flat-shaded version of the same office.

Results revealed that schema consistent elements of the scene were more likely to be recognized than inconsistent information. Overall, higher confidence ratings were assigned to consistent rather than inconsistent items. Total object recognition was better for the scene including rough shadows (mid-quality condition) compared to the flat-shaded scene. The presence of accurate shadow information, though, did not affect recognition of consistent or inconsistent objects, therefore lower quality of rendering was adequate for better memory recognition of consistent objects. This study was limited to the investigation of subtle shadow variations. Another experimental study employed a more extreme set of rendering types: wireframe with added color, and full radiosity [87] (Figure 3). The proportion of inconsistent/consistent objects was varied, and object recognition tests ensured that all objects were easily recognized in all conditions. The results showed a significant interaction between rendering type, object type, and consistency ratio. This suggests that inconsistent objects are only preferentially remembered if the scene looks normal or if there are many such objects in an abnormal scene such as in the wireframe condition. It was also shown that memory performance is better for the inconsistent objects in the radiosity rendering condition compared to the wireframe condition. We conclude that memory for objects can be used to assess the degree to which the context of a VE appears close to expectations.

Despite contradictory results in literature as detailed above, it seems that perceptual information can be complemented by involuntary knowledge based on past experience. Experimental studies in synthetic scenes have revealed that consistent objects which are expected to be found in a scene can be rendered in lower quality without affecting information uptake taking advantage of such expectations, whereas inconsistent items which are salient would require a high level of rendering detail in order for them to be perceptually acknowledged [71]. Therefore, by exploiting schema theory, it is possible to reduce computational complexity, producing scenes from a cognitive point of view without affecting information uptake and resulting in an entirely novel and interdisciplinary approach which is gaze, task and saliency-model independent. A novel selective rendering system has been presented that exploits schema theory by identifying the perceptual importance of scene regions [134]. Objects that have been rendered in low and high quality are incorporated in a scene based on schema expectations. The rendered quality of these objects will change in real time, dependent on user navigation and interaction [87].

High level visual cognition is that which takes place late in the HVS, that is parietal and temporal cortex and into the frontal lobes when decisions based on visual information need to be made. Low level visual cognition is that which takes place in the occipital lobe, early on in the visual processing stream, e.g. the visual signal is received in the retinae, and initially passed through the Lateral Geniculate Nucleus to the occipital lobe at the back. Thus, these higher decision processes are unaffected by changes in the experiments, whilst normal (visual) cognition occurs as long as the scene is realistic. Taken together, the results of previous studies investigating the effect of schemas on object recognition suggest that high-level visual cognition is generally unaffected by ubiquitous

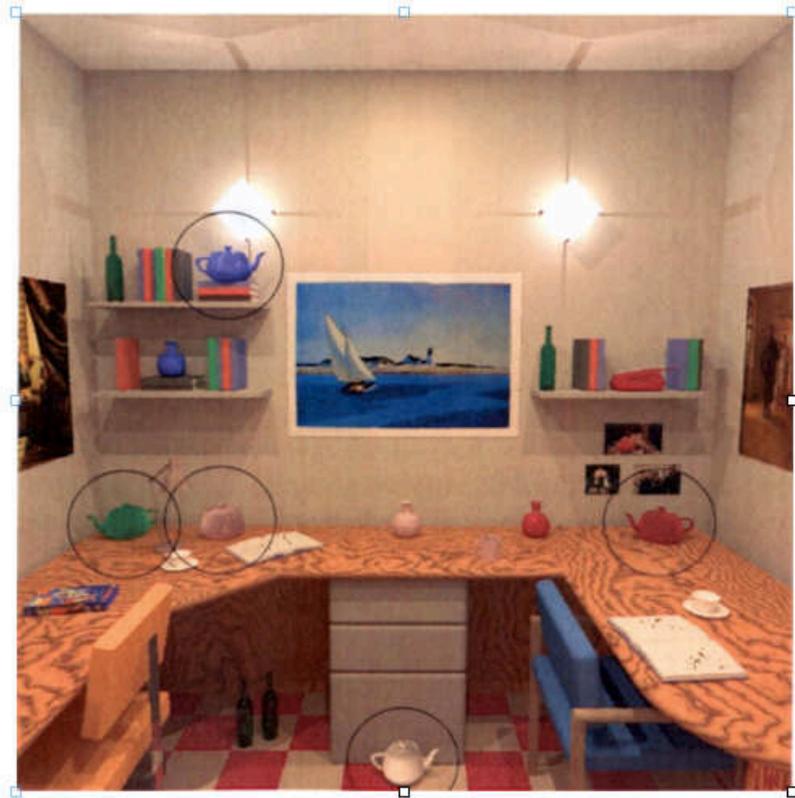


Figure 10: Task-based rendering [19].



Figure 11: Gaze-based rendering [70].

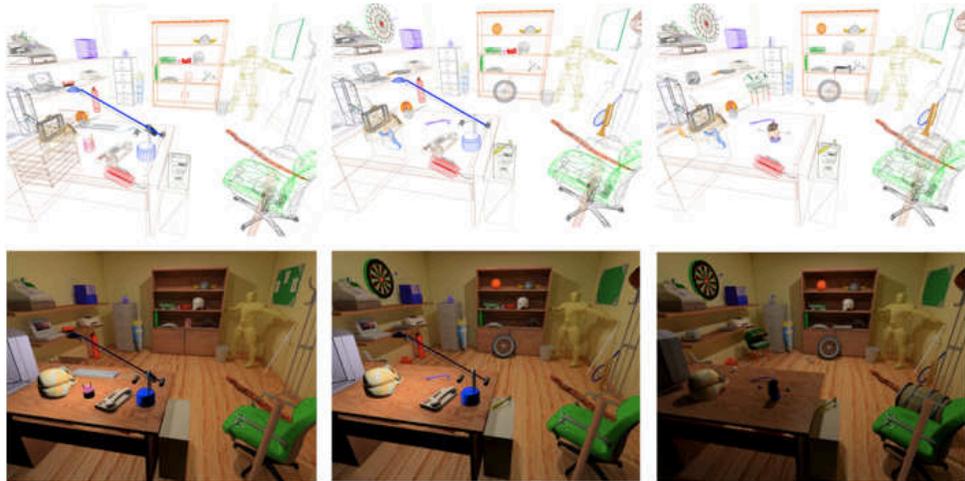


Figure 12: Schema memory experimental studies [87].

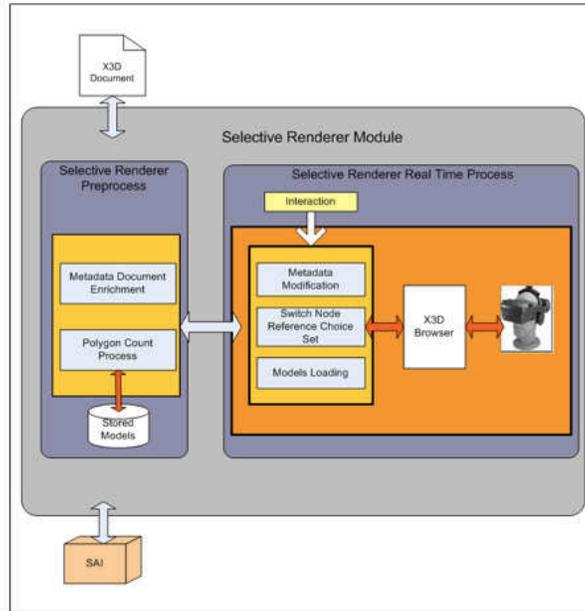


Figure 13: Selective Renderer Architecture [134].

graphics manipulations such as polygon count and depth of shadow rendering; normal cognition operates as long as the scenes look acceptably realistic. On the other hand, when the overall realism of the scene is greatly reduced, such as in the wireframe condition, then visual cognition becomes abnormal. Specifically, effects that distinguish schema-consistent from schema-inconsistent objects change because the whole scene now looks incongruent, and we have shown that this effect is not due to a failure of basic recognition. Thus, a recipe for anyone wishing to use such displays in studies of visual cognition is to construct environments which look acceptably realistic in terms of polygon count but need not be of very high quality. This relates to the kinds of high-level visual cognitive effects we have studied here, such as object congruence. Lower-level effects, such as recognition, can be dissociated from these high-level effects (but make their presence felt when the scene is further degraded, e.g. when color is removed from the wireframe scenes). Thus, high-level processes need somewhat greater realism than low-level ones.

### 5.3 Behavioral Fidelity of Simulations based on Space Memory

The entertainment world appears to consider highly realistic visual quality one of the keys to success, with cinematic quality graphics claimed for the next generation of gaming consoles. On the other hand, when interactive immersive

Virtual Environments (VEs) are implemented for training rather than entertainment purposes, visual quality might not be as significant. If the training is to be effective, the skills acquired must transfer into the real world at appropriate levels of performance. A VE with maximum visual and interaction fidelity would result in a transfer of information equivalent to real world training since the environments would be indistinguishable [73]. Visual fidelity refers to the degree to which visual features in the VE conform to visual features in the equivalent real environment [122, 106]. Functional realism refers to the communication of similar information in the real and virtual world rather than the aesthetics or physics, in the sense that users are able make the same judgments and perform the same tasks as in the real world [37, 91].

It is tempting to replicate the real world as accurately as possible in order to provide equivalent experiences [67]. Whilst arguably ideal, it is not yet computationally feasible for this to occur. Trade-offs between visual/interaction fidelity and computational complexity should be applied to a simulation system without detracting from its training effectiveness [72, 124]. There is, therefore, a call for efficient techniques assessing the fidelity of a VE and determine its relationship with performance in order to economize on rendering computation without compromising the level of information transmitted (functional realism) [37].

The utility of Virtual Environment (VE) technologies for training systems such as flight simulators is predicated upon the accuracy of the spatial representation formed in the VE. Spatial memory tasks, therefore, are often incorporated in benchmarking processes when assessing the fidelity of a VE simulation for training. Spatial awareness is significant for human performance efficiency of such tasks as it is dependent on spatial knowledge of an environment [64, 28, 129]. A central research issue, therefore, is how an interactive synthetic scene is cognitively encoded and how recognition and memory of such worlds transfer to real world conditions [72, 1, 38]. Previous research has examined the variables that communicate transfer of spatial knowledge acquired in a simulation environment, in the real-world and discuss the form and development of spatial awareness in VE training compared to either real-world training or training with maps, photographs and blueprints [8, 6]. The suitability of VE systems as effective training mediums was examined and was concluded to be as effective as map or blueprint training. Configurational knowledge acquisition based on estimation of absolute distances and directions between known points could yield training effects similar to training with photographs and real world training [8]. Furthermore, estimation of travel distance from optic flow is subject to scaling when compared to static intervals in the environment, irrespective of additional depth cues [40]. Past research often aims to identify the minimum system characteristics relevant to rendering computations and interaction interfaces that would yield the maximum performance on a task or the greatest sense of presence. For example, search objects rendered in global or ambient illumination have been shown to take significantly longer to identify than those rendered through a local illumination model [133]. What if the visual fidelity of a system should be assessed across a range of applications and

tasks? Could we interrogate the human cognitive systems that are activated when training within VE scenes of varied visual or interaction fidelity in order to identify whether such responses are transferable to the real-world task situation simulated? Which simulation characteristics should we optimize in order to match the capabilities of the VE system to the requirements of these cognitive systems?

Because of the wide-range of VE applications and differences in participants across their backgrounds, abilities and method of processing information, an understanding of how spatial knowledge is acquired within a VE, complementing spatial memory performance per se, is significant. Common strategies may be revealed across a range of applications and tasks. Recent research focuses upon the effect of rendering quality (flat-shaded vs radiosity) on object-location recognition memory and its associated awareness states while spatial knowledge is transferred from a synthetic training environment into a real-world situation. The main premise of this work is that accuracy of performance per se is an imperfect reflection of the cognitive activity that underlies performance on memory tasks [74]. The framework to be presented has been drawn from traditional memory research adjusted to form an experimental procedure [118, 11, 27].

Accurate recognition memory can be supported by: a specific recollection of a mental image or prior experience (remembering); reliance on a general sense of knowing with little or no recollection of the source of this sense (knowing); guesses. Gardiner and Richardson-Klavehn [41] explained the remembering as personal experiences of the past that are recreated mentally. Meanwhile knowing refers to other experiences of the past but without the sense of reliving it mentally. The work of Tulving [118] first suggested that remembering and knowing were measurable constructs. Through a series of experiments, Tulving [118] reported that participants find it easy to distinguish between experiences of remembering and knowing when self-reporting their experiences. The sense of knowing has since been further divided into two related concepts. The correct answer may be just known without the associated recollection of contextual detail associated with remembering or the answer may feel more familiar than a uninformed guess, but cannot be considered as being known (familiar).

According to this theoretical framework derived from memory psychology, measures of the accuracy of memory can be complemented by self-report of states of awareness such as remember, know, familiar and guess during recognition [24]. Previous studies have investigated the relationship between recognition memory and simulation environments of varied visual and interaction fidelity. Such work by the authors of this paper revealed varied distribution of awareness states whilst overall accuracy remained the same across experimental conditions suggesting that measurement of awareness states acts as a useful additional measure to supplement the information provided by accuracy [75, 72].

A different study employing the same methodology aims to interrogate the mental processes associated with obtaining spatial knowledge during exposure to a simulated scene while transferring such knowledge in the real-world scene simulated [74]. An object-memory task was performed in the simulated real-world

environment immediately after VE training and a retention test was conducted one week after the VE exposure. The virtual scene was rendered with one of two levels of visual fidelity (flat shaded vs. radiosity rendering) and displayed on a stereo Head Mounted Display (HMD). The experimental scene consisted of a room depicting an academics office. Central to this work is identifying whether high fidelity or low fidelity scenes are associated with stronger visually induced recollections represented by self-report of the remember awareness state. A secondary, exploratory goal is to investigate the effect of schemas on memory recognition post VE exposure. Memory recognition studies in synthetic scenes have demonstrated that low interaction fidelity interfaces such as the mouse compared to head tracking as well as low visual fidelity scenes provoked a higher proportion of visually-induced recollections associated with the remember awareness state, while there was no effect of condition upon memory recognition performance [75, 72].

Broadly, desirable influences on recognition memory and the associated cognitive states may be ultimately identified and generalized to aid specific applications. It could be true, for instance, that for flight simulation applications it is crucial for trainees to refer to mental images associated with instruments as opposed to recollections that are confident but not accompanied by mental imagery when training is transferred into a real-world flight situation. The study presented, therefore, explores the effect of training in immersive environments of varied visual fidelity on the distribution of memory awareness states measured in a real-world task [74]. The fact that it has been shown that interfaces of low interaction or visual fidelity induce a higher number of recollections based on mental imagery when compared with systems of high visual or interaction fidelity, may relate to attentional resources directed to systems that vary strongly from the real-world. The results demonstrated that participants who trained in the low fidelity simulation reported a larger proportion of correct remember responses while conducting the memory recognition task in the real-world situation compared to participants trained in the high fidelity simulation. These results were consistent with previous findings that associated a larger proportion of correct remember responses with low visual and interaction fidelity simulations [75, 72]. The results observed consistently in previous studies was also observed in this study despite the fact that participants physically performed the task in the real-world room after training in its simulated counterpart consisting of an ecologically plausible training scenario.

Recent developments in psychological research have shown that distinctive information or experiences generate more awareness states associated with remembering. For example, participants who are shown typical and distinctive faces are more likely to recognize the distinctive faces in a later memory test with an accompanying experience of remembering [12]. Similar results have also been found using other stimuli such as forenames [11]. In the current context, a low fidelity rendered simulation could be considered as being more distinctive than a high fidelity rendered simulation because of its variation from real. Given that these are immersive environments, distinctiveness in this instance would be judged relative to reality. The less real the environment is, the more

distinctive it can be considered. It would be expected that a more distinctive immersive environment, e.g. a low fidelity one would result in more remember responses than a less distinctive immersive environment, e.g. a high fidelity one. It is worth noting that distinctiveness in this sense may not only refer to visual distinctiveness but to motor responses to the environments [72]. The important variable therefore appears to be differentiation relative to multiple aspects of reality, e.g. visual appearance of, and, motor responses within. Here, higher confidence scores associated with the flat-shaded condition compared to confidence of recollections after training in the radiosity condition further support this suggestion.

Whilst the relationship between distinctiveness and memory may prove useful in explaining these effects it is important to consider what cognitive processes may underlie such a relationship. Previous psychological research has indicated that remember responses require more attentional processing in the first instance than those based on familiarity [93, 12]. A tentative claim would therefore be: immersive environments that are distinctive recruit more attentional resources. This additional attentional processing may bring about a change in participants subjective experiences of remembering when they later recall the environment. This change would therefore lead to an increase in the experience of remembering. Interestingly, this effect was not observed during the retest that revealed similar proportions of awareness states distributed across the viewing conditions. It is likely that the fidelity of the training environment only affects awareness states when transfer of training is tested immediately. As time goes by, the enhanced attentional resources associated with low fidelity environments do not influence the long-term memories associated with the training simulation.

Moreover, it is found here that more correct know responses are reported after training in the high fidelity rendered simulation than in the low fidelity rendered simulation. This would suggest a shift from remember responses to know responses. Memories that are accompanied with a feeling of remembering for participants in the low fidelity simulation are only accompanied with a feeling of knowing in the high fidelity simulation. In line with suggestions made above, this could be explained on the basis of reduced attentional processing of these items in the high fidelity simulation.

## 5.4 Investigating Perceptual Sensitivity to Head Tracking Latency

Virtual Environment (VE) latency is the time lag between a users action in a Virtual Environment and the systems response to this action. This lag typically takes the form of a transport delay and arises from the sum of times associated with measurement processes of the various input devices, computation of the VE contents and interaction dynamics, graphics rendering, and finite data transmission intervals between these various components. The VE and human factors literature has established that these delays have a significant impact on user performance [33], [32] and user impressions of simulation fidelity [31, 33, 60, 78, 1]. Latency negatively affects user performance in 3D object

placement tasks [66, 127].

Excessive latency has long been known to hinder operator adaptation to other display distortions such as static displacement offset [51]. Latency also degrades manual performance, forcing users to slow down to preserve manipulative stability, ultimately driving them to adopt a move and wait strategy [107], [108]. Operator compensation for a delay usually requires the ability to predict the future state of a tracked element.

Interest has more recently been directed toward the subjective impact of system latency relevant to virtual reality simulations. Latency as well as update rate have been considered as factors affecting the operators sense of presence in the environment [128, 120]. In a recent study, lower latencies were associated with a higher self-reported sense of presence and a statistically higher change in heart rate for users while in a stress-inducing (fear of heights), photorealistic environment involving walking around a narrow pit [83].

Since the combination of sensing, computation, rendering, and transmission delay is unavoidable in most VE, tele-operation, and augmented reality applications, interest naturally is directed to how detectable differing levels of latency might be. Both the quantification of perceptual sensitivity to latency and description of the mechanism by which VE latency is perceived will be essential to guide system design in the development of countermeasures such as predictive compensation [5, 60].

Previous research has also focused on the precision, stability, efficiency and complexity of operation interaction and performance with latency-plagued systems [78]. Additionally, the first measures of human operators discrimination of the consequences of latency during head or hand tracked movements have been provided [34, 31]. Related investigations have explored the hypothesis that observers do not explicitly detect time delay, but rather detect the consequences of latency, i.e., they use the artifact motion of the VE scene (away from its normally expected spatially stable location) caused by system time lags [1]. Relevant perceptual thresholds (i.e., Just Noticeable Difference or JND) were identified to average 8-17ms, depending on viewing condition. This psychometric quantity appeared to be invariant across different pedestals (33, 100 and 200ms, standard stimuli). The apparent invariance of the detection function in [34, 31, 1] demonstrated that the classic Webers Law of psychophysics (that JND is linearly proportional to the magnitude of the standard stimulus) did not hold for latency. In other words, observers of long latency VEs will be as sensitive to changes in latency as those who use prompter, more advanced systems. It can also be inferred that the same sensitivity would also apply for comparisons against zero latency pedestal.

Regan et al. [102] found 70.7% latency thresholds averaging 15ms for a specialized non-immersing CRT display. By making assumptions of Gaussian psychometric functions and zero response bias for two-interval forced-choice judgments with balanced presentation order, the 70.7% threshold from [102] can be equated with a JND of 18.6ms.

Allison et al. 2001 observed on the other hand that with large virtual objects occupying the full Head Mounted Display (HMD) Field-of-View (FOV), 50%

thresholds for perceived image instability (oscillopsia) were found to be 180-320ms depending on head motion velocity [3]. This threshold indicates the latency level at which observers were equally likely as not to say the image was unstable and represents their average response bias or preference. Such response biases may be attributable to, among other things, the amount of observer training before the data was collected and the type of judgment task required. In the case of Allison et. al. [3] participants performed single interval judgments i.e., they did not compare each presentation against a standard stimulus but relied on their own internal notion of when an image was no longer stable. Data from citeellis99a, [31], [1] shows their participants response bias ranged between 40 and 70ms for a two-interval judgment of whether the stimulus was the same as or different than the pedestal standard. In contrast, the participants in [102] were forced to choose which of the two stimulus intervals was actually the one with added latency, which though not reported, leads to a presumption of zero bias.

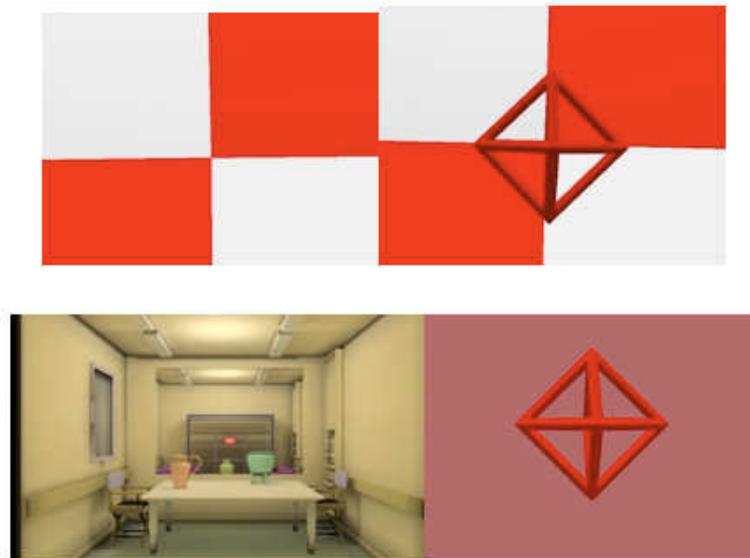


Figure 14: Experimental conditions [30], [73].

The much higher threshold reported by Allison et al. [3] might also be attributable to the fact that their participants viewed a textured virtual background (the inside surface of red and white faceted sphere) that completely enveloped their head and thus always occupied their entire FOV. Surrounding observers with such a geometrically structured environment contributes to the phenomenon of visual capture. The term visual capture implies that when concurrent multi-sensory spatial information is available, the observer will weight the visual channel more heavily in constructing a percept. It has been demon-

strated that, even with very simple VE graphics in an HMD, visually discrepant information will bias proprioceptive and vestibular feedback of static head pitch angle [91]. Since awareness of VE image instability relies on visual, vestibular, and proprioceptive information, the full structured background viewed in Allison et al.'s study [3] may have diminished their observers sensitivity to latency induced oscillopsia. Furthermore, without the inclusion of nearer objects in their environment, participant head movement does not trigger motion of scene contents relative to the background and thus does not provide cues through internal image shear.

One aim of ongoing research on latency perception at NASA Ames Research Centre has been to quantify the latency that a VE system can exhibit without being perceptible to the user. In our prior studies, we employed very sparse environments containing only a single simple object such as a faceted sphere [31], [34] or a hollow-framed octahedron [1] against an empty black background. Synthetic environments with differing levels of graphical complexity with the goal of extending the generality of our results for participant sensitivity latency in VEs have been also employed by Mania et al. [73].

The focus of the [73] study is in describing observer sensitivity to latency differences during head movements in an immersive VE representing a real-world space (room, building, etc.) sensitivity that has not been measured in previous research. On the one hand, because there could be an inherent association with how the real world is perceived, we might expect observers to be more sensitive to the visual consequences of latency when viewing a scene representing what could be a real-world space rather than a sparse, simplified scene with only one or two artificial objects. On the other hand, an enveloping structured scene could promote visual capture, thereby degrading observers sensitivity to VE latency.

During an earlier study more fully reported in [30], a simple white-red checker sphere surrounding the observer, such as that used in [3] and/or a hollow-frame octahedron in front of the observer, as in [1] served as the VEs visual content. Participants were asked to compare two sequential stimulus presentations while moving their head in a constant pattern and report whether the stimuli differed in the visible consequences of the experimentally manipulated VE latency. The study presented here employed the same experimental methodology, but instead, the visual scene was a pre-computed radiosity rendering of two interconnected rooms that include real-world objects. Here, we also statistically compare sensitivity results derived from Ellis et al. [30] and the study presented in this paper. Both studies also explore whether relative motion shear between more than one artificial object in the VE could be a mechanism contributing to observer perception of head tracking latency.

In summary, results from these studies conducted at NASA Ames Research Centre suggest that virtual environment system designers should expect observers who are not burdened with any other performance tasks to generally be able to notice differences in latency as low as 15ms, regardless of the relative location of objects in the scene, the meaningfulness of the scene context in relation to the real world, or possibly even the degree of photorealism in their

rendering. These results will also serve as performance guidelines to aid in the design of predictive compensation algorithms.

## APGV

### 6 Trends from the 7th annual ACM/Eurographics Symposium on Applied Perception in Graphics and Visualization 2010

#### 6.1 APGV Proceedings and Resources

In 2001 a small group of researchers gathered in Snowbird for a “campfire” on Graphics and Perception, there was a certain air of excitement that things were getting started. The growing interest in this area, and the realization that perception is playing an increasingly important role in graphics and visualization lead to the establishment of a symposium dedicated to perceptual research in graphics and visualization, called APGV - Applied Perception in Graphics and Visualization. The goal for this symposium is to have it serve as an inclusive forum where researchers working at the intersection of perception, graphics and visualization can come together to share ideas and results. APGV hopes to provide a great opportunity for people not only to acquire new knowledge, but also to seek new partnerships and collaborations. Now in it’s 7th year papers of the years have represented active interdisciplinary efforts. A wide range of topics have treated . including color, shape, motion, distance judgments, virtual reality, and haptics, as well as application areas such as product design and medicine. Please see <http://www.apgv.org> and this years proceedings for further information, [22, 20, 21, 23, 96, 95].

## References

- [1] B. D. Adelstein, T. G. Lee, and S. R. Ellis. Head tracking latency in virtual environments: Psychophysics and a model. *In Proc. of the 47th Annual Human Factors and Ergonomics Society meeting*, pages 2083–2087, 2003.
- [2] Stéphane Albin, Gilles Rougeron, Bernard Péroche, and Alain Trémeau. Quality image metrics for synthetic images based on perceptual color differences. *IEEE Transactions on Image Processing*, 11(9):961–971, 2002.
- [3] Robert S. Allison, Laurence R. Harris, Michael Jenkin, Urszula Jasiodedzka, and James E. Zacher. Tolerance of temporal delay in virtual environments. *In VR '01: Proceedings of the Virtual Reality 2001 Conference (VR'01)*, page 247, Washington, DC, USA, 2001. IEEE Computer Society.
- [4] D. Ariely. Seeing sets: Representation by statistical properties. *Psychological Science*, 12(2):157–162, 2001.
- [5] Ronald Azuma and Gary Bishop. Improving static and dynamic registration in an optical see-through hmd. *In SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 197–204, New York, NY, USA, 1994. ACM.
- [6] J. H. Bailey and B. G. Witmer. Learning and transfer of spatial knowledge in a virtual environment. *Proc. of the Human Factors and Ergonomics Society 38th Annual Meeting*, pages 1158–1162, 1994.
- [7] Dirk Bartz, Douglas Cunningham, Jan Fischer, and Christian Wallraven. Star state of the art report the role of perception for computer graphics, 2008.
- [8] James P. Bliss, Philip D. Tidwell, and Michael A. Guest. The effectiveness of virtual reality for administering spatial navigation training to firefighters. *Presence*, 6(1):73–86, 1997.
- [9] Mark R. Bolin and Gary W. Meyer. A perceptually based adaptive sampling algorithm. *Computer Graphics*, 32(Annual Conference Series):299–309, August 1998.
- [10] M.R. Bolin and G.W. Meyer. A frequency based ray tracer. *In ACM SIGGRAPH '95 Conference Proceedings*, pages 409–418, 1995.
- [11] K. R. Brandt, J.M. Gardiner, and C. N. Macrae. The distinctiveness effect in forenames: The role of subjective experiences and recognition memory. *British Journal of Psychology*, pages 269–280, 2006.
- [12] K. R. Brandt, C.N. Macrae, A.M. Schloerscheidt, and A. B. Milne. Do i know you? target typicality and recollective experience. *Memory*, 11(1):89–100, 2003.

- [13] W. Brewer and J. Treyns. Role of schemata in memory for places. *Cognitive Psychology*, 13:207–230, 1981.
- [14] D. E. Broadbent and M. H. P. Broadbent. From detection to identification: Response to multiple targets in rapid serial visual presentation. *Perception and Psychophysics*, 42(4):105–113, 1987.
- [15] T. C. Callaghan. Dimensional interaction of hue and brightness in preattentive field segregation. *Perception & Psychophysics*, 36(1):25–34, 1984.
- [16] T. C. Callaghan. Interference and domination in texture segregation: Hue, geometric form, and line orientation. *Perception & Psychophysics*, 46(4):299–311, 1989.
- [17] T. C. Callaghan. Interference and dominance in texture segregation. In D. Brogan, editor, *Visual Search*, pages 81–87. Taylor & Francis, New York, New York, 1990.
- [18] J. Cataliotti and A. Gilchrist. Local and global processes in lightness perception. In *Perception and Psychophysics*, volume 57(2), pages 125–135. Perception, 1995.
- [19] K. Cater, A. Chalmers, and G. Ward. Detail to attention: exploiting visual tasks for selective rendering. In *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering*, pages 270–280, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association.
- [20] Heinrich Conference Chair-Bülthoff and Tom Conference Chair-Troschianko. Apgv '05: Proceedings of the 2nd symposium on applied perception in graphics and visualization. *A Coruña, Spain*, 2005.
- [21] Roland W. Conference Chair-Fleming and Sunghee Conference Chair-Kim. Apgv '06: Proceedings of the 3rd symposium on applied perception in graphics and visualization. *Boston, Massachusetts*, 2006.
- [22] Victoria Conference Chair-Interrante, Ann Conference Chair-McNamara, Heinrich Program Chair-Bulthoff, and Holly Program Chair-Rushmeier. Apgv '04: Proceedings of the 1st symposium on applied perception in graphics and visualization. *Los Angeles, California*, 2004.
- [23] Christian Conference Chair-Wallraven and Veronica Conference Chair-Sundstedt. Apgv '07: Proceedings of the 4th symposium on applied perception in graphics and visualization. *Tubingen, Germany*, 2007.
- [24] et AL. Conway, M. A. Changes in memory awareness during learning: The acquisition of knowledge by psychology undergraduates. *Journal of Experimental Psychology : General*, 126(4):393–413, 1997.
- [25] Scott Daly. The visible differences predictor: an algorithm for the assessment of image fidelity. In *Digital images and human vision*, pages 179–206, Cambridge, MA, USA, 1993. MIT Press.

- [26] Kate Devlin, Alan Chalmers, Alexander Wilkie, and Werner Purgathofer. Star: Tone reproduction and physically based spectral rendering. In Dieter Fellner and Roberto Scopigno, editors, *State of the Art Reports, Eurographics 2002*, pages 101–123, Vienna, September 2002. The Eurographics Association.
- [27] S. A. Dewhurst, S. J. Holmes, K. R. Brandt, and G. M. Dean. Measuring the speed of the conscious components of recognition memory: Remembering is faster than knowing. *Consciousness and Cognition*, 15:147–162, 2006.
- [28] Huong Q. Dinh, Neff Walker, Chang Song, Akira Kobayashi, and Larry F. Hodges. Evaluating the importance of multi-sensory input on memory and the sense of presence in virtual environments. In *VR '99: Proceedings of the IEEE Virtual Reality*, page 222, Washington, DC, USA, 1999. IEEE Computer Society.
- [29] George Drettakis, Nicolas Bonneel, Carsten Dachsbacher, Sylvain Lefebvre, Michael Schwarz, and Isabelle Viaud-Delmon. An interactive perceptual rendering pipeline using contrast and spatial masking. In *Rendering Techniques (Proceedings of the Eurographics Symposium on Rendering)*. Eurographics, June 2007.
- [30] S. R. Ellis, K. Mania, B. D. Adelstein, and M. I. Hill. Generalizability of latency detection in a variety of virtual environments. In *Proc. of the 48th Annual Human Factors and Ergonomics Society meeting*, 2004.
- [31] S. R. Ellis, M. J. Young, S. M. Ehrlich, and B. D. Adelstein. Discrimination of changes of rendering latency during voluntary hand movement. In *Proc. of the 43th Annual Human Factors and Ergonomics Society meeting*, pages 1182–1186, 1999.
- [32] S.R. Ellis, A. Wolfram, and B.D. Adelstein. Large amplitude three-dimensional tracking in augmented environments: a human performance trade-off between system latency and update rate. In *Proc. of the 46th Annual Human Factors and Ergonomics Society meeting*, 2002.
- [33] Stephen R. Ellis, F. Breant, Brian M. Menges, Richard H. Jacoby, and Bernard D. Adelstein. Operator interaction with visual objects: Effect of system latency. *HCI (2)*, pages 973–976, 1997.
- [34] Stephen R. Ellis, Mark J. Young, Bernard D. Adelstein, and Sheryl M. Ehrlich. Discrimination of changes in latency during head movement. *HCI (2)*, pages 1129–1133, 1999.
- [35] Land. M. F. Motion and vision: why animals move their eyes. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioural Physiology*, 185(4):341–352, 1999.

- [36] Jean-Philippe Farrugia and Bernard Péroche. A progressive rendering algorithm using an adaptive perceptually based image metric. *Comput. Graph. Forum*, 23(3):605–614, 2004.
- [37] J. Ferwerda. Hi-fi rendering. *ACM Siggraph Eurographics campfire on perceptually adaptive graphics*. <http://isg.cs.tcd.ie/campfire/jimferwerda2.html>, 2001.
- [38] Philip W. Fink, Patrick S. Foo, and William H. Warren. Obstacle avoidance during walking in real and virtual environments. *ACM Trans. Appl. Percept.*, 4(1):2, 2007.
- [39] R. Flannery, K.A. Wallis. How does schema theory apply to real versus virtual memories? *Cyberspsychology and Behavior*, 6(2):151–159, 2003.
- [40] Harald Frenz, Markus Lappe, Marina Kolesnik, and Thomas Bührmann. Estimation of travel distance from visual motion in virtual environments. *ACM Trans. Appl. Percept.*, 4(1):3, 2007.
- [41] J. M. Gardiner and A. Richardson-klavehn. Remembering and knowing. In: *Tulving, E. and Craik, F. I. M., eds. Handbook of Memory*, 1992.
- [42] S. Gibson and R.J. Hubbard. Efficient hierarchical refinement and clustering for radiosity in complex environments. *Computer Graphics Forum*, 15(5):297–310, 1996.
- [43] A. Gilchrist, S. Delman, and A. Jacobsen. The classification and integration of edges as critical to the perception of reflectance and illumination. *Perception and Psychophysics*, 33(5):425–436, 1983.
- [44] A. L. Gilchrist. The perception of surface blacks and whites. *Scientific American*, 240(3):88–97, March 1979.
- [45] Alan Gilchrist and Alan Jacobsen. Perception of lightness and illumination in a world of one reflectance. *Perception*, 13:5–19, 1984.
- [46] Baining Guo. Progressive radiance evaluation using directional coherence maps. In *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 255–266, New York, NY, USA, 1998. ACM.
- [47] Jörg Haber, Karol Myszkowski, Hitoshi Yamauchi, and Hans-Peter Seidel. Perceptually guided corrective splatting. *Computer Graphics Forum*, 20(3), 2001.
- [48] Christopher G. Healey and James T. Enns. Building perceptual textures to visualize multidimensional datasets. In *Proceedings Visualization '98*, pages 111–118, Research Triangle Park, North Carolina, 1998.

- [49] Christopher G. Healey and James T. Enns. Large datasets at a glance: Combining textures and colors in scientific visualization. *IEEE Transactions on Visualization and Computer Graphics*, 5(2):145–167, 1999.
- [50] David Hedley, Adam Worrall, and Derek Paddon. Selective culling of discontinuity lines. In Julie Dorsey and Philipp Slusallek, editors, *Rendering Techniques '97 (Proceedings of the Eighth Eurographics Workshop on Rendering)*, pages 69–80. Springer Wien, 1997. ISBN 3-211-83001-4.
- [51] R. Held, A. Efsathiouand, and M Greene. Adaptation to displaced and delayed visual feedback from the hand. *Journal of Experimental Psychology*, 72(6):887–891, 1966.
- [52] L. Huang and H. Pashler. A boolean map theory of visual attention. *Psychological Review*, 114(3):599–631, 2007.
- [53] L. Huang, A. Treisman, and H. Pashler. Characterizing the limits of human visual awareness. *Science*, 317:823–825, 2007.
- [54] Laurent Itti, Christof Koch, and Ernst Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. Pattern Anal. Mach. Intell.*, 20(11):1254–1259, 1998.
- [55] B. Julesz. *Foundations of Cyclopean Perception*. University of Chicago Press, Chicago, Illinois, 1971.
- [56] B. Julesz. Experiments in the visual perception of texture. *Scientific American*, 232:34–43, 1975.
- [57] B. Julesz. A theory of preattentive texture discrimination based on first-order statistics of textons. *Biological Cybernetics*, 41:131–138, 1981.
- [58] B. Julesz, E. N. Gilbert, and L. A. Shepp. Inability of humans to discriminate between visual textures that agree in second-order statistics—revisited. *Perception*, 2:391–405, 1973.
- [59] B. Julesz, E. N. Gilbert, and J. D. Victor. Visual discrimination of textures with identical third-order statistics. *Biological Cybernetics*, 31:137–140, 1978.
- [60] J. Y. Jung, B. D. Adelstein, and S. R. Ellis. Discriminability of prediction artifacts in a time delayed virtual environment. In *Proc. of the 44th Annual Human Factors and Ergonomics Society meeting*, pages 499–502, 2000.
- [61] Grzegorz Krawczyk, Karol Myszkowski, and Hans-Peter Seidel. Lightness perception in tone reproduction for high dynamic range images. In *The European Association for Computer Graphics 26th Annual Conference EUROGRAPHICS 2005*, volume 24 of *Computer Graphics Forum*, pages xx–xx, Dublin, Ireland, 2005. Blackwell.

- [62] Grzegorz Krawczyk, Karol Myszkowski, and Hans-Peter Seidel. Computational model of lightness perception in high dynamic range imaging. In Bernice E. Rogowitz, Thrasyvoulos N. Pappas, and Scott J. Daly, editors, *Human Vision and Electronic Imaging XI, IS&T SPIE's 18th Annual Symposium on Electronic Imaging (2006)*, volume xxxx, pages xxx–xxx, 2006.
- [63] Gregory Ward Larson, Holly Rushmeier, and Christine Piatko. A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):291–306, October 1997.
- [64] William B. Lathrop and Mary K. Kaiser. Perceived orientation in physical and virtual environments: changes in perceived orientation as a function of idiothetic information available. *Presence: Teleoper. Virtual Environ.*, 11(1):19–32, 2002.
- [65] Guillaume Lavoué. A roughness measure for 3d mesh visual masking. In *APGV '07: Proceedings of the 4th symposium on Applied perception in graphics and visualization*, pages 57–60, New York, NY, USA, 2007. ACM.
- [66] A. Liu, S. Tharp, L. Lai, French, and L. Stark. Some of what one needs to know about using head-mounted displays to improve teleoperator performance. *IEEE Transactions on Robotics and Automation*, 9(5):638–648, 1995.
- [67] G. Liu, E. Austen, K. Booth, B. Fischer, M. Rempel, and J. Enns. Multiple object tracking is based on scene, not retinal, coordinates. *Journal of Experimental Psychology: Human Perception and Performance*, 31(2):235–247, 2005.
- [68] G. R. Loftus and N. H. Mackworth. Cognitive determinants of fixation location during picture viewing. *Educational Psychology: Human Perception and Performance*, 4(4):565–572, 1978.
- [69] A. Mack and I. Rock. *Inattentional Blindness*. MIT Press, Menlo Park, California, 2000.
- [70] G. W. Maconcie and L.C. Loschky. Human performance with a gaze linked multi-resolutional display. In *Proceedings of the Advanced Displays and Interactive Displays First Annual Symposium 1997*, pages 25–34, 1997.
- [71] K. Mania, A. Robinson, and K. Brandt. The effect of memory schemata on object recognition in virtual environments. *Presence, Teleoperators and Virtual Environments*, 14(5):606–615, 2005.
- [72] K. Mania, T. Troscianko, R. Hawkes, and A. Chalmers. Fidelity metrics for virtual environment simulations based on human judgments of spatial

- memory awareness states. *Presence, Teleoperators and Virtual Environments*, 12(3):296–310, 2003.
- [73] Katerina Mania, Bernard D. Adelstein, Stephen R. Ellis, and Michael I. Hill. Perceptual sensitivity to head tracking latency in virtual environments with varying degrees of scene complexity. In *APGV '04: Proceedings of the 1st Symposium on Applied perception in graphics and visualization*, pages 39–47, New York, NY, USA, 2004. ACM.
- [74] Katerina Mania, Shahrul Badariah, Matthew Coxon, and Phil Watten. Cognitive transfer of spatial awareness states from immersive virtual environments to reality. *ACM Trans. Appl. Percept.*, 7(2):1–14, 2010.
- [75] Katerina Mania, Dave Wooldridge, Matthew Coxon, and Andrew Robinson. The effect of visual and interaction fidelity on spatial cognition in immersive virtual environments. *IEEE Trans. Vis. Comput. Graph.*, 12(3):396–404, 2006.
- [76] Rafal Mantiuk, Karol Myszkowski, and Hans-Peter Seidel. A perceptual framework for contrast processing of high dynamic range images. In *APGV '05: Proceedings of the 2nd symposium on Applied perception in graphics and visualization*, pages 87–94, New York, NY, USA, 2005. ACM.
- [77] G Marmitt and A. Duchowski. Modeling visual attention in vr: measuring the accuracy of predicted scanpaths. *Proceedings of Eurographics 2002, Short Presentations*, pages 217–226, 2002.
- [78] Jeffrey McCandless, Stephen R. Ellis, and Bernard D. Adelstein. Localization of a time-delayed, monocular virtual object superimposed on a real environment. *Presence*, 9(1):15–24, 2000.
- [79] B. H. McCormick, T. A. DeFanti, and M. D. Brown. Visualization in scientific computing. *Computer Graphics*, 21(6):1–14, 1987.
- [80] A. McNamara. Evaluating image quality metrics vs. human evaluation. ACM SIGGRAPH 2000 Sketches Program, 2000.
- [81] Ann McNamara. Visual perception in realistic image synthesis. *Comput. Graph. Forum*, 20(4):211–224, 2001.
- [82] Ann McNamara. Visual perception in realistic image synthesis. *Comput. Graph. Forum*, 20(4):211–224, 2001.
- [83] Michael Meehan, Sharif Razzaque, Mary C. Whitton, and Frederick P. Brooks, Jr. Effect of latency on presence in stressful virtual environments. In *VR '03: Proceedings of the IEEE Virtual Reality 2003*, page 141, Washington, DC, USA, 2003. IEEE Computer Society.
- [84] G. W. Meyer and A. Liu. Color spatial acuity control of a screen subdivision image synthesis algorithm. *Human Vision, Visual Processing, and Digital Display*, 1666(3):387–399, 1992.

- [85] Don P. Mitchell. Generating antialiased images at low sampling densities. *Computer Graphics*, 21(4):65–72, July 1987.
- [86] Yann Morvan and Carol O’Sullivan. A perceptual approach to trimming unstructured lumigraphs. In *APGV ’07: Proceedings of the 4th symposium on Applied perception in graphics and visualization*, pages 61–68, New York, NY, USA, 2007. ACM.
- [87] N. Mourkoussis, F. Rivera, T. Troscianko, T. Dixon, R. Hawkes, and K. Mania. Quantifying fidelity for virtual environment simulations employing memory schema assumptions. *ACM Transactions on Applied Perception*, 2010.
- [88] K. Myszkowski. The visible differences predictor: Applications to global illumination problems. In G. Drettakis and N. Max, editors, *Rendering Techniques ’98 (Proceedings of Eurographics Rendering Workshop ’98)*, pages 233–236, New York, NY, 1998. Springer Wien.
- [89] K. Myszkowski, A. B. Khodulev, and E. A. Kopylov. Validating global illumination algorithms and software. In *Visual Proceedings, Technical Sketch at ACM Siggraph 1997*, 1997.
- [90] U. Neisser. The control of information pickup in selective looking. In A. D. Pick, editor, *Perception and its Development: A Tribute to Eleanor J. Gibson*, pages 201–219. Lawrence Erlbaum and Associates, Hillsdale, New Jersey, 1979.
- [91] K. Nemire, R. H. Jacoby, and S. R. Ellis. Simulation fidelity of a virtual environment display. *Human Factors*, 36(1):1994, 1994.
- [92] Carol O’Sullivan, Sarah Howlett, Yann Morvan, Rachel McDonnell, and Keith O’Conor. Perceptually Adaptive Graphics. In Christophe Schlick and Werner Purgathofer, editors, *STAR-Proceedings of Eurographics 2004*, number STAR-6 in State of the Art Reports, pages 141–164. INRIA and the Eurographics Association, 2004.
- [93] A. J. Parkin, J. M. Gardiner, and R. Rosser. Functional aspects of recollective experience in face recognition. *Consciousness and Cognition*, 4(4):387–398, 1995.
- [94] S. N. Pattanaik, J. A. Ferwerda, D. A. Greenberg, and M. D. Fairchild. A multiscale model of adaptation and spatial vision for realistic imaging. In *Computer Graphics (ACM SIGGRAPH ’98 Proceedings)*, pages 287–298, 1998.
- [95] Bobby Program Chair-Bodenheimer, Carol Program Chair-O’Sullivan, Katerina Conference Chair-Mania, and Bernhard Conference Chair-Riecke. *Apgv ’09: Proceedings of the 6th symposium on applied perception in graphics and visualization. Chania, Crete, Greece*, 2009.

- [96] Sarah Program Chair-Creem-Regehr and Karol Program Chair-Myszkowski. Apgv '08: Proceedings of the 5th symposium on applied perception in graphics and visualization. *Los Angeles, California*, 2008.
- [97] P. Rademacher, J. Lengyel, E. Cutrell, and T. Whitted. Measuring the perception of visual realism in images, 2001.
- [98] Ganesh Ramanarayanan, Kavita Bala, and James A. Ferwerda. Perception of complex aggregates. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, pages 1–10, New York, NY, USA, 2008. ACM.
- [99] Ganesh Ramanarayanan, James Ferwerda, Bruce Walter, and Kavita Bala. Visual equivalence: towards a new standard for image fidelity. In *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, page 76, New York, NY, USA, 2007. ACM.
- [100] Mahesh Ramasubramanian, Sumanta N. Pattanaik, and Donald P. Greenberg. A perceptually based physical error metric for realistic image synthesis. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 73–82, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.
- [101] J. E. Raymond, K. L. Shapiro, and K. M. Arnell. Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception & Performance*, 18(3):849–860, 1992.
- [102] Matthew J. P. Regan, Gavin S. P. Miller, Steven M. Rubin, and Chris Kogelnik. A real-time low-latency hardware light-field renderer. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 287–290, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.
- [103] Erik Reinhard, Michael Stark, Peter Shirley, and James Ferwerda. Photographic tone reproduction for digital images. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 267–276, New York, NY, USA, 2002. ACM Press.
- [104] R. A. Rensink, J. K. O'Regan, and J. J. Clark. To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, 8:368–373, 1997.
- [105] Ronald A. Rensink. Seeing, sensing, and scrutinizing. *Vision Research*, 40(10-12):1469–1487, 2000.
- [106] James C. Rodger and Roger A. Browse. Choosing rendering parameters for effective communication of 3d shape. *IEEE Computer Graphics and Applications*, 20(2):20–28, 2000.

- [107] Thomas B. Sheridan. Remote manipulative control with transmission delay. *IEEE Transactions on Human Factors in Electronics*, 4(1):25–29, 1963.
- [108] Thomas B. Sheridan. Musings on telepresence and virtual presence. *Presence*, 1(1):120–125, 1992.
- [109] D. J. Simons and R. A. Rensink. Change blindness: Past, present, and future. *Trends in Cognitive Science*, 9(1):16–20, 2005.
- [110] Daniel J. Simons. Current approaches to change blindness. *Visual Cognition*, 7(1/2/3):1–15, 2000.
- [111] P. H. Smith and J. Van Rosendale. Data and visualization corridors report on the 1998 CVD workshop series (sponsored by DOE and NSF). Technical Report CACR-164, Center for Advanced Computing Research, California Institute of Technology, 1998.
- [112] William A. Stokes, James A. Ferwerda, Bruce J. Walter, and Donald P. Greenberg. Perceptual illumination compone high quality global illumination rendering. *ACM Transactions on Graphics*, 23(3):742–749, 2004.
- [113] Veronica Sundstedt, Efstathios Stavrakis, Michael Wimmer, and Erik Reinhard. A psychophysical study of fixation behavior in a computer game. In *APGV '08: Proceedings of the 5th symposium on Applied perception in graphics and visualization*, pages 43–50, New York, NY, USA, 2008. ACM.
- [114] J. J. Thomas and K. A. Cook. *Illuminating the Path: Research and Development Agenda for Visual Analytics*. IEEE Press, Piscataway, New Jersey, 2005.
- [115] A. Treisman. Preattentive processing in vision. *Computer Vision, Graphics, and Image Processing*, 31:156–177, 1985.
- [116] A. Treisman. Search, similarity, and integration of features between and within dimensions. *Journal of Experimental Psychology: Human Perception & Performance*, 17(3):652–676, 1991.
- [117] A. Treisman and S. Gormican. Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, 95(1):15–48, 1988.
- [118] E. Tulving. *Elements of Episodic Memory*. Oxford : Oxford Science Publications, 1992.
- [119] Jack Tumblin and Holly E. Rushmeier. Tone Reproduction for Realistic Images. *IEEE Computer Graphics and Applications*, 13(6):42–48, November 1993.

- [120] S. Uno and M. Slater. The sensitivity of presence to collision response. In *VRAIS '97: Proceedings of the 1997 Virtual Reality Annual International Symposium (VRAIS '97)*, page 95, Washington, DC, USA, 1997. IEEE Computer Society.
- [121] C.J. van den Branden Lambrecht. *Perceptual models and architectures for video coding applications*. PhD thesis, Ecole Polytechnique Federal de Lausanne, 1996.
- [122] David Waller, Earl Hunt, and David Knapp. The transfer of spatial knowledge in virtual environment training. *Presence: Teleoper. Virtual Environ.*, 7(2):129–143, 1998.
- [123] B.J. Walter. *Density estimation techniques for global illumination*. Ph.D. thesis, Cornell University, 1998.
- [124] Leonard Wanger. The effect of shadow quality on the perception of spatial relationships in computer generated imagery. In *I3D '92: Proceedings of the 1992 symposium on Interactive 3D graphics*, pages 39–42, New York, NY, USA, 1992. ACM.
- [125] Colin Ware. *Information Visualization: Perception for Design, 2nd Edition*. Morgan Kaufmann Publishers, Inc., San Francisco, California, 2004.
- [126] Benjamin Watson, Neff Walker, Larry F. Hodges, and Martin Reddy. An evaluation of level of detail degradation in head-mounted display peripheries. *Presence*, 6(6):630–637, 1997.
- [127] Benjamin Watson, Neff Walker, Peter Woytiuk, and William Ribarsky. Maintaining usability during 3d placement despite delay. In *VR '03: Proceedings of the IEEE Virtual Reality 2003*, page 133, Washington, DC, USA, 2003. IEEE Computer Society.
- [128] R. B. Welch, T. T. Blackmon, A. Liu, B. A. Mellers, and L. W. Stark. The effects of pictorial realism, delay of visual feedback and observer interactivity on the subjective sense of presence. *Presence, Teleoperators and Virtual Environments*, 5(3):263–273, 1996.
- [129] Betsy Williams, Gayathri Narasimham, Claire Westerman, John Rieser, and Bobby Bodenheimer. Functional similarities in spatial representations between real and virtual environments. *ACM Trans. Appl. Percept.*, 4(2):12, 2007.
- [130] Jeremy M. Wolfe. Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review*, 1(2):202–238, 1994.
- [131] Jeremy M. Wolfe and Kyle R. Cave. Deploying visual attention: The Guided Search model. In T. Troscianko and A. Blake, editors, *AI and the Eye*, pages 79–103. John Wiley & Sons, Inc., Chichester, United Kingdom, 1989.

- [132] Jeremy M. Wolfe, Kyle R. Cave, and Susan L. Franzel. Guided Search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception & Performance*, 15(3):419–433, 1989.
- [133] P Zimmons. *The Influence of Lighting Quality on Presence and task Performance in Virtual Environment*, Unpublished PhD thesis. PhD thesis, University of North Carolina, <ftp://ftp.cs.unc.edu/pub/publications/techreports/04-017.pdf>, 2005.
- [134] Alexandros Zotos, Katerina Mania, and Nicholas Mourkoussis. A schema-based selective rendering framework. In *APGV '09: Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization*, pages 85–92, New York, NY, USA, 2009. ACM.

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

Christopher G. Healey and James T. Enns

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

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

## I. INTRODUCTION

THIS paper investigates the problem of visualizing multivariate data elements arrayed across an underlying height field. We seek a flexible method of displaying effectively large and complex datasets that encode multiple data values at a single spatial location. Examples include visualizing geographic and environmental conditions on topographical maps, representing surface locations, orientations, and material properties in medical volumes, or displaying rigid and rotational velocities on the surface of a three-dimensional object. Currently, features like hue, intensity, orientation, motion, and isocontours are used to represent these types of datasets. We are investigating the simultaneous use of perceptual textures and colors for multivariate visualization. We believe an effective combination

of these features will increase the number of data values that can be shown at one time in a single display. To do this, we must first design methods for building texture and color patterns that support the rapid, accurate, and effective visualization of multivariate data elements.

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

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

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

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

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

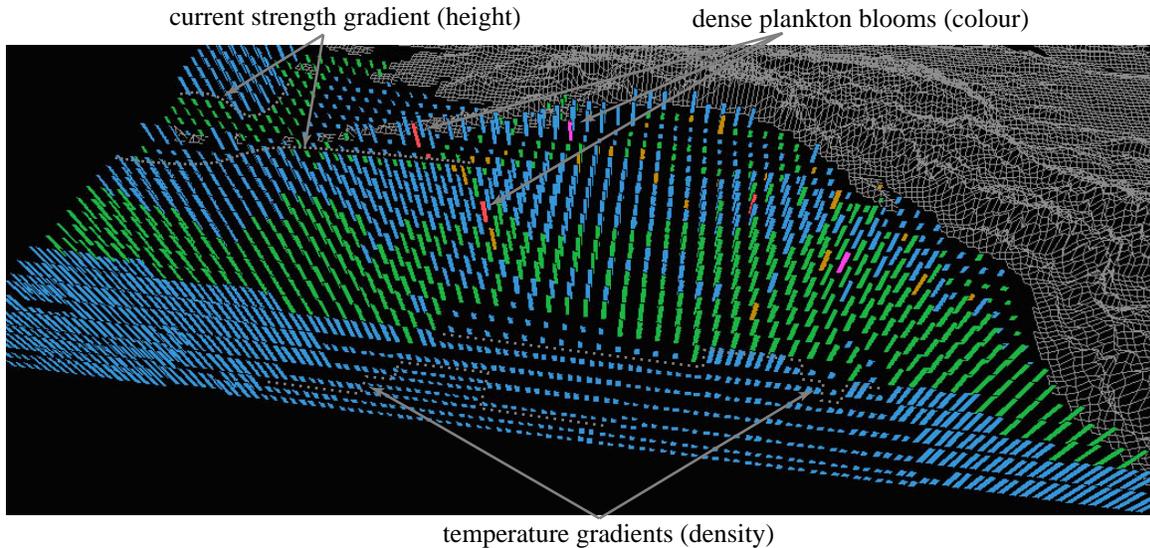


Fig. 1. Color, height, and density used to visualize open-ocean plankton density, ocean current strength, and sea surface temperature, respectively; low to high plankton densities represented with blue, green, brown, red, and purple, stronger currents represented with taller pixels, and warmer temperatures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

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

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

visually explore a multivariate dataset in a rapid, accurate, and relatively effortless fashion.

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

## II. RELATED WORK

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

### A. Texture

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

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

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

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

Work in computer vision is also interested in how viewers segment images, in part to try to develop automated texture classification and segmentation algorithms. Tamura et al. and Rao and Lohse identified texture dimensions by conducting experiments that asked observers to divide pictures depicting different types of textures (Brodatz images) into groups [39], [40], [50]. Tamura et al. used their results to propose methods for measuring coarseness, contrast, directionality, line-likeness, regularity, and roughness. Rao and Lohse used multidimensional scaling to identify the primary texture dimensions used by their observers to group images: regularity, directionality, and complexity. Haralick built grayscale spatial dependency matrices to identify

features like homogeneity, contrast, and linear dependency [21]. These features were used to classify satellite images into categories like forest, woodlands, grasslands, and water. Liu and Picard used Wold features to synthesize texture patterns [35]. A Wold decomposition divides a 2D homogeneous pattern (*e.g.*, a texture pattern) into three mutually orthogonal components with perceptual properties that roughly correspond to periodicity, directionality, and randomness. Malik and Perona designed computer algorithms that use orientation filtering, nonlinear inhibition, and computation of the resulting texture gradient to mimic the discrimination ability of the low-level human visual system [37].

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

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

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

### B. Color

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

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

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

Levkowitz and Herman designed a locally optimal color scale that maximizes the just-noticeable color difference between neighboring pairs of colors [34]. The result is a significantly larger number of just-noticeably different colors in their color scales, compared to standard scales like red-blue, rainbow, or luminance.

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

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

### C. Combined Texture and Color

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

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

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

### III. PERCEPTUAL VISUALIZATION

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

Bottom-up factors in the control of attention include the limited set of features that psychophysicists have identified as being detected very quickly by the human visual system, without the need for search. These features are often called preattentive, because their detection occurs rapidly and accurately, usually in an amount of time independent of the total number of elements being displayed. When

applied properly, preattentive features can be used to perform different types of exploratory analysis. Examples include searching for data elements with a unique visual feature, identifying the boundaries between groups of elements with common features, tracking groups of elements as they move in time and space, and estimating the number of elements with a specific feature. Preattentive tasks can be performed in a single glance, which corresponds to 200 milliseconds (ms) or less. As noted above, the time required to complete the task is independent of the number of data elements being displayed. Since the visual system cannot choose to refocus attention within this timeframe, users must complete their task using only a “single glance” at the image.

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

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

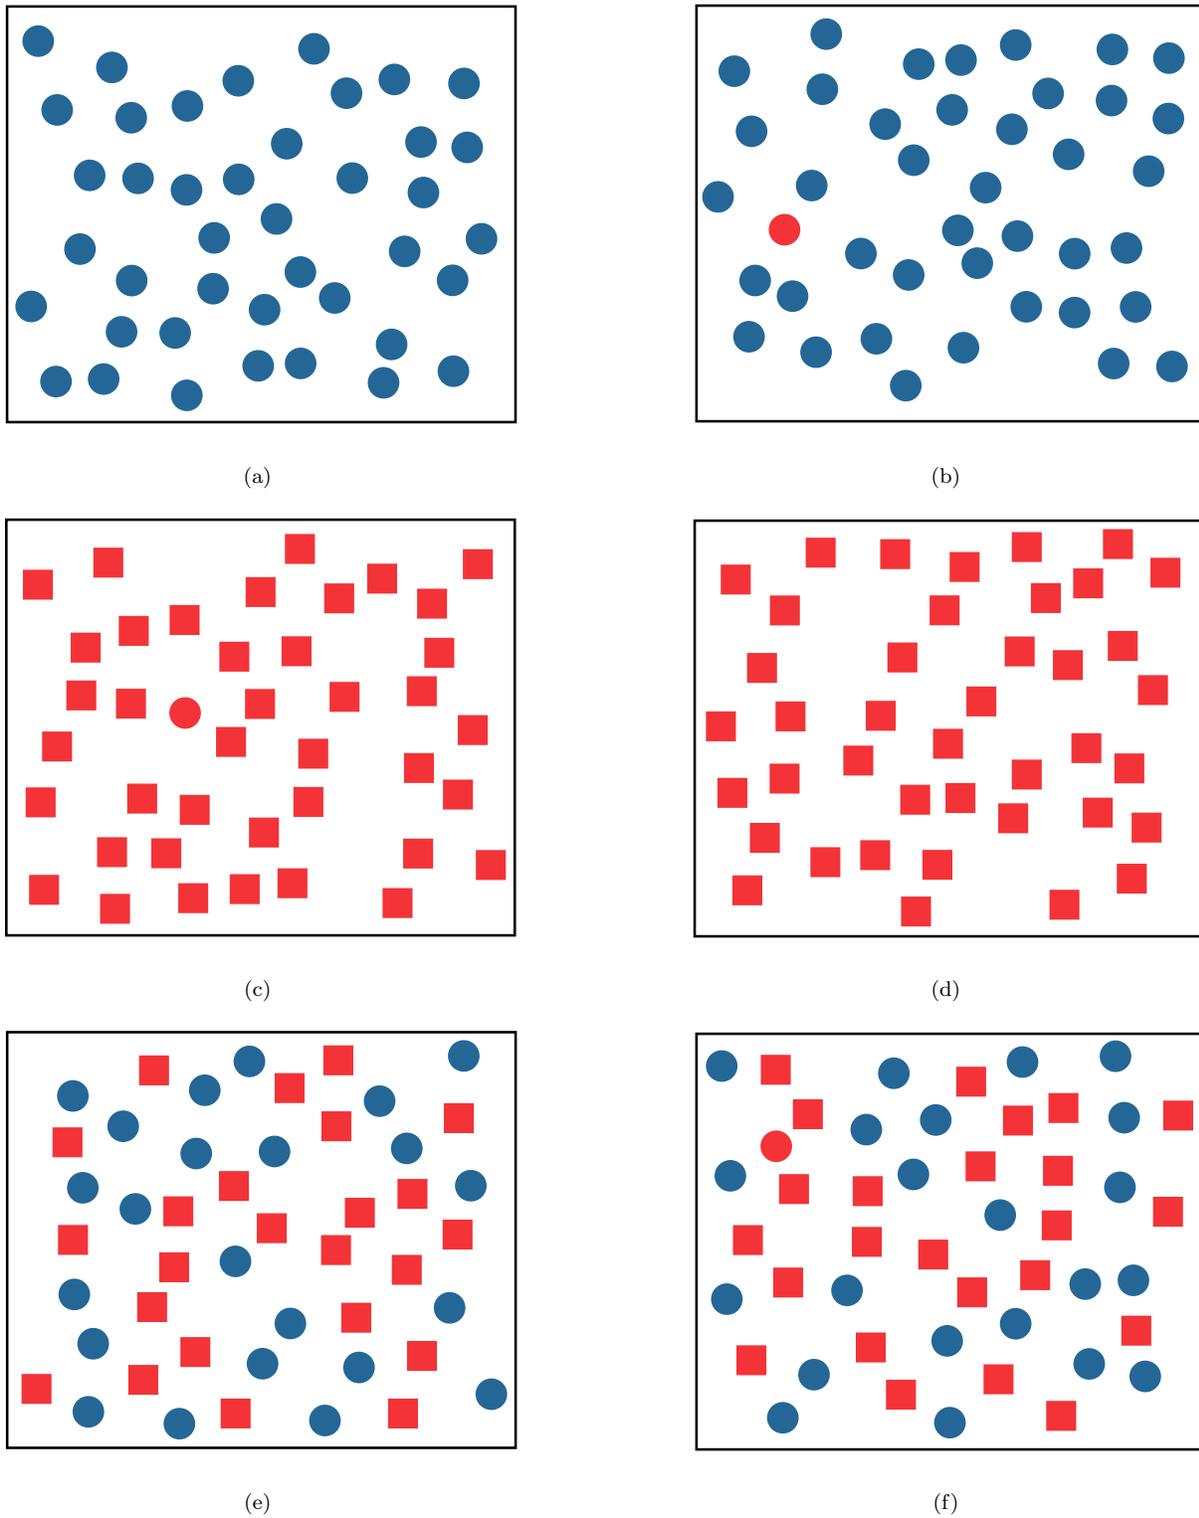


Fig. 2. Examples of target search: (a, b) identifying a red target in a sea of blue distracters is rapid and accurate, target absent in (a), present in (b); (c, d) identifying a red circular target in a sea of red square distracters is rapid and accurate, target present in (c), absent in (d); (e, f)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

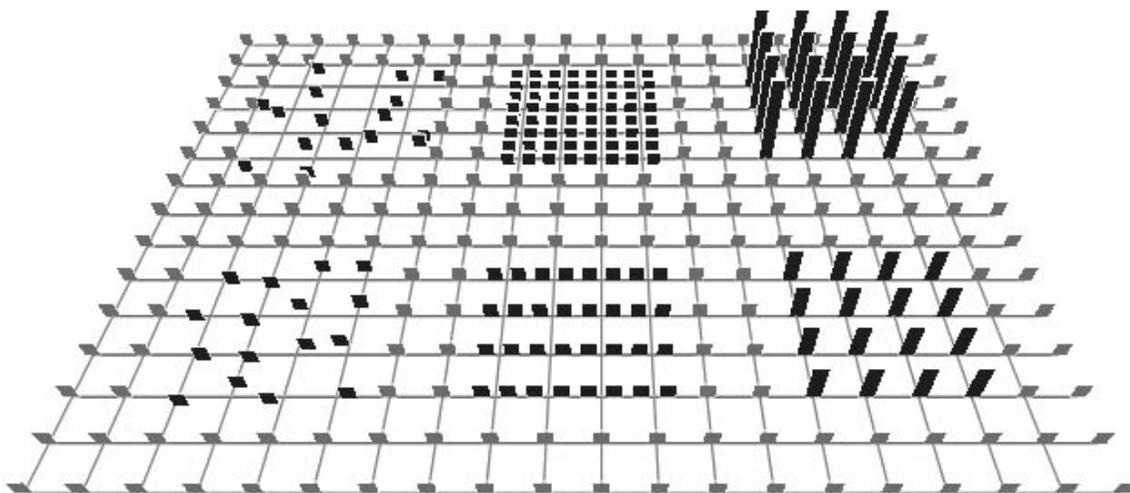


Fig. 3. A background array of short, sparse, regular pexels; the lower and upper groups on the left contain irregular and random pexels, respectively; the lower and upper groups in the center contain dense and very dense pexels; the lower and upper groups to the right contain medium and tall pexels

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

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

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

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

We are currently investigating the combined use of two important and commonly used visual features: texture

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

#### IV. PEXELS

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

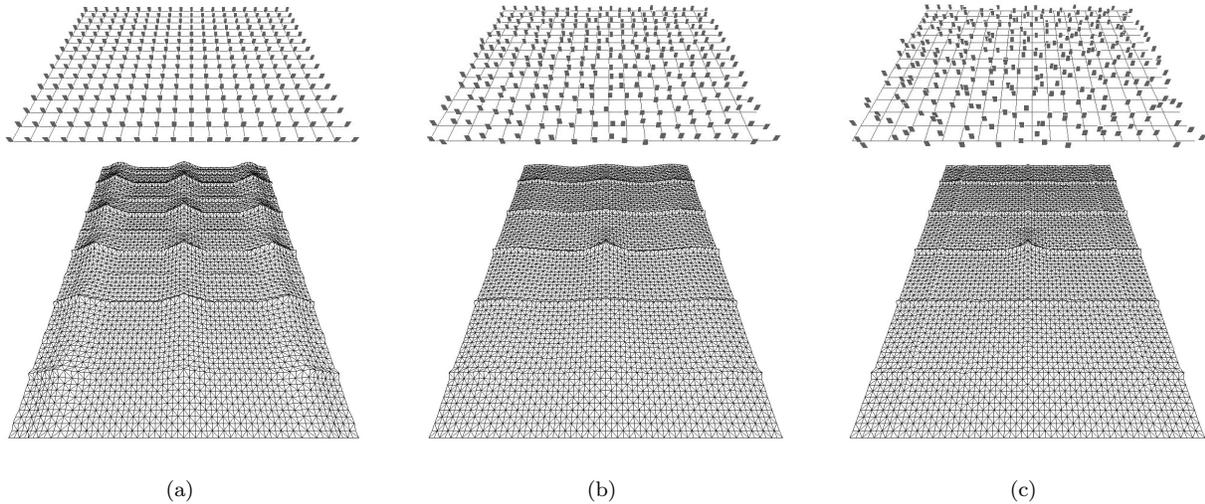


Fig. 4. Three displays of pexels with different regularity and a  $5 \times 3$  patch from the center of the corresponding autocorrelation graphs: (a) a completely regular display, resulting in sharp peaks of height 1.0 at regular intervals in the autocorrelation graph; (b) a display with irregularly-spaced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

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

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

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

#### A. Ordering Texture Dimensions

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

Although viewers can often order regularity intuitively, we required a more formal method for measurement. We chose to use autocorrelation to rank regularity. This technique measures the second-order statistic of a texture pat-

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

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

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

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

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

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

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

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

### B. Pexel Saliency and Interference

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

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

### C. Experiments

We designed texture displays to test the detectability of six different target types: taller, shorter, denser, sparser,

more regular, and more irregular. For each target type, a number of parameters were varied, including exposure duration, texture dimension saliency, and visual interference. For example, during the “taller” experiment, each display showed a  $20 \times 15$  array of pexels rotated  $45^\circ$  about the X-axis. Observers were asked to determine whether the array contained a group of pexels that were taller than all the rest. The following conditions varied:

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

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

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

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

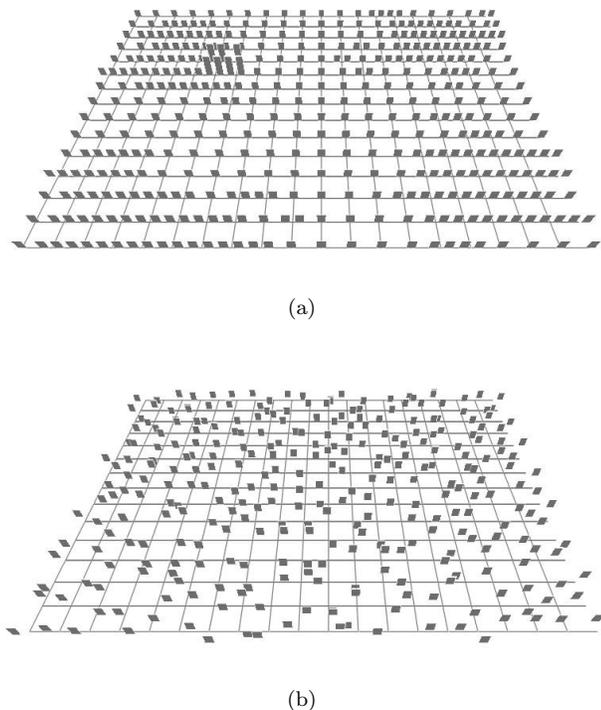


Fig. 5. Two display types from the taller and regular pexel experiments: (a) a target of medium pexels in a sea of short pexels with a background density pattern ( $2 \times 2$  target group located left of center); (b) a target of regular pexels in a sea of irregular pexels with no background texture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

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

#### D. Results

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

1. Taller target regions were identified rapidly (*i.e.*, 150 ms or less) with very high accuracy (approximately 93%); background density and regularity patterns produced no significant interference.
2. Shorter, denser, and sparser targets were more difficult to identify than taller targets, although good results were obtained at both 150 and 450 ms (82.3%, 94.0%, and 94.7% for shorter, denser, and sparser targets with no background variation at 150 ms). This was not surprising, since similar

results have been documented by [51] and [1] using displays of texture on a two-dimensional plane.

3. Background variation in non-target attributes produced small, but statistically significant, interference effects. These effects tended to be largest when target detectability was lowest. For example, density and regularity interfered with searching for shorter targets; height and regularity interfered with searching for sparser targets; but only height interfered with the (easier to find) denser targets.

4. Irregular target regions were difficult to identify at 150 and 450 ms, even with no secondary texture pattern (approximately 76%). Whether this accuracy level is sufficiently high will depend on the application. In contrast, regular regions were invisible under these conditions; the percentage of correct responses approached chance (*i.e.*, 50%) in every condition.

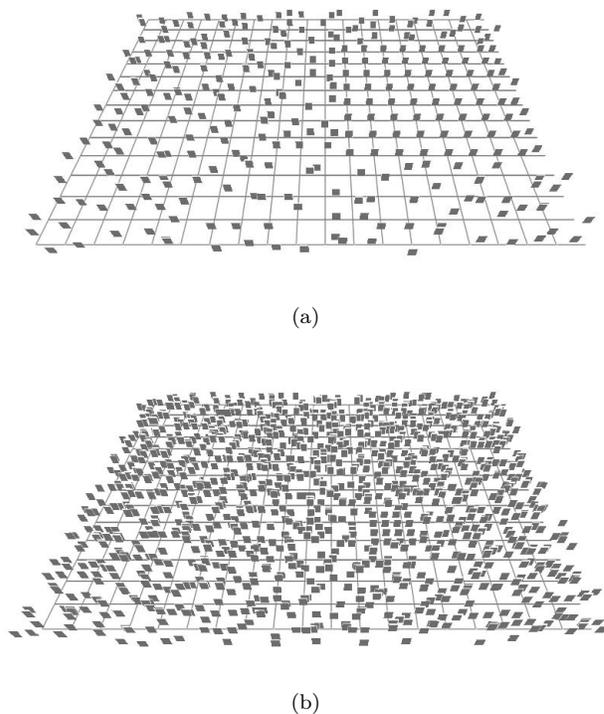


Fig. 6. Two displays with a regular target, both displays should be compared with the target shown in Fig. 5b: (a) larger target, an  $8 \times 8$  target in a sea of sparse, random pexels; (b) denser background, a  $2 \times 2$  target in a sea of dense, random pexels (target group located right of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

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

Regularity	88.3%	66.5%	80.4%	68.8%		
Density	87.4%	75.9%			55.9%	68.6%
Height			64.1%	77.2%	53.7%	58.5%
None	93.1%	83.7%	93.8%	93.4%	49.3%	76.8%
<b>Background:</b>	Taller	Shorter	Denser	Sparser	Regular	Random
<b>Target:</b>						

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

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

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

## V. PERCEPTUAL COLORS

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

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

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

### A. Color Distance

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

$$\Delta E_{xy}^* = \sqrt{(\Delta L_{xy}^*)^2 + (\Delta u_{xy}^*)^2 + (\Delta v_{xy}^*)^2} \quad (5)$$

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

### B. Linear Separation

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

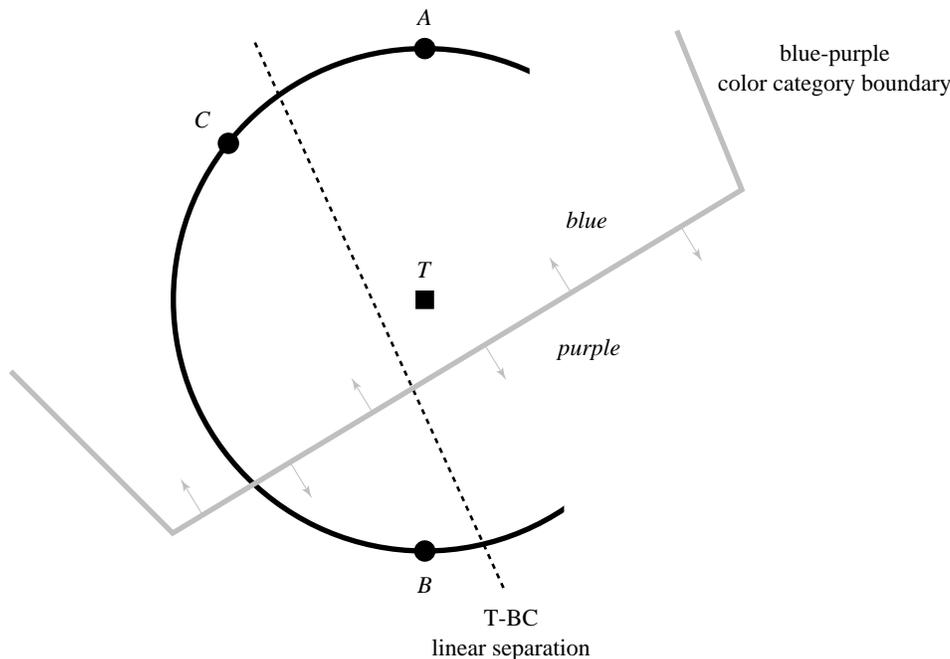


Fig. 8. A small, isoluminant patch within the CIE LUV color model, showing a target color  $T$  and three background distracter colors  $A$ ,  $B$ , and  $C$ ; note that  $T$  is collinear with  $A$  and  $B$ , but can be separated by a straight line from  $B$  and  $C$ ; note also that  $T$ ,  $A$ ,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

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

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

### C. Color Category

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

Consider again the target color  $T$  and two background distracter colors  $A$  and  $B$  shown in CIE LUV space in Fig. 8. Notice also that this region of color space has been divided into two named color categories. As before, the Euclidean distances  $TA$  and  $TB$  are equal, yet finding an

element colored  $T$  in a sea of background elements colored  $A$  is significantly more difficult than finding  $T$  in a sea of elements colored  $B$ . Kawai et al. suggest this is because both  $T$  and  $A$  lie within a color category named “blue”, while  $B$  lies within a different category named “purple”. Colors from different named categories have a higher perceived color difference, even when a perceptual color model is used to try to control that difference.

### D. Color Selection Experiments

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

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

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

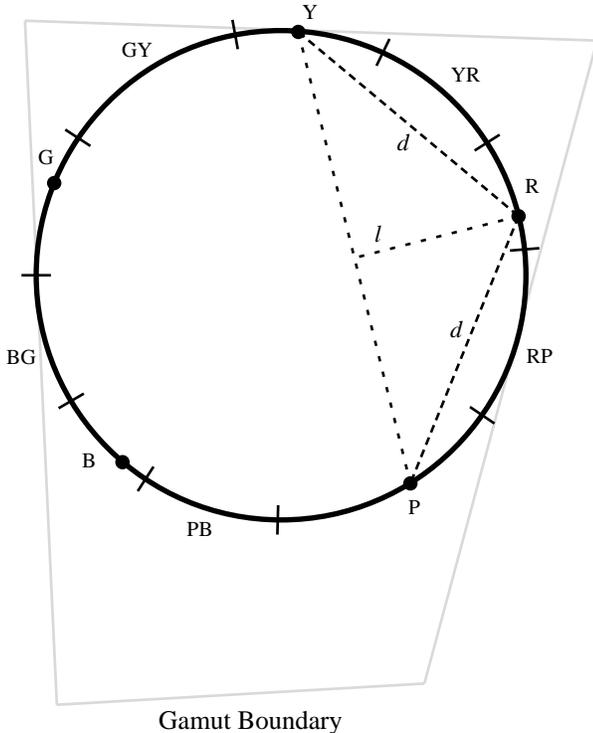


Fig. 9. Choosing colors from the monitor’s gamut, the boundary of the gamut at  $L^* = 61.7$  represented as a quadrilateral, along with the maximum inscribed circle centered at  $(L^*, u^*, v^*) = (67.1, 13.1, -0.98)$ , radius  $70.5\Delta E^*$ ; five colors chosen around the circle’s circumference;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

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

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

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

### E. Results

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

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

### F. Color Category Experiments

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

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

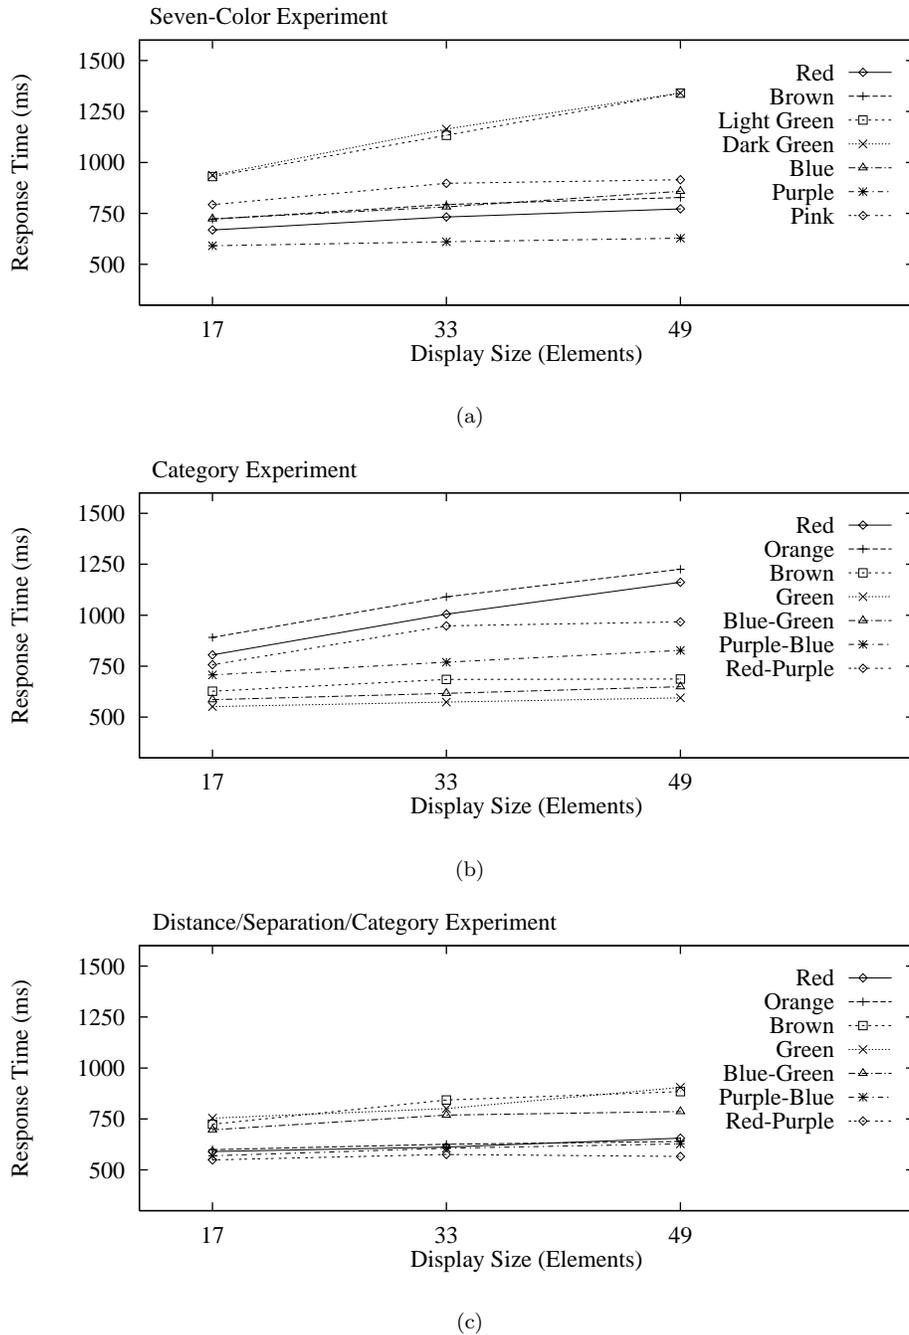


Fig. 10. Graphs showing averaged subject response times for three of the six studies: (a) response time as a function of display size (*i.e.*, total number of elements shown in each display) for each target from the seven-color study; (b) response times for each target from the color category study;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

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

Consider Table I, which lists the percentage of observers who selected a particular name for six of the representative colors. For example, the table shows that representative colors from P and R overlap only at the “pink” name. Their overlap is not that strong, since neither P nor R

are strongly classified as pink. The amount of overlap is computed by multiplying the percentages for the common name, giving a P-R overlap of  $5.2\% \times 26.3\% = 0.014$ . A closer correspondence of user-chosen names for a pair of regions results in a stronger category similarity. For example, nearly all observers named the representative colors from the G and GY regions as “green”. This resulted in an overlap of 0.973. Representative colors that overlap over multiple names are combined using addition, for example, YR and Y overlapped in both the “orange” and “brown”

TABLE I

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

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

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

### G. Color Category Results

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

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

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

cluded that up to seven isoluminant colors can be displayed simultaneously while still allowing for rapid and accurate identification, but only if the colors satisfy proper color distance, linear separation, and color category guidelines.

## VI. COMBINING TEXTURE AND COLOR

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

### A. Experiments

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

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

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

For color targets, we varied:

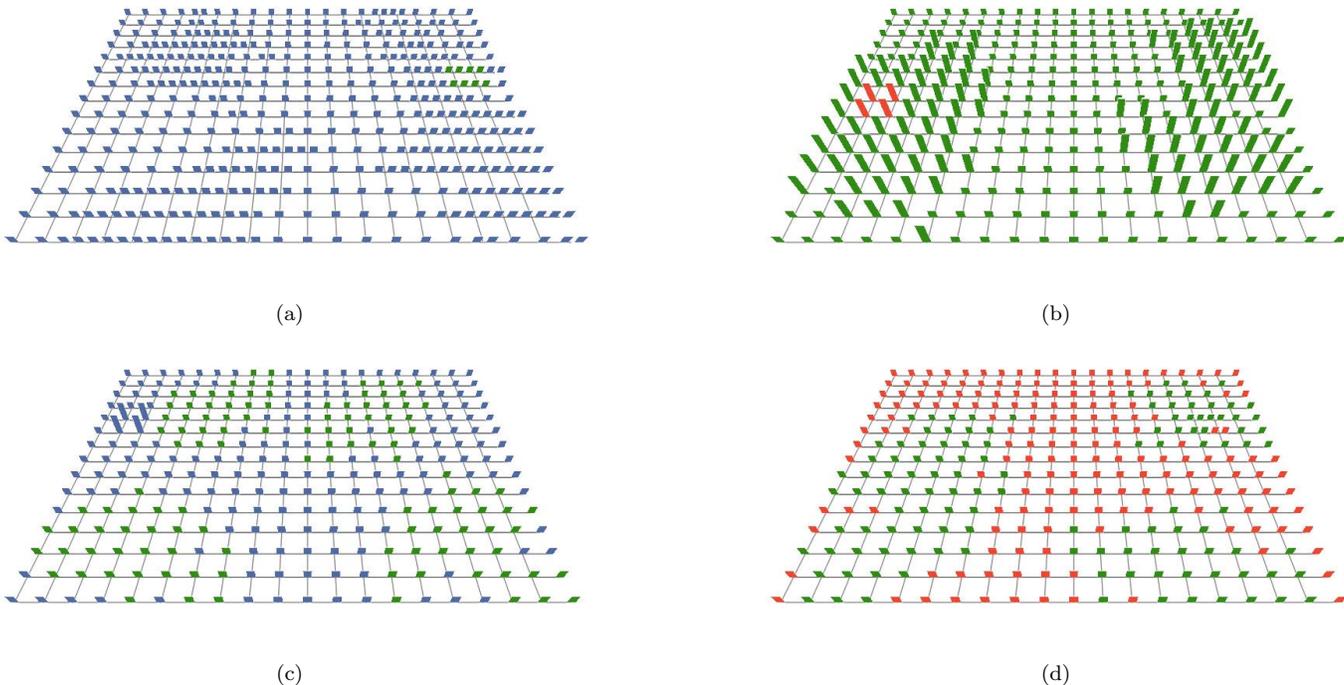


Fig. 11. Four displays from the combined color-texture experiments, printed colors may not match exactly on-screen colors used during our experiments: (a) a green target in a sea of blue pixels with background density variation; (b) a red target in a sea of green pixels with background height variation;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

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

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

- *target-background pairing*: half the displays contained a target region of medium pixels in a sea of short pixels, while the other half contained a target region of tall pixels in a sea of medium pixels; two different pairings were used to increase the generality of the results,

- *secondary dimension*: the displays contained pixels that were either a constant gray or that varied randomly between two colors; when color was varied, half the displays contained blue and green pixels, while the other half of the displays contained green and red pixels (Fig. 11c),
- *exposure duration*: displays were shown for 50, 150, or 450 ms, and
- *target patch size*: target groups were either  $2 \times 2$  pixels or  $4 \times 4$  pixels in size.

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

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

Color	96.5%	75.7%	89.1%	85.8%	
Density					95.5%
Height					95.4%
None	93.1%	83.7%	93.8%	93.4%	93.8%
<b>Background:</b>	Taller	Shorter	Denser	Sparser	Color
<b>Target:</b>					

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

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

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

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

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

Ten users participated as observers in each of the two color and four texture experiments. Each observer had normal or corrected acuity. Observers who completed the

color experiments were also tested for color blindness [12]. Observers were provided with an opportunity to practice before each experiment. This helped them become familiar with the task and the duration of the displays. Before each testing session began, observers were told that half the displays would contain a target, and half would not. We used a Macintosh computer with an 8-bit color lookup table to run our experiments. Responses (either “target present” or “target absent”) for each display an observer was shown were recorded for later analysis.

## B. Results

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

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

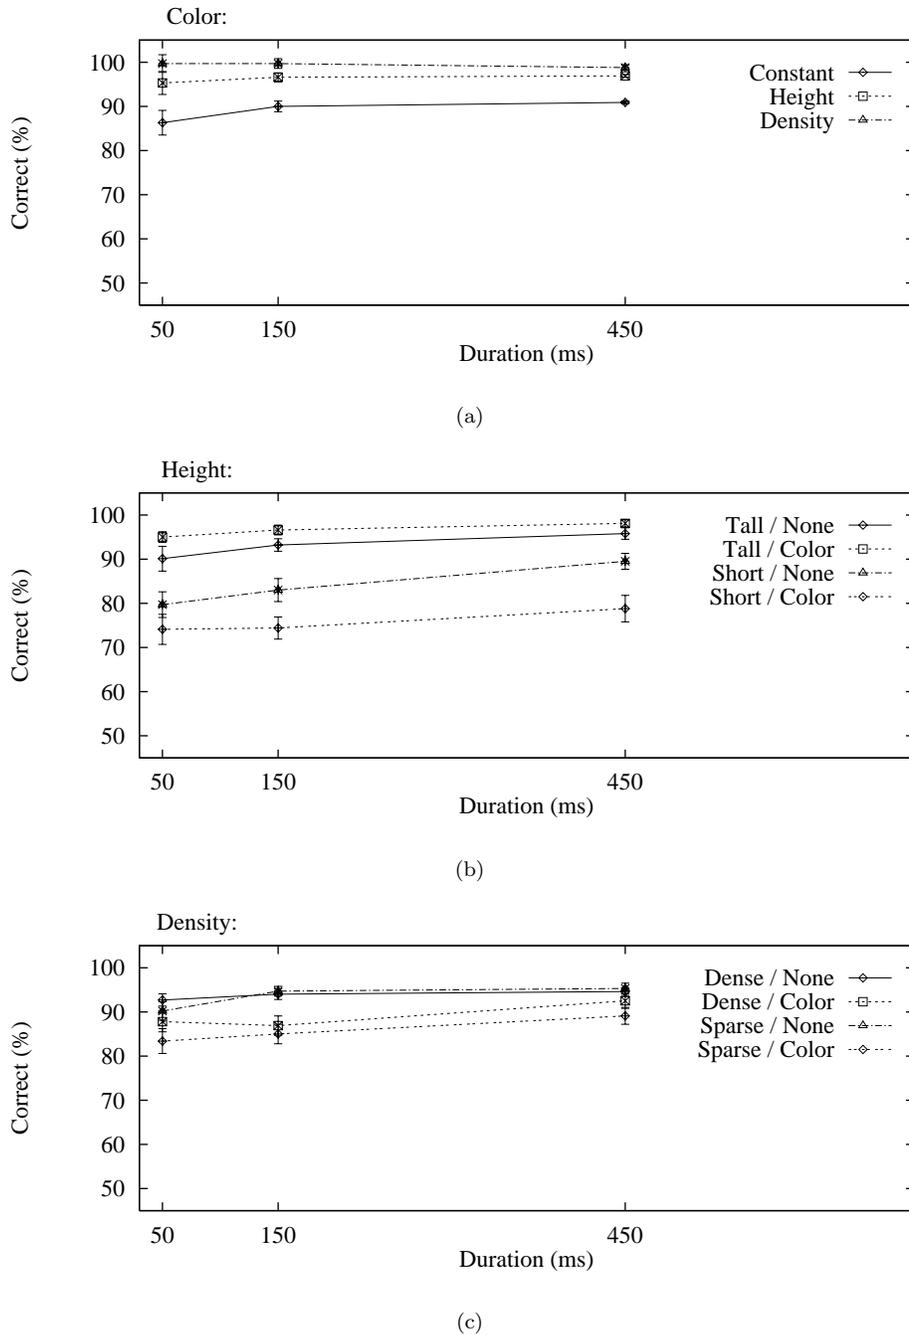


Fig. 13. Graphs showing averaged subject results for color, height, and density trials: (a) results for color trials, horizontal axis plots exposure duration, vertical axis plots percentage of correct responses, each line corresponds to one of the three different background conditions (no variation,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

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

Fig. 12 shows average subject performance as a table representing each target-background pair. Target type varies along the horizontal axis, while background type varies

along the vertical axis. Darker squares represent target-background pairings with highly accurate subject performance. The number in the center of each square reports the percentage of correct responses averaged across all subjects.

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

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

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

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

cases where the detection task is relatively effortless prior to the addition of color variation. As can be seen in Fig. 13b and 13c, the interference effect of color variation tends to be greatest when the target detection task is most difficult.

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

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

## VII. PRACTICAL APPLICATIONS

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

### A. Visualizing Typhoons

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

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

<sup>1</sup><http://grid2.cr.usgs.gov/dem/>

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

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

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

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

Precipitation is represented with color. We used our perceptual color selection technique to choose five perceptually uniform colors. Daily precipitation readings of zero, 0–0.03 inches, 0.03–0.4 inches, 0.4–1.0 inches, and 1.0–10.71 inches were colored green, yellow, orange, red, and purple, respectively (each precipitation range had an equal number of entries in our typhoon dataset). Pexels on the open

ocean or at weather stations where no precipitation values were reported were colored blue-green. Our experimental results showed no texture-on-color interference. Moreover, our color selection technique is designed to produce colors that are equally distinguishable from one another. Our mapping uses red and purple to highlight the high-precipitation areas associated with typhoon activity.

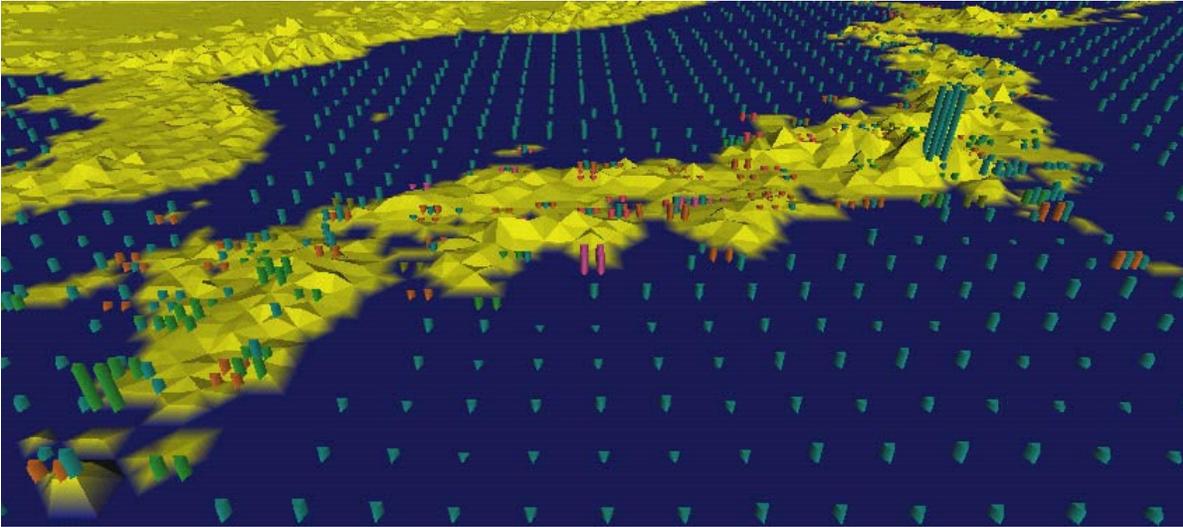
We should note that our data-feature mapping is designed to allow viewers to rapidly and accurately identify and track the locations of storms and typhoons as spatial collections of tall, dense, red and purple pexels. Our visualization system is not meant to allow users to determine exact values of windspeed,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

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

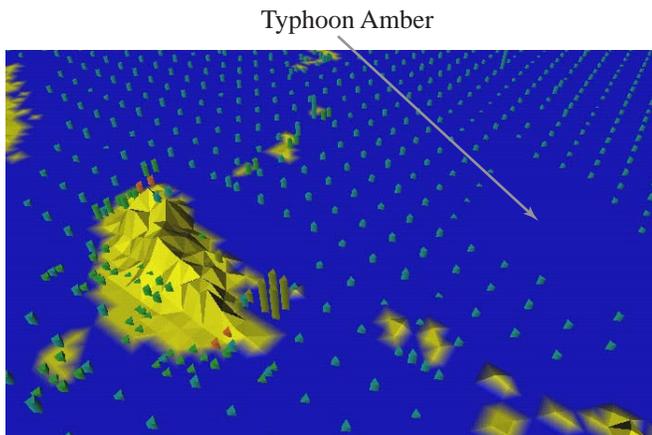
Fig. 14 shows windspeed, pressure, and precipitation around Japan, Korea, and Taiwan during August 1997. Fig. 14a looks north, and displays normal summer conditions across Japan on August 7, 1997. Fig. 14b, looking northeast, tracks typhoon Amber (one of the region's major typhoons) approaching along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

<sup>2</sup><http://www.ncdc.noaa.gov/ol/climate/online/g sod.html>

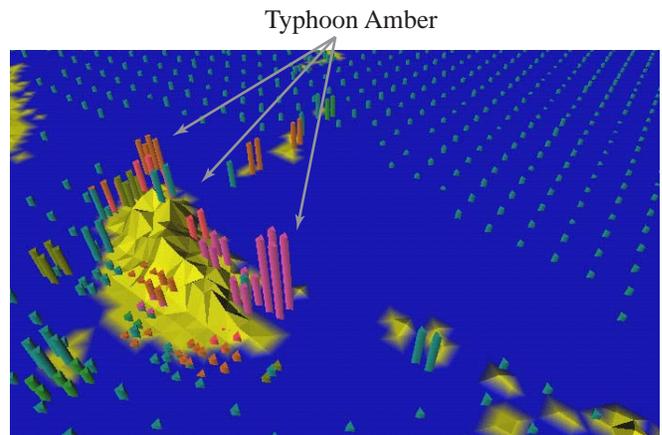
<sup>3</sup><http://ghrc.msfc.nasa.gov/ghrc/list.html>



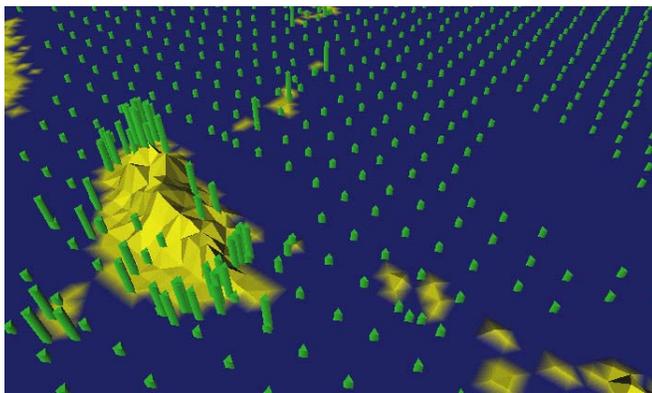
(a)



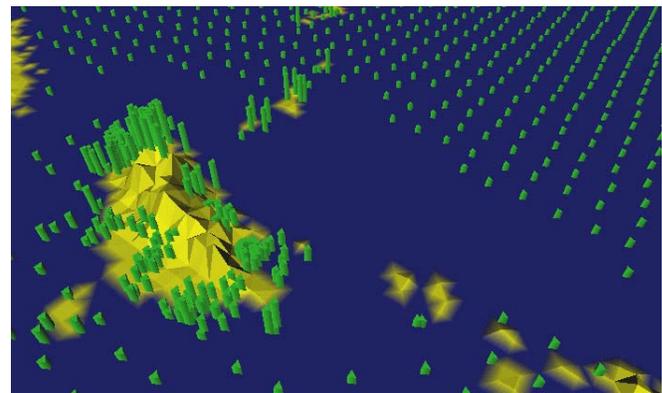
(b)



(c)



(d)



(e)

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

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

### VIII. CONCLUSIONS AND FUTURE WORK

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

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

One important area of future work is a comparison of our visualization techniques against other methods that might be used to represent information in our real-world applications. For example, it would be useful to test a user's ability to track storm activity in our visualization environment against other standard techniques for representing weather

activity. Although we have yet to conducted these kinds of practical experiments, we hope to initiate them in the near future as part of our perceptual visualization studies.

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

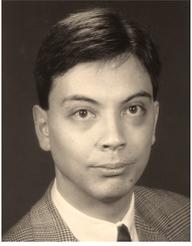
### ACKNOWLEDGMENTS

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

### REFERENCES

- [1] AKS, D. J., AND ENNS, J. T. Visual search for size is influenced by a background texture gradient. *Journal of Experimental Psychology: Perception and Performance* 22, 6 (1996), 1467–1481.
- [2] BANKS, W. P., AND PRINZMETAL, W. Configurational effects in visual information processing. *Perception & Psychophysics* 19 (1976), 361–367.
- [3] BAUER, B., JOLICOEUR, P., AND COWAN, W. B. Visual search for colour targets that are or are not linearly-separable from distractors. *Vision Research* 36 (1996), 1439–1446.
- [4] BAUER, B., JOLICOEUR, P., AND COWAN, W. B. The linear separability effect in color visual search: Ruling out the additive color hypothesis. *Perception & Psychophysics* 60, 6 (1998), 1083–1093.
- [5] BERGMAN, L. D., ROGOWITZ, B. E., AND TREINISH, L. A. A rule-based tool for assisting colormap selection. In *Proceedings Visualization '95* (Atlanta, Georgia, 1995), pp. 118–125.
- [6] BRUCKNER, L. A. On Chernoff faces. In *Graphical Representation of Multivariate Data*, P. C. C. Wang, Ed. Academic Press, New York, New York, 1978, pp. 93–121.
- [7] CALLAGHAN, T. C. Dimensional interaction of hue and brightness in preattentive field segregation. *Perception & Psychophysics* 36, 1 (1984), 25–34.
- [8] CALLAGHAN, T. C. Interference and domination in texture segregation: Hue, geometric form, and line orientation. *Perception & Psychophysics* 46, 4 (1989), 299–311.
- [9] CALLAGHAN, T. C. Interference and dominance in texture segregation. In *Visual Search*, D. Brogan, Ed. Taylor & Francis, New York, New York, 1990, pp. 81–87.
- [10] CHERNOFF, H. The use of faces to represent points in k-dimensional space graphically. *Journal of the American Statistical Association* 68, 342 (1973), 361–367.
- [11] CIE. *CIE Publication No. 15, Supplement Number 2 (E-1.3.1): Official Recommendations on Uniform Color Spaces, Color-Difference Equations, and Metric Color Terms*. Commission Internationale de L'Éclairage, 1976.

- [12] COREN, S., AND HAKSTIAN, A. R. Color vision screening without the use of technical equipment: Scale development and cross-validation. *Perception & Psychophysics* 43 (1988), 115–120.
- [13] CUTTING, J. E., AND MILLARD, R. T. Three gradients and the perception of flat and curved surfaces. *Journal of Experimental Psychology: General* 113, 2 (1984), 198–216.
- [14] D'ZMURA, M. Color in visual search. *Vision Research* 31, 6 (1991), 951–966.
- [15] EGETH, H. E., AND YANTIS, S. Visual attention: Control, representation, and time course. *Annual Review of Psychology* 48 (1997), 269–297.
- [16] FOLEY, J., AND RIBARSKY, W. Next-generation data visualization tools. In *Scientific Visualization: Advances and Challenges*, L. Rosenblum, Ed. Academic Press, San Diego, California, 1994, pp. 103–127.
- [17] FOLEY, J. D., VAN DAM, A., FEINER, S. K., AND HUGHES, J. F. *Computer Graphics: Principles and Practice*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1990.
- [18] GLASSNER, A. S. *Principles of Digital Image Synthesis*. Morgan Kaufmann Publishers, Inc., San Francisco, California, 1995.
- [19] GRINSTEIN, G., PICKETT, R., AND WILLIAMS, M. EXVIS: An exploratory data visualization environment. In *Proceedings Graphics Interface '89* (London, Canada, 1989), pp. 254–261.
- [20] HALLETT, P. E. Segregation of mesh-derived textures evaluated by resistance to added disorder. *Vision Research* 32, 10 (1992), 1899–1911.
- [21] HARALICK, R. M., SHANMUGAM, K., AND DINSTEN, I. Textural features for image classification. *IEEE Transactions on System, Man, and Cybernetics SMC-3*, 6 (1973), 610–621.
- [22] HEALEY, C. G. Choosing effective colours for data visualization. In *Proceedings Visualization '96* (San Francisco, California, 1996), pp. 263–270.
- [23] HEALEY, C. G. Building a perceptual visualisation architecture. *Behaviour and Information Technology (in press)* (1998).
- [24] HEALEY, C. G., BOOTH, K. S., AND ENNS, J. T. Real-time multivariate data visualization using preattentive processing. *ACM Transactions on Modeling and Computer Simulation* 5, 3 (1995), 190–221.
- [25] HEALEY, C. G., AND ENNS, J. T. Building perceptual textures to visualize multidimensional datasets. In *Proceedings Visualization '98* (Research Triangle Park, North Carolina, 1998), pp. 111–118.
- [26] INTERRANTE, V. Illustrating surface shape in volume data via principle direction-driven 3d line integral convolution. In *SIGGRAPH 97 Conference Proceedings* (Los Angeles, California, 1997), T. Whitted, Ed., pp. 109–116.
- [27] JULÉSZ, B. Textons, the elements of texture perception, and their interactions. *Nature* 290 (1981), 91–97.
- [28] JULÉSZ, B. A theory of preattentive texture discrimination based on first-order statistics of textons. *Biological Cybernetics* 41 (1981), 131–138.
- [29] JULÉSZ, B. A brief outline of the texton theory of human vision. *Trends in Neuroscience* 7, 2 (1984), 41–45.
- [30] JULÉSZ, B., AND BERGEN, J. R. Textons, the fundamental elements in preattentive vision and perception of textures. *The Bell System Technical Journal* 62, 6 (1983), 1619–1645.
- [31] KAWAI, M., UCHIKAWA, K., AND UJIKE, H. Influence of color category on visual search. In *Annual Meeting of the Association for Research in Vision and Ophthalmology* (Fort Lauderdale, Florida, 1995), p. #2991.
- [32] LAIDLAW, D. H., AHRENS, E. T., KREMERS, D., AVALOS, M. J., JACOBS, R. E., AND READHEAD, C. Visualizing diffusion tensor images of the mouse spinal cord. In *Proceedings Visualization '98* (Research Triangle Park, North Carolina, 1998), pp. 127–134.
- [33] LEVKOWITZ, H. Color icons: Merging color and texture perception for integrated visualization of multiple parameters. In *Proceedings Visualization '91* (San Diego, California, 1991), pp. 164–170.
- [34] LEVKOWITZ, H., AND HERMAN, G. T. Color scales for image data. *IEEE Computer Graphics & Applications* 12, 1 (1992), 72–80.
- [35] LIU, F., AND PICARD, R. W. Periodicity, directionality, and randomness: Wold features for perceptual pattern recognition. In *Proceedings 12th International Conference on Pattern Recognition* (Jerusalem, Israel, 1994), pp. 1–5.
- [36] MACK, A., AND ROCK, I. *Inattentive Blindness*. MIT Press, Menlo Park, California, 1998.
- [37] MALIK, J., AND PERONA, P. Preattentive texture discrimination with early vision mechanisms. *Journal of the Optical Society of America A* 7, 5 (1990), 923–932.
- [38] MEIER, B. J. Painterly rendering for animation. In *SIGGRAPH 96 Conference Proceedings* (New Orleans, Louisiana, 1996), H. Rushmeier, Ed., pp. 477–484.
- [39] RAO, A. R., AND LOHSE, G. L. Identifying high level features of texture perception. *CVGIP: Graphics Models and Image Processing* 55, 3 (1993), 218–233.
- [40] RAO, A. R., AND LOHSE, G. L. Towards a texture naming system: Identifying relevant dimensions of texture. In *Proceedings Visualization '93* (San Jose, California, 1993), pp. 220–227.
- [41] REED, T. R., AND HANS DU BUF, J. M. A review of recent texture segmentation and feature extraction techniques. *CVGIP: Image Understanding* 57, 3 (1993), 359–372.
- [42] RENSINK, R. A., O'REGAN, J. K., AND CLARK, J. J. To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science* 8 (1997), 368–373.
- [43] RHEINGANS, P., AND TEBBS, B. A tool for dynamic explorations of color mappings. *Computer Graphics* 24, 2 (1990), 145–146.
- [44] ROBERTSON, P. K. Visualizing color gamuts: A user interface for the effective use of perceptual color spaces in data displays. *IEEE Computer Graphics & Applications* 8, 5 (1988), 50–64.
- [45] ROGOWITZ, B. E., AND TREINISH, L. A. An architecture for rule-based visualization. In *Proceedings Visualization '93* (San Jose, California, 1993), pp. 236–243.
- [46] SALISBURY, M., WONG, M. T., HUGHES, J. F., AND SALESIN, D. H. Orientable textures for image-based pen-and-ink illustration. In *SIGGRAPH 97 Conference Proceedings* (Los Angeles, California, 1997), T. Whitted, Ed., pp. 401–406.
- [47] SCHWEITZER, D. Artificial texturing: An aid to surface visualization. *Computer Graphics (SIGGRAPH 83 Conference Proceedings)* 17, 3 (1983), 23–29.
- [48] SIMON, D. J., AND LEVIN, D. T. Change blindness. *Trends in Cognitive Science* 1 (1997), 261–267.
- [49] SNOWDEN, R. J. Texture segregation and visual search: A comparison of the effects of random variations along irrelevant dimensions. *Journal of Experimental Psychology: Human Perception and Performance* 24, 5 (1998), 1354–1367.
- [50] TAMURA, H., MORI, S., AND YAMAWAKI, T. Textural features corresponding to visual perception. *IEEE Transactions on Systems, Man, and Cybernetics SMC-8*, 6 (1978), 460–473.
- [51] TRIESMAN, A. Preattentive processing in vision. *Computer Vision, Graphics and Image Processing* 31 (1985), 156–177.
- [52] TRIESMAN, A. Search, similarity, and integration of features between and within dimensions. *Journal of Experimental Psychology: Human Perception & Performance* 17, 3 (1991), 652–676.
- [53] TRIESMAN, A., AND GORMICAN, S. Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review* 95, 1 (1988), 15–48.
- [54] TURK, G., AND BANKS, D. Image-guided streamline placement. In *SIGGRAPH 96 Conference Proceedings* (New Orleans, Louisiana, 1996), H. Rushmeier, Ed., pp. 453–460.
- [55] WARE, C. Color sequences for univariate maps: Theory, experiments, and principles. *IEEE Computer Graphics & Applications* 8, 5 (1988), 41–49.
- [56] WARE, C., AND BEATTY, J. C. Using colour dimensions to display data dimensions. *Human Factors* 30, 2 (1988), 127–142.
- [57] WARE, C., AND KNIGHT, W. Using visual texture for information display. *ACM Transactions on Graphics* 14, 1 (1995), 3–20.
- [58] WOLFE, J. M. Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review* 1, 2 (1994), 202–238.
- [59] WOLFE, J. M., YU, K. P., STEWART, M. I., SHORTER, A. D., FRIEDMAN-HILL, S. R., AND CAVE, K. R. Limitations on the parallel guidance of visual search: Color  $\times$  color and orientation  $\times$  orientation conjunctions. *Journal of Experimental Psychology: Human Perception & Performance* 16, 4 (1990), 879–892.
- [60] WYSZECKI, G., AND STILES, W. S. *Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd Edition*. John Wiley & Sons, Inc., New York, New York, 1982.



**Christopher G. Healey** received a PhD in computer science in 1996 from the University of British Columbia, an MSc in 1992 from the University of British Columbia, and a BMath in 1990 from the University of Waterloo. Following graduation he completed a two-year postdoctoral fellowship in computer graphics with Dr. Carlo Séquin at the University of California, Berkeley. He is currently working as an assistant professor in the Department of Computer Science at North Carolina State University.

His dissertation studied methods for displaying effectively large, multivariate datasets during scientific visualization. This work investigated techniques for exploiting the low-level human visual system for information representation. His current research focuses on the use of visual features like color, texture, and apparent motion for visually exploring multivariate data. He is also investigating automated data-feature mapping techniques and data management issues in an effort to design a flexible, robust perceptual visualization architecture.



**James T. Enns** received a PhD in psychology from Princeton University (1984) and a BA (honours) from the University of Winnipeg (1980). Following graduation, he was appointed an assistant professor at Dalhousie University, before moving to the University of British Columbia in 1987, where he is now appointed as a professor in both the Department of Psychology and the Graduate Program in Neuroscience. A central focus of his research over the years has been the role of attention in

perception. This has included studies of how perception and attention change with development, how the visual world is represented outside the focus of attention, and how attention changes the perceptions that form the basis for consciousness. Along with the publication of these studies in *Science*, *Psychological Review*, *Perception & Psychophysics*, and *The Journal of Experimental Psychology*, he has edited two research volumes (*The Development of Attention*, 1990; *Attention, Development, & Psychopathology*, 1997) and coauthored two textbooks (*Analysis of Variance*, 1986; *Sensation & Perception*, fifth edition, 1999).

# 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 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.

---

This work was supported in part by the National Science Foundation (NSF) grant CISE-ACI-0083421, and by the National 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; lgta-teos@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/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2004 ACM 0730-0301/2004/0100-0001 \$5.00

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 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

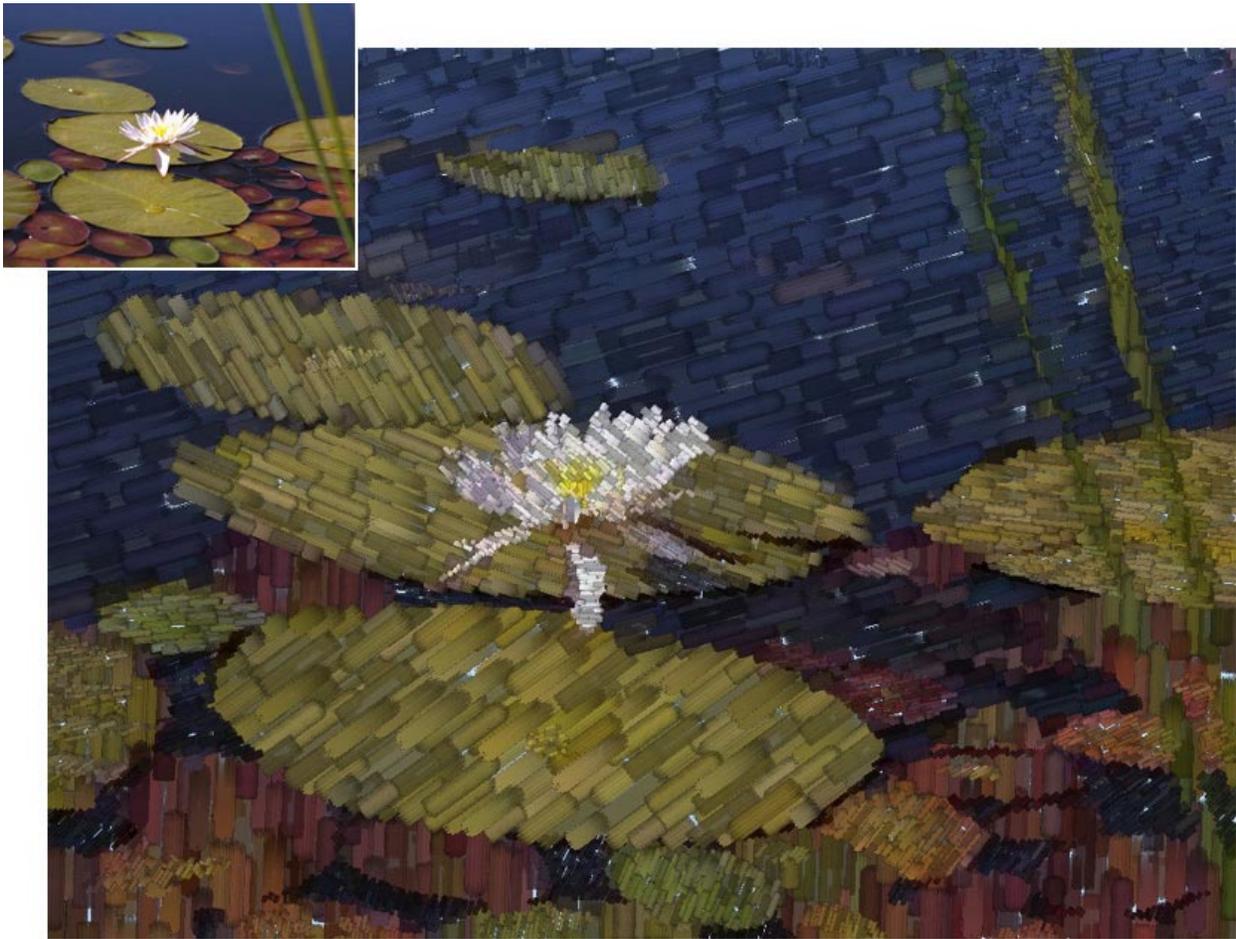


Fig. 1. A simple nonphotorealistic rendering of a collection of water lilies; the original image is shown in the upper-left corner

[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 non-photorealistic 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.
- (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.

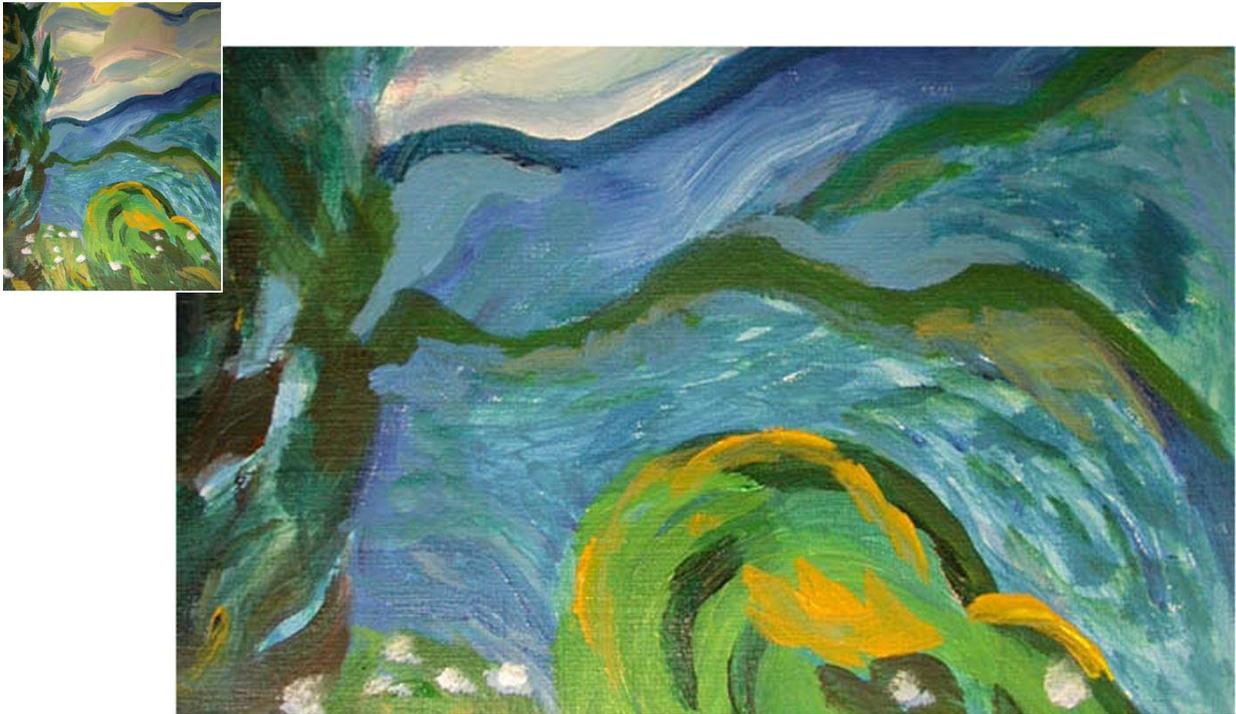


Fig. 2. A close-up view of a small section of an oil painting that demonstrates various stroke properties such as color, path, size, and density; the entire painting is shown in the upper-left corner

- (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 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



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)

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.

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



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)

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 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

<sup>1</sup>Hint: Look at the bushes immediately behind the back of the Sphinx

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 [1988] constructed a color scale that spirals up around the luminance axis to maintain a uniform simultaneous contrast error along its length. 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 CIELUV 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; Julész 1975; 1984; Julész 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 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  $15^\circ$  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.

## 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

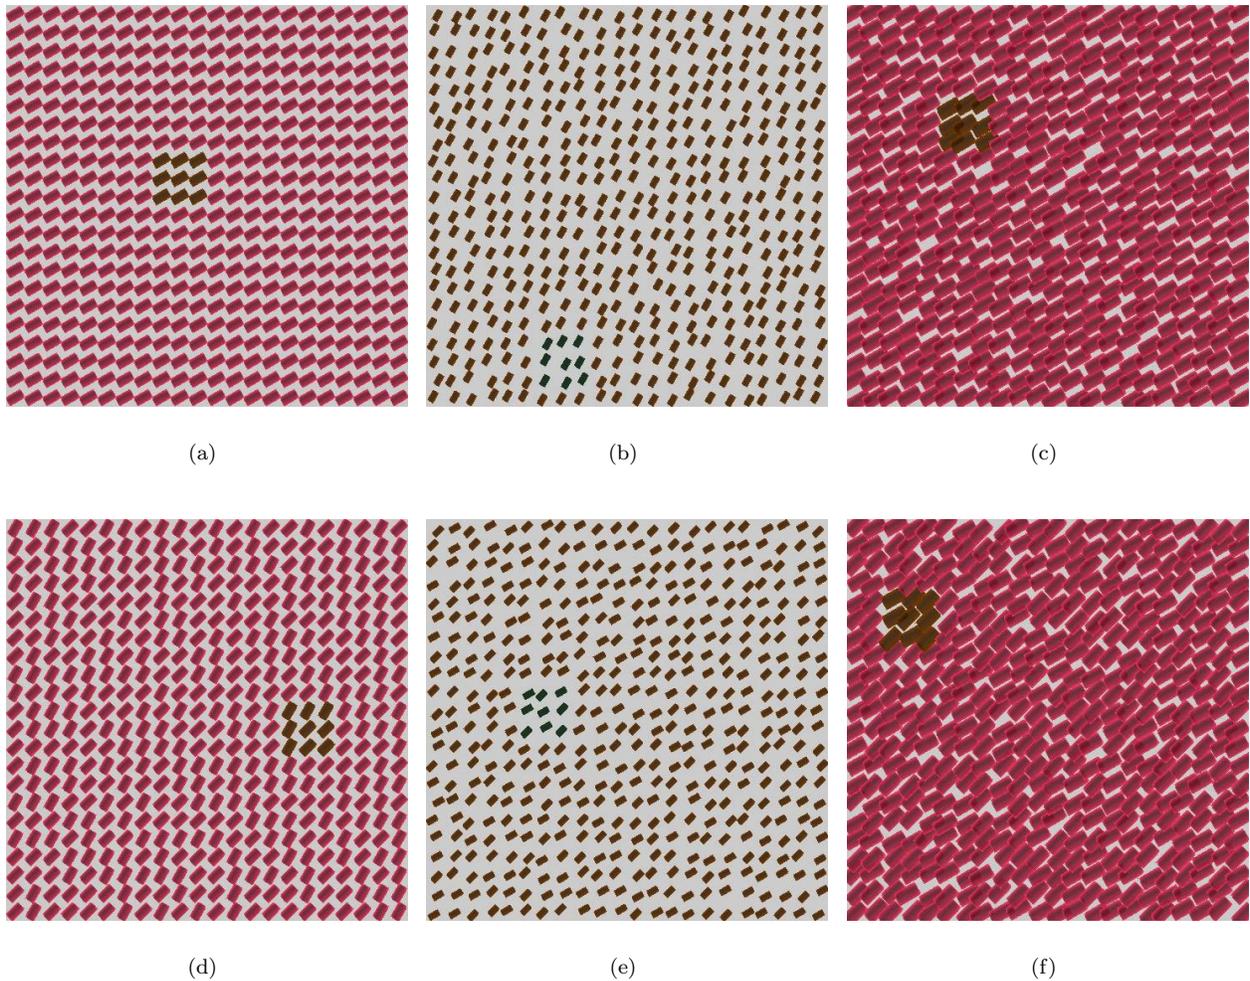


Fig. 5. Examples of target detection, color targets with constant orientation (top) and random orientation (bottom); (a) orange target in pink strokes, constant 45° background orientation; (b) green target in orange strokes, constant 60° background orientation; (c) green target in orange strokes, constant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

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

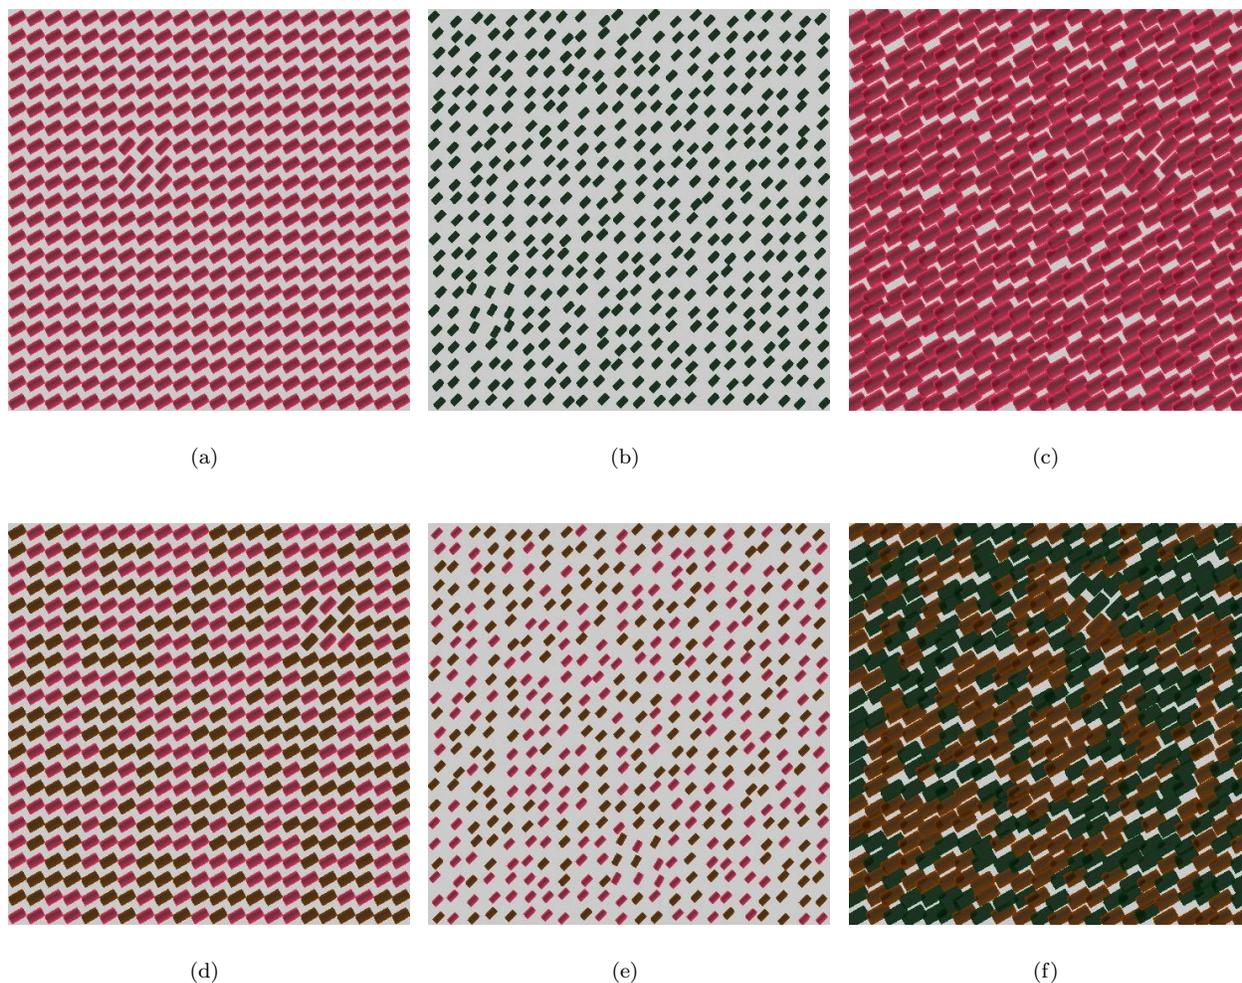


Fig. 6. Examples of target detection, orientation targets with constant color (top) and random color (bottom); (a)  $45^\circ$  target in  $30^\circ$  strokes, constant pink background color; (b)  $60^\circ$  target in  $45^\circ$  strokes, constant green background color; (c)  $45^\circ$  target in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

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,  $3 \times 3$  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 of strokes tilted  $60^\circ$  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  $22 \times 22$  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  $30^\circ$  or  $60^\circ$ ), or random (strokes are randomly oriented  $30^\circ$  and  $45^\circ$ , or  $45^\circ$  and  $60^\circ$ ); 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 5(a), 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, who found a general color on texture interference

pattern, but no corresponding texture on color effect [Healey and Enns 1999]. 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  $\varepsilon_j$  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.

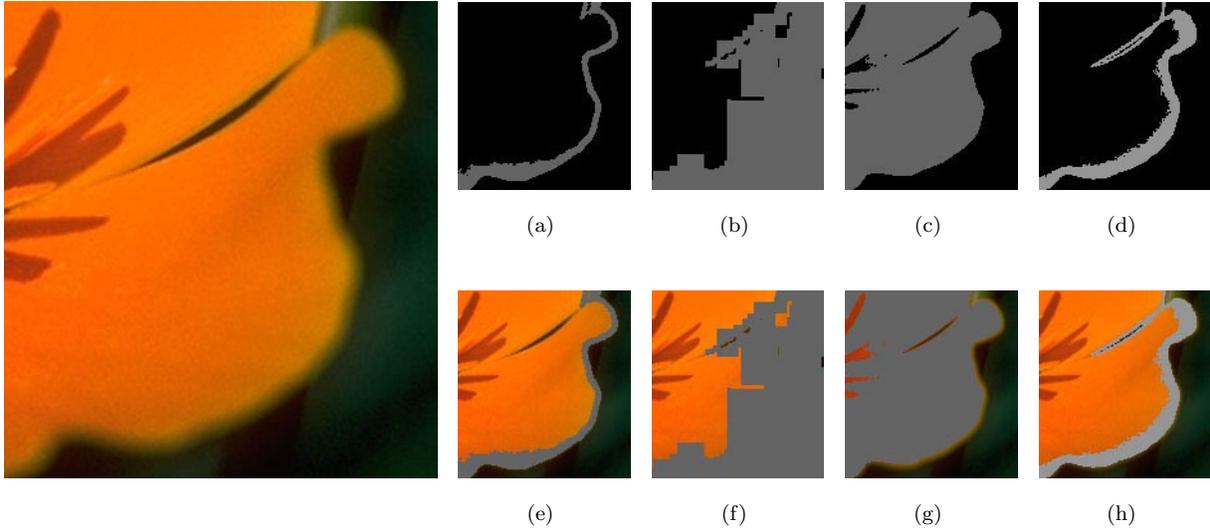


Fig. 7. Examples of different segmentation algorithms applied to an RGB image of a golden poppy; (a, b, c, d) segments with a fixed, running average, weighted average with  $w = 1$ , and weighted average with  $w = \frac{7}{8}$ , respectively; (e, f, g, h) segments overlaid on the original RGB image

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.

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:



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

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

$w$  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 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 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  $M$ :

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

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

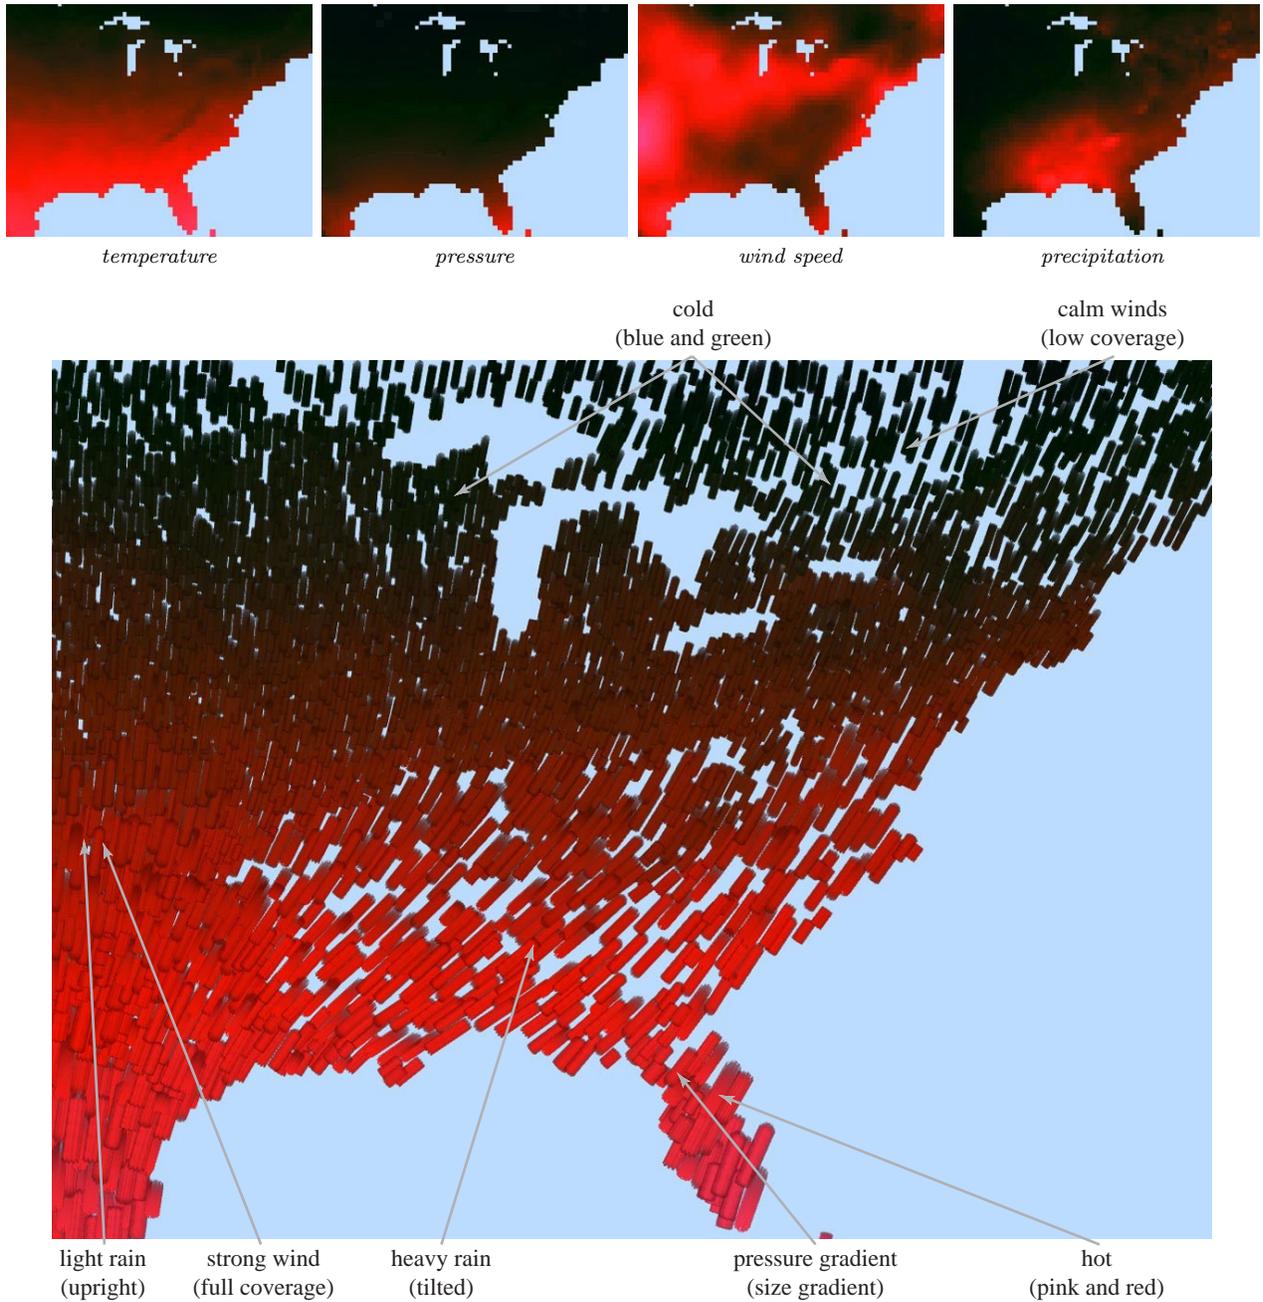
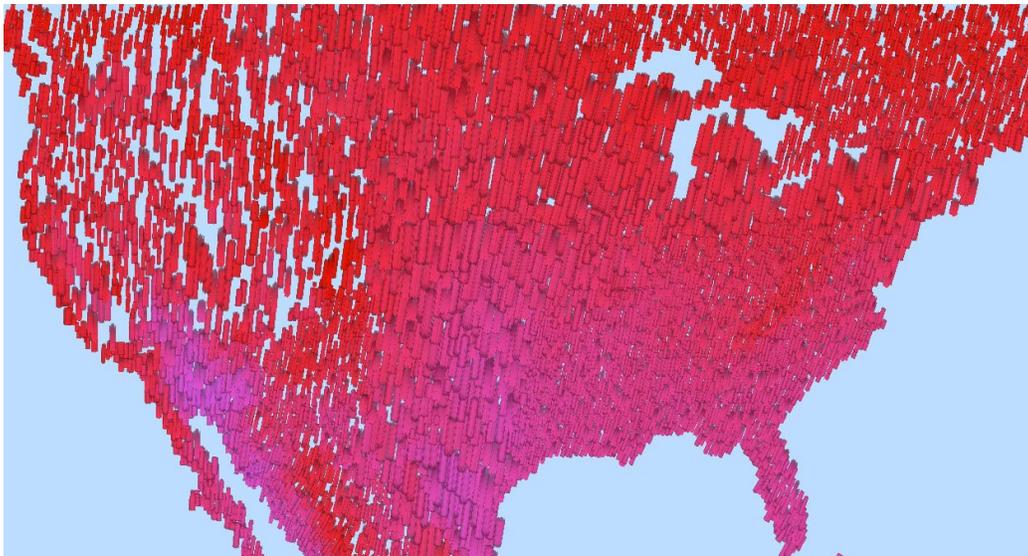


Fig. 9. Nonphotorealistic visualization of weather conditions for February over the eastern United States: (top row) perceptual color ramps (dark blue for low to bright pink for high) of mean *temperature*, *pressure*, *wind speed*, and *precipitation* in isolation; (bottom row) combined visualization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



(a)



(b)

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

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 a 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.
- (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  $A = (\textit{temperature}, \textit{wind speed}, \textit{wind direction}, \textit{precipitation})$ . 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  $V_s = (\textit{color}, \textit{luminance}, \textit{directed contours}, \textit{semi-transparent color})$ ,  $\Phi_s = (\textit{green} \cdots \textit{yellow}, \textit{dark} \cdots \textit{bright}, 0^\circ \cdots 360^\circ, \textit{green} \cdots \textit{red})$ . Certain modifications were needed to combine the data into a single image. For example, *temperature*, *wind speed*, and *precipitation*

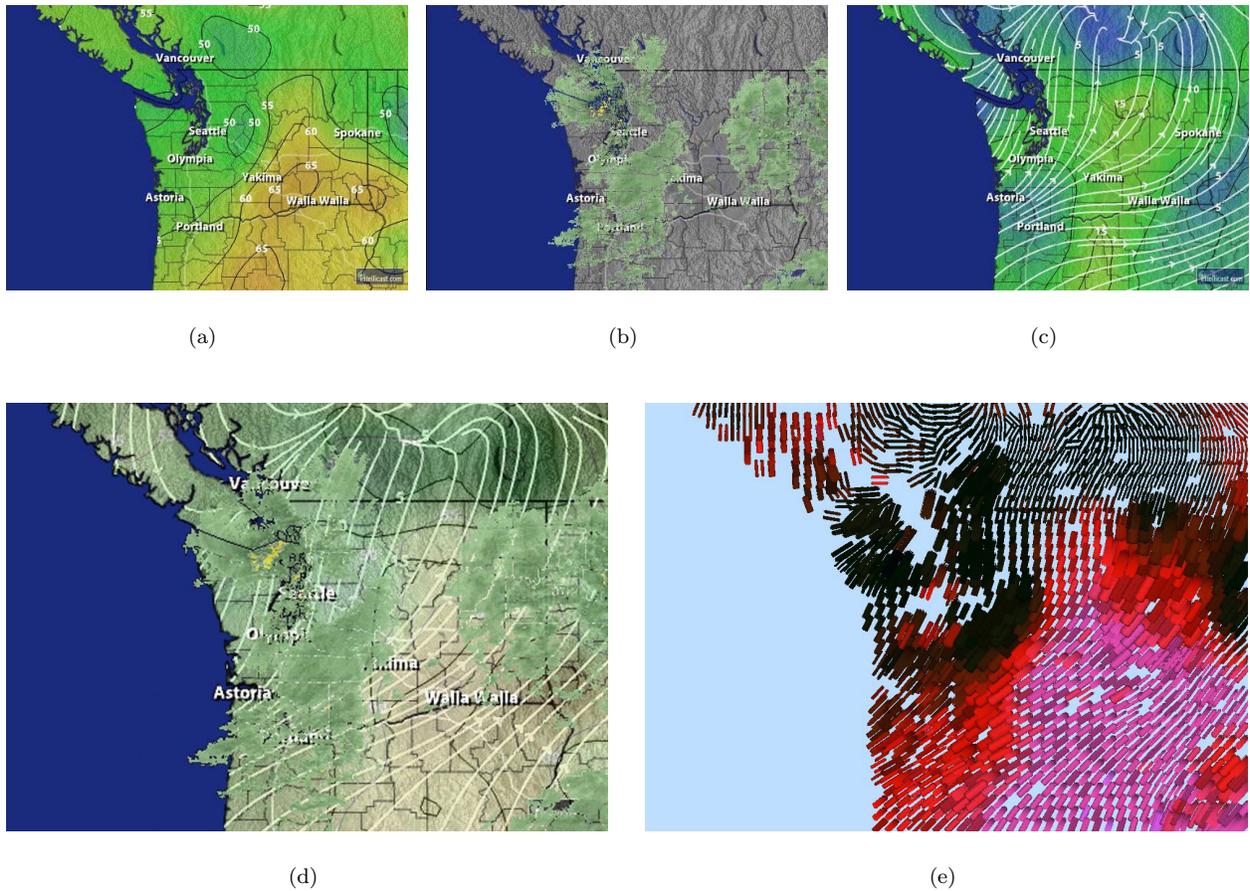


Fig. 11. Example multidimensional visualizations: (a) standard visualization of *temperature* with color (dark green for cooler to yellow for warmer); (b) standard visualization of *precipitation* with Doppler radar (green for light rainfall to red for heavy rainfall); (c) standard visualization of *wind 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

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  $V_n = (\text{color, coverage, orientation, size})$ ,  $\Phi_n = (\text{dark blue} \cdots \text{bright pink, low} \cdots \text{high, } 0^\circ \cdots 360^\circ, \text{small} \cdots \text{large})$ . 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:

- (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 7.1 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 7.1 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

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

Visualization	<i>temperature</i>	<i>precipitation</i>	<i>wind speed</i>	<i>wind direction</i>
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

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 7.1), this did not produce any improvement in identifying regions of high *precipitation* ( $\chi_2^2 = 0.37$ ,  $p < 0.90$ ).

Table 7.1 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 semi-transparent 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 7.1 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 ( $\chi_1^2 = 8.572$ ,  $p < 0.01$ ).

## 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.

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

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 7.1 and 7.1 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., semitransparent 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.<sup>2</sup> 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*,

<sup>2</sup>In 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).

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 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. Experiment. 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 colours 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 the 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. *Comput. Graph. (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. Experiment. Psych.: Human Percept. Perf.* 15, 4, 635–649.
- POSNER, M. I. AND RAICHEL, 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. of 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. ACM, New York, 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. Perf.* 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 SCHLECHTWEG, 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. Sys., 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. Perf.* 26, 2, 693–716.
- ZEKI, S. 1999. *Inner Vision*. Oxford University Press, Oxford, U. K.

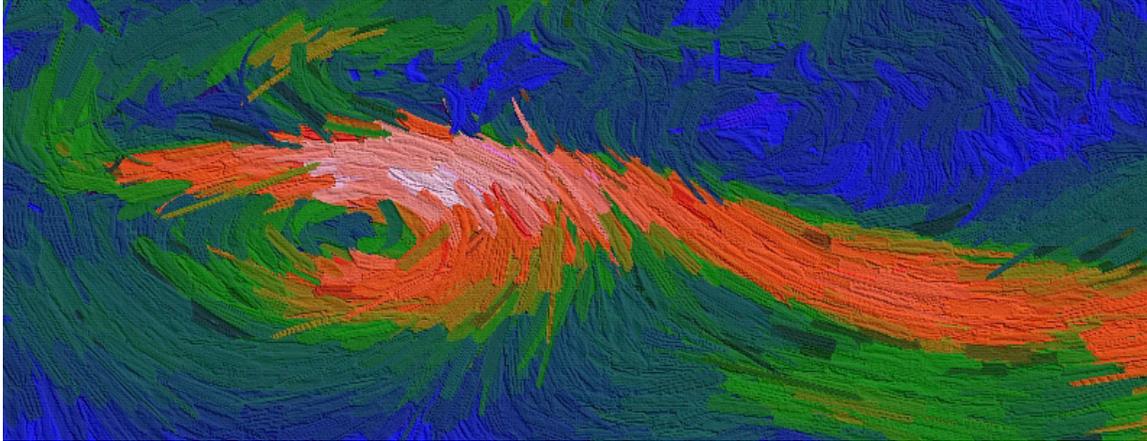
Received January 2003; revised May 2003; accepted August 2003

# Engaging Viewers Through Nonphotorealistic Visualizations

Laura G. Tateosian\*  
Department of Computer Science  
North Carolina State University

Christopher G. Healey†  
Department of Computer Science  
North Carolina State University

James T. Enns  
Department of Psychology  
University of British Columbia



**Figure 1:** A visual complexity style visualization of flow patterns in a 2D slice through a simulated supernova collapse, using the mappings: flow orientation  $\rightarrow$  stroke orientation, magnitude  $\rightarrow$  color (dark blue to bright pink for low to high), and pressure  $\rightarrow$  stroke size.

## Abstract

Research in human visual cognition suggests that beautiful images can engage the visual system, encouraging it to linger in certain locations in an image and absorb subtle details. By developing aesthetically pleasing visualizations of data, we aim to engage viewers and promote prolonged inspection, which can lead to new discoveries within the data. We present three new visualization techniques that apply painterly rendering styles to vary interpretational complexity (IC), indication and detail (ID), and visual complexity (VC), image properties that are important to aesthetics. Knowledge of human visual perception and psychophysical models of aesthetics provide the theoretical basis for our designs. Computational geometry and nonphotorealistic algorithms are used to preprocess the data and render the visualizations. We demonstrate the techniques with visualizations of real weather and supernova data.

**Keywords:** nonphotorealistic rendering, visualization, aesthetics, mesh simplification, Voronoi diagrams, NPR applications

## 1 Introduction

Visualizations enable scientists to inspect, interpret, and analyze large multi-dimensional data sets. We believe effective visualizations should both orient and engage (attract *and* hold) viewer attention, directing the viewer’s gaze in response to a visual stimulus,

\*e-mail: lauratateosian@yahoo.com

†e-mail: healey@csc.ncsu.edu

and then encouraging it to linger at a given image location. Research on human visual perception describes how to attract attention, using salient visual features, such as color and texture. Less is known about how to engage viewers with visualizations.

Recently, researchers have applied techniques from nonphotorealistic rendering in computer graphics to enhance and highlight important details in a visualization. We believe that nonphotorealism can also be used to increase the aesthetic merit of a visualization, which in turn may increase its ability to engage a viewer. To this end, we present three new visualization techniques designed with the goal of both orienting and engaging viewers. Orienting is controlled by painterly brush stroke glyphs whose visual properties vary to represent attributes of the underlying data. Human perception studies showed that the information-carrying capabilities of this painterly technique are consistent with more traditional glyph-based visualizations [Healey et al. 2004; Tateosian 2002]. To engage viewers, we incorporate visual qualities important to aesthetics into our visualizations. These visual qualities, interpretational complexity, indication and detail, and visual complexity, were chosen based on psychophysical models of aesthetics. Experiments conducted using these models provide strong evidence that varying image complexity can significantly impact a viewer’s aesthetic judgment.

We apply our visualization techniques to real meteorological and supernova data to explore their capabilities in a real-world setting. Anecdotal feedback from domain experts is strongly positive, supporting the hypothesis that enhancing the artistic merit of a visualization can result in more effective and more productive visual analysis.

The remainder of this paper describes our visualization techniques. Section 2 discusses related work in artistic visualization, nonphotorealistic rendering, and aesthetics. Section 3 introduces the visual qualities that form the basis of our visualizations. Sections 4 and 5 describe the algorithms that were used to preprocess and render the data. Section 6 shows examples of real world applications visualized with our techniques. Finally, Section 7 discusses future work.

## 2 Related Work

We reviewed previous research on aesthetics from the visualization, computer graphics, and human perception communities, then extended and combined results from each domain during the design of our aesthetic visualization algorithms.

### 2.1 Nonphotorealistic Visualization

Scientists in visualization are using ideas from NPR to inspire new visualization techniques. Artistic techniques, such as using abstraction to eliminate unimportant distractions, and sharpening details to draw attention to important areas, can help to convey information more effectively. [Laidlaw et al. 1998; Kirby et al. 1999; Healey et al. 2004; Tateosian 2002] use painting concepts to visualize data. [Laidlaw et al. 1998] visualizes diseased and healthy mouse spinal cords using a tonal underpainting for the anatomy of the brain matter and textured elliptical brush strokes for diffusion tensor data. [Kirby et al. 1999] visualizes air flow past a propeller with an underpainting, ranging from blue to yellow for clockwise to counter clockwise vorticity, and elliptical and arrow-shaped strokes representing additional flow information. The mapping is chosen to display velocity and vorticity prominently, while still showing deformation of the fluid elements. [Healey et al. 2004; Tateosian 2002] create painterly visualizations by mapping data attributes to the visual features of rectangular brush stroke glyphs. In these techniques, brush strokes are “painted” atop an underpainting and stroke features such as shape, color, transparency, orientation, and texture display local data attributes.

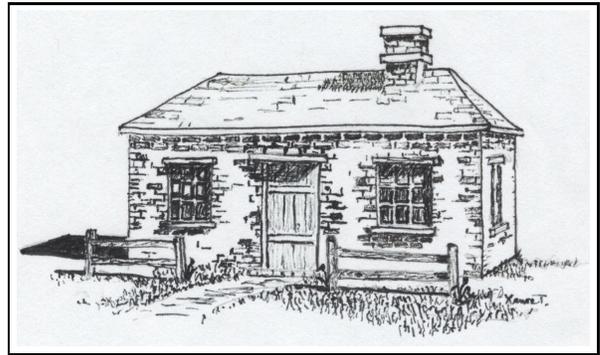
[Grigoryan and Rheingans 2002; Lu et al. 2002] stipple-render volumes. [Grigoryan and Rheingans 2002] uses multi-colored stippling, to show uncertainty on the surface of growing tumors as an intuitively uncertain, fuzzy surface of sparsely placed points. [Lu et al. 2002] uses pen-and-ink stippling, placing and sizing stipples to control both local and global tone. [Rheingans and Ebert 2001; Stoppel et al. 2002] use artistic techniques like silhouetting and outline sketches to enhance volume visualizations.

These visualization methods use artistic techniques to enhance their expressiveness, as do our new algorithms. One important contribution of our visualizations is a basis in psychological literature to vary complexity in ways specifically designed to improve a visualization’s aesthetic merit.

### 2.2 Nonphotorealistic Rendering Techniques

Existing work in NPR provides a rich vocabulary of artistic expression in computer graphics. We use several of the NPR concepts and algorithms described here. Sophisticated pen-and-ink illustration and stippling simulation systems have been developed [Winkenbach and Salesin 1994; Salisbury et al. 1994; Winkenbach and Salesin 1996; Salisbury et al. 1997; Secord 2002; Wilson and Ma 2004]. In our work, we make use of a classic pen-and-ink technique, called, “indication and detail,” that draws more detail in areas of interest and reduces detail in homogeneous regions [Winkenbach and Salesin 1994]. In [Winkenbach and Salesin 1994], users draw line segments on an image where high detail is desired. The level of detail is reduced as the distance from the line segments increases. The hand-drawn house in Fig. 2 shows results this system would produce if detail segments were drawn around the door and front windows.

Painting simulation techniques are of particular interest because of our brush-stroke based visualization style. [Lewis 1984; Strassmann 1986; Pham 1991; Hsu and Lee 1994] present approaches for modeling sophisticated individual strokes. Several systems use



**Figure 2:** A hand-drawn recreation of an image from Winkenbach and Salesin’s paper on computer-generated pen-and-ink drawings demonstrating the use of “indication and detail” [Winkenbach and Salesin 1994]. This technique avoids monotonous repetition by presenting a small amount of detail to indicate patterns that continue in the surrounding regions.

lists of brush strokes to create Impressionist style imagery [Haeberli 1990; Meier 1996; Litwinowicz 1997; Hertzmann 1998; Shiraishi and Yamaguchi 2000; Hertzmann 2002]. [Meier 1996] creates animated paintings and [Litwinowicz 1997] generates painterly videos. [Hertzmann 1998; Shiraishi and Yamaguchi 2000] simulate paintings of source photographs. [Hertzmann 1998] layers curved B-splined strokes, increasing detail on each layer. [Shiraishi and Yamaguchi 2000] use image moment functions to size and orient rectangular strokes. [Hertzmann 2002] extends [Hertzmann 1998] with a fast paint texture method to simulate the thickness of paint using a height field and lighting, a technique we implement to create textured visualizations.

### 2.3 Characterization of Beauty

To choose visual properties to vary in our visualizations, we reviewed models of aesthetics that identify parameters that are believed to affect aesthetic judgment. In an early attempt to model beauty with mathematical equations, Birkhoff calculated the aesthetic measure,  $M$ , of objects, such as polygons, tiles, and vases as a ratio of order to complexity [Birkhoff 1932]. Birkhoff theorized that heightened complexity and disorder induce feelings of discomfort and so aesthetic measure decreases as the ratio of order to complexity decreases. Observer rankings of Birkhoff’s polygons in a study conducted by Davis, suggest that aesthetic value is curvilinearly related to  $M$ , peaking when  $M$  is moderate [Davis 1936].

The same pattern is seen in Berlyne’s model of aesthetic pleasure as a function of arousal, which plots as an inverted U [Berlyne 1971]. Viewing beautiful images generally evokes a feeling of pleasure. The image provides an activating stimulus and the arousal induces pleasure. Thus, artistic value is identified with pleasure. In this model, as arousal increases, pleasure increases until it peaks and decreases toward indifference (zero pleasure) and displeasure (negative pleasure values). Structural properties, such as complexity, novelty, conflict, expectations, ambiguity, and instability increase arousal; whereas, familiarity, dominance, grouping and pattern, and expectations moderate arousal.

Arousal and pleasure axes also comprise Barrett and Russell’s affective model, a tool for specifying a wide range of emotional states. This does not imply that aesthetic judgment is solely an emotional process. Factor analysis on how paintings are judged revealed that a cognitive factor, made up of complexity, meaningfulness, interestingness, pleasantness, and familiarity, explained a

majority (51.4%) of the variance in judgment of paintings and an emotional factor, made up of warmth, emotionality, arousal, and dominance, explained 13.7% of the variance [Baltissen and Ostermann 1998]. The converse was true for emotional pictures (48.1% of the variance explained by the emotional scale and 14% explained by the cognitive scale). Viewing paintings may be largely a cognitive process, because it requires the viewer to identify the meaning, and emotions may arise merely as an after-effect associated with the cognitive process.

### 3 Designing Aesthetic Visualizations

Birkhoff, Berlyne, and Barrett and Russell reveal that complexity and closely related properties such as dominance play an important role in aesthetic judgment. Our three nonphotorealistic visualization techniques are designed to vary these properties.

**Interpretational Complexity (IC).** IC uses a layered approach to introduce complexity. Paintings are often developed in layers with an undercoating broadly defining the shapes in the image and details refined with subsequent layers. This aspect of paintings adds interpretational complexity, because the information provided by additional layers requires interpretation. Unlike the layering in [Tateosian 2002; Hertzmann 1998; Shiraishi and Yamaguchi 2000; Hertzmann 2002], we wanted our underpainting to contrast in display style and level of detail from the top layer, because extracting contrast reinforces attention [Ramachandran 2000]. Hence the underpainting, a colored canvas with sparse faint strokes, provides a lower level of detail than the highly detailed top level painted with wet strokes. The data to be visualized is filtered to detect regions of rapid change, as described in Section 4. Then distinct strokes are laid on a second layer to provide detail in rapid change areas, but not elsewhere so that the underpainting is not fully covered.

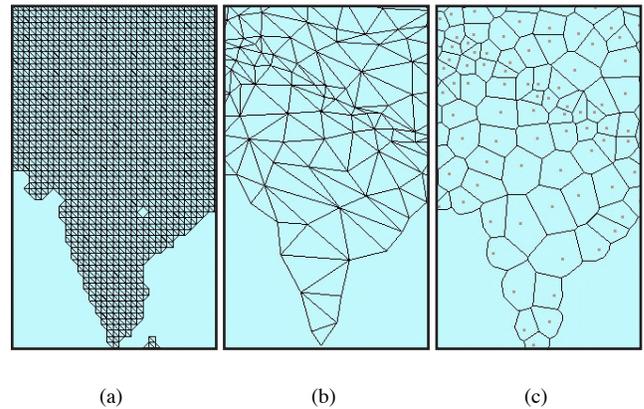
**Indication and Detail (ID).** ID provides focal points with high detail areas and omits monotonous details in homogeneous regions (Fig. 2). Tracking the eye movements of people viewing paintings has shown that viewers fixate on areas of high detail, such as faces [Wooding 2002; Holman et al. 2004]. The abstraction of detail in some areas allows other areas to be visually dominant, resolving the conflict created by competing stimuli. Short-term familiarity develops when a visual element appears, disappears, and reappears in an image after a short hiatus. Overly repetitive visual displays may be too simplistic to be engaging. But *variation*, repeating patterns that are alike in one respect and varying in another, is a powerful aesthetic device [Berlyne 1971]. We can reuse the filtering technique we developed for IC to preprocess the data. Then highly variable regions are drawn with high detail and clearly distinguishable brush strokes, and homogeneous areas are drawn with low detail and subtle strokes. Transitional strokes, something like the partially drawn bricks in Fig. 2, are drawn on the borders between low and high detail regions. In ID, stroke textures are outlined with a marker to create a congruous style, so that the detailed areas gradually evolve into abstracted areas, unlike IC where the layers are designed to contrast in style.

**Visual Complexity (VC).** In two patterns with the same number of visual elements, the one with the most similarity among its elements will be considered less visually complex [Berlyne 1971]. Inspecting master Impressionist artist paintings reveals a great deal of variation amongst the individual brush strokes. There are often differences in the color and thickness of paint even within a single stroke. These visual properties contribute to the complexity of the painting, but in a way that differs from interpretational complexity. VC involves local visual variations, where IC creates global trends. Psychologists consistently cite complexity as one of the components affecting the judgment of aesthetics. Thus, we decided

to study visual complexity, as well as interpretational complexity.

To inject visual complexity, we introduce new stroke properties that correspond to observable characteristics of brush strokes in master Impressionist paintings. To identify these properties, we consulted an artist who has a background in Fine Arts and is a practicing oil and watercolor painter. We told her we were seeking salient and aesthetically pleasing ways to vary our basic brush stroke style and asked her to look at several Van Gogh reproductions. We mapped her observations of the kinds of variations she saw in these paintings to six properties: varied stroke contour, variegation (paint color varying within one stroke), curved strokes, embossed strokes (paint thickness varying within one stroke, sometimes dramatically), varied aspect ratio (some strokes are much longer or fatter than others), and contrast strokes (sets of dark and light strokes moving in different directions, juxtaposed in places).

### 4 Data Preprocessing



**Figure 3:** Processing Indian weather data: a) Triangulate the data element locations. b) Simplify the mesh based on data attribute values. c) Find the corresponding Voronoi diagram.

The IC and ID techniques first identify areas of high spatial frequency within the data to be visualized. To do this we apply mesh simplification followed by spatial analysis with Voronoi regions.

Feature preserving mesh simplification reduces a mesh by removing polygon vertices in ways that minimize the surface error. For example, a simplified terrain map mesh can describe a plateau with just a few large polygons, whereas, many more polygons remain for craggy mountains. Mesh simplification algorithms that handle surface properties, such as RGB triples, surface normals, and texture coordinates, can be used to process a multidimensional data set, with each data element’s attribute values representing surface properties at the corresponding vertex [Walter and Healey 2001]. After simplification, the plateaus of our example, are characterized by a sparse distribution of data elements, while in regions of rapid attribute value change (craggy mountains) the data elements remain dense. We used a popular implementation called qslim, which applies iterative vertex contraction with a quadric error metric to simplify meshes with associated surface properties [Garland and Heckbert 1997; Garland and Heckbert 1998; Garland 1999].

To use qslim, we triangulated our data sets, set the surface properties to data attributes, and chose a desired face count. Fig. 3a shows triangulated data elements from an Indian weather data set. The output face count is reduced to 3% of the original face count (Fig. 3b). The data attributes, *radiation*, *wet day frequency*, and *diurnal temperature range* were set as surface properties, and used to determine the quadric error.

To measure the denseness of the mesh vertices, we use a nearest point Voronoi diagram, which is a partitioning of space around a set,  $S$ , of  $n$  sites,  $S = p_1, \dots, p_n$ , into  $n$  regions, such that each point inside the region containing  $p_i$  is closer to  $p_i$  than any of the other  $n - 1$  sites. Since the size of the Voronoi region decreases as density increases, the Voronoi diagram of our reduced mesh identifies regions of rapid change. We used an available implementation, qhull, to find our Voronoi diagrams with the results from qslim acting as input sites to qhull [Barber et al. 1993; Barber et al. 1996; Brown 1979]. Using our simplified Indian weather example in Fig. 3b as input, the resulting Voronoi diagram is shown in Fig. 3c. The region sizes are correlated with the density of the input mesh. For example, because the input attributes, *radiation*, *wet day frequency*, and *diurnal temperature range*, vary rapidly in the Himalayas, the Voronoi regions are much smaller in these mountainous areas.

## 5 Visualization Algorithms

Our three visualization techniques all create painted regions of strokes that vary in color, size, and orientation based on underlying data values. IC, ID, and VC share a common painting algorithm to lay strokes in a random manner while controlling coverage. As region  $S_k$  is painted, coverage is controlled by tracking the percentage,  $c$ , of coverage of the total area of  $S_k$ . A stroke is centered at a randomly chosen unpainted position within  $S_k$ . It is transformed and scan converted, and then retained or rejected based on whether there is too much overlap (when too much of a stroke falls outside  $S_k$  or overlaps previously laid strokes). Stroke properties such as size, orientation, and color, are determined by a data sample at the center of the stroke. Each sample is a vector of data attribute values at that position in the underlying data set, which are mapped to a vector of visual features. Stroke placement continues until the region’s desired coverage,  $c$ , has been met. The output is a list of the accepted strokes and their properties. Each stroke in the list is drawn as a textured quad. Stroke textures were mostly generated from scanned hand-painted strokes or strokes drawn with marker. The algorithm for creating a stroke list with coverage  $V_k$  of region  $S_k$  is:

```

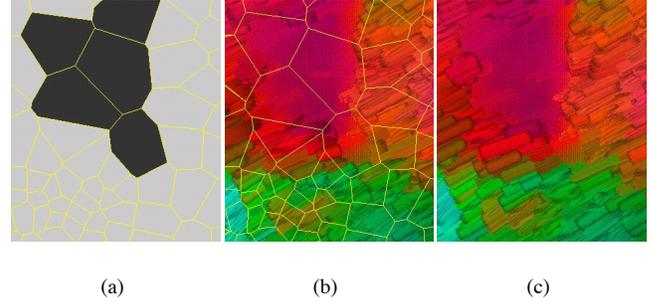
while  $c < V_k$  do
  Randomly pick an unpainted position  $p$  within  $S_k$ 
  Identify data element  $e_i$  associated with  $p$ 
  Set size and orientation of a new stroke,  $s$ , based on  $e_i$ 
  Center  $s$  at  $p$  and scan convert
  Compute amount of  $s$  outside of  $S_k$ , outside
  Compute amount of  $s$  overlapping existing strokes, overlap
  if outside or overlap are too large then
    Shrink  $s$  until it fits or it cannot be shrunk
  end if
  if  $s$  fits then
    Add to stroke list and update  $c$ 
  end if
end while

```

### 5.1 The Interpretational Complexity Algorithm

For IC, we wanted to build two layers, each with a distinct style. Master artists often use an underpainting to broadly define forms. The underpainting, varying in color, sometimes peeks through to higher detail layers. The individual strokes in the underpainting are difficult to discern. These patches of underpainting appear where there is little change or detail in the image.

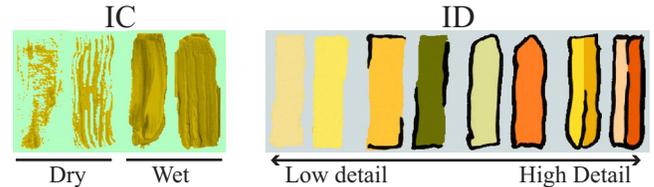
The Voronoi regions determine where the underpainting should show by partitioning the data into two sets:  $S_0$ , representing low spatial frequency changes in the data (i.e. large Voronoi regions), and  $S_1$ , representing high spatial frequency changes (i.e. small



**Figure 4:** The IC algorithm: a) Partition Voronoi regions into sets  $S_0$  (dark) and  $S_1$  (light) based on size. b) Paint layer  $S_0$ , then  $S_1$  (Voronoi regions overlaid). c) Render the results.

Voronoi regions). For each region, if its area exceeds a *threshold* percentage of the largest Voronoi region, it is placed in  $S_0$ , otherwise it is placed in  $S_1$  (Fig. 4a).

The Voronoi regions in  $S_0$  are tiled with colored canvas texture to establish the broad color changes of the underpainting. The tiles are smoothly shaded, with the nearest data element’s attribute values determining the color of each tile corner. A sparse collection of strokes are placed on the tiles by centering a stroke on each Voronoi vertex, and in the center of each Voronoi region (Figs. 4b and c). The strokes are textured with a dry brush appearance (Fig. 5), and seem faint since some of the underlying canvas shows through.



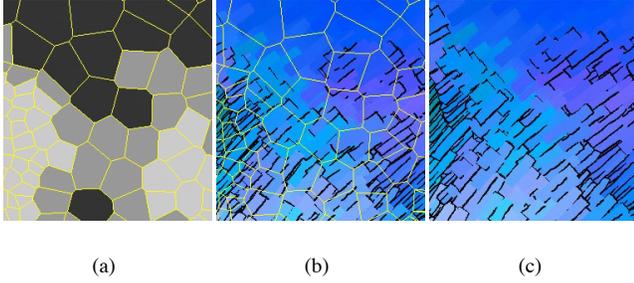
**Figure 5:** Sample stroke textures: left) dry and wet brush strokes for the IC visualizations. right) strokes with increasing detail for the ID visualizations.

Finally, Voronoi regions in  $S_1$  are painted with our painting algorithm set to produce full coverage. Each stroke is textured in a wet brush fashion (Fig. 5). This produces the necessary contrast with the dry brush strokes in regions from  $S_0$ . Fig. 4 demonstrates the IC process.

### 5.2 The Indication and Detail Algorithm

To create our stylized ID visualizations, we use the Voronoi regions to partition the data into three sets  $S_0$ ,  $S_1$ , and  $S_2$ , then paint with textured strokes like those in Fig. 5. We begin by identifying all Voronoi regions whose area  $A \geq \text{threshold}$ , and placing them in  $S_0$  (large Voronoi regions). For the remaining regions if the region has at least one neighbor in  $S_0$ , we place it in  $S_1$  (small Voronoi regions bordering large ones), otherwise we place it in  $S_2$  (small Voronoi regions).

Voronoi regions in  $S_0$  are drawn as smooth shaded polygons, with vertex colors based on the data element at their corresponding spatial position. The regions are then painted with one stroke at their center and one stroke at each vertex, rendered with low detail texture maps that have no outline (Fig. 5).



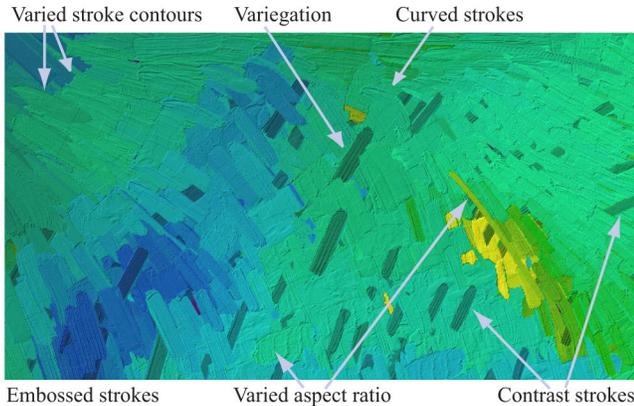
**Figure 6:** Generating increase in detail for the ID algorithm: a) Partition Voronoi regions into sets  $S_0$ ,  $S_1$ , and  $S_2$  (dark to light). b) Paint each region (Voronoi regions overlaid). c) Render the results

Voronoi regions in  $S_1$  and  $S_2$  are painted with our painting algorithm. Coverage for regions in  $S_1$  is set to  $V_1$ ,  $0 \ll V_1 < 1$ . The strokes are texture mapped to have partial outlines (Fig. 5).

Regions in  $S_2$  are painted with full coverage. Strokes in these regions are further partitioned based on local variations of the attribute mapped to color. First, the stroke is rotated based on the attribute value mapped to orientation. Next, the data is sampled  $\frac{1}{4}w$  from the stroke's center (where  $w$  represents stroke width) to locate flank colors  $(R_1, G_1, B_1)$  and  $(R_2, G_2, B_2)$ . If the difference between the flank colors  $\sqrt{\Delta R^2 + \Delta B^2 + \Delta G^2} \leq \delta$  for a small positive constant  $\delta$ , the stroke is textured with a complete outline but no center vein (Fig. 5). Otherwise, the stroke is rendered as two quads colored  $(R_1, G_1, B_1)$  and  $(R_2, G_2, B_2)$ , and textured with a complete outline and a center vein.

With this design, the stroke count increases from  $S_0$  to  $S_1$  and again from  $S_1$  to  $S_2$ . The addition of stroke outlines reinforces this gradual increase in level of detail. Fig. 6 shows an example of the ID algorithm.

### 5.3 The Visual Complexity Algorithm



**Figure 7:** Stroke properties in our VC algorithm

VC varies properties of the individual strokes. The VC algorithm places every data element in a single region  $S_0$ , then applies our painting algorithm with full coverage to create a list of strokes, varying the stroke properties as the list is created.

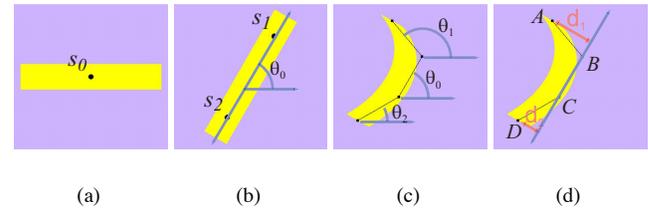
**Varied Stroke Contours.** We scanned hand-painted strokes to construct our texture maps. This introduces more variability to the

footprint of the stroke, compared to rectangular, computer generated strokes. For example, some of the strokes have a rounded end, while others are angular (Fig. 7).

**Variegation.** Strokes are drawn as three quadrilaterals  $q_1$ ,  $q$ , and  $q_2$ , of equal length. The data is sampled at the center of the stroke to choose a primary stroke color  $C = (R, G, B)$ , and the length  $l$  and width  $w$  of the stroke.  $C$  is varied randomly to produce  $C_1 = (R_1, G_1, B_1)$  and  $C_2 = (R_2, G_2, B_2)$  in a small interval around  $C$ 's RGB values.  $w$  is also varied to generate widths  $w_1$  and  $w_2$  such that  $0 \leq w_1, w_2 < \frac{1}{8}w$ .

$q$  is then drawn centered on the sample point with length  $l$ , width  $w - (w_1 + w_2)$ , and color  $C$ .  $q_1$  is drawn next to  $q$  with length  $l$ , width  $w_1$ , and color  $C_1$ , and  $q_2$  is drawn opposite of  $q_1$  with length  $l$ , width  $w_2$ , and color  $C_2$ . This simulates different paint colors being pushed to the margins of a stroke, as often occurs when the brush has residual paint from a previous dip. The randomly colored flanks provides important color variations.

**Curved Strokes.** We render strokes as Bézier curves with four control points,  $A$ ,  $B$ ,  $C$ , and  $D$ . The configuration of a curve is controlled by the data attribute mapped to orientation. We query the attribute values at three sample points,  $s_0$ ,  $s_1$ , and  $s_2$ , then compute the associated orientations to position the stroke's control points.



**Figure 8:** Placing control points,  $A$ ,  $B$ ,  $C$ , and  $D$ : a) Get data sample,  $s_0$ . b) Rotate by  $\theta_0$  and get samples,  $s_1$  and  $s_2$ . c) Place control points, so that  $\overline{BC}$  has orientation  $\theta_0$ ,  $\overline{AB}$  has orientation  $\theta_1$ , and  $\overline{CD}$  has orientation  $\theta_2$ . d) If  $\text{MAX}\{d_1, d_2\}$  is large, draw stroke as a curve.

$s_0$  is taken at the center of the stroke to define its orientation  $\theta_0$  (Figs. 8a).  $s_0$  also determines the stroke's color and size. This is consistent with how rectangular strokes are colored and sized. Since both the curve and  $s_0$  will lie in the convex hull of the control points, the curve will either pass through  $s_0$  or close to it. After the stroke is rotated by  $\theta_0$ , samples  $s_1$  and  $s_2$  are taken  $\frac{1}{4}l$  from the ends of the stroke, yielding orientations  $\theta_1$  and  $\theta_2$ , respectively (Fig. 8b). The control points are then positioned such that that line segment  $\overline{BC}$  has orientation  $\theta_0$ , line segment  $\overline{AB}$  has orientation  $\theta_1$ , and line segment  $\overline{CD}$  has orientation  $\theta_2$  (Fig. 8c). Since Bézier curves are approximating, not interpolating, the curve may not pass through  $B$  and  $C$ , but it will be pulled in their direction. Curved strokes are only drawn where the underlying data attribute mapped to orientation varies rapidly (Fig. 8d).

**Embossed Strokes.** We simulate variations in the thickness of the paint on each stroke using Hertzmann's fast paint texture technique [2002]. The algorithm assigns a height map to each stroke, then generates an overall height field of the image as it is painted. When the scene is lit, shadows generated by the height field simulate the undulations of paint on a canvas. This procedure is meant to avoid rendering expensive three-dimensional mesh models of the strokes, while producing many of the same visual results.

Hertzmann's height field takes into account both the texture of each individual stroke produced by the bristles as they spread the paint and the overlap of paint as strokes are laid atop each other. We

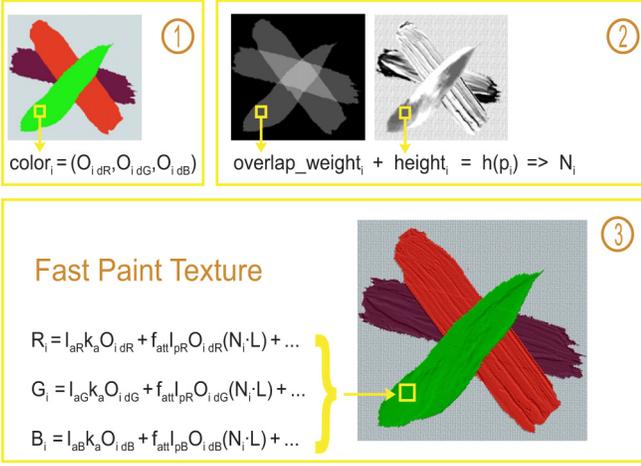


Figure 9: Implementation of Hertzmann’s fast paint texture.

use texture maps paired with opacity maps to calculate these values. We refer to the texture maps as “height maps” here, since the greyscale values are used as height values (dark to light for low to high). The contrast on wet brush stroke texture maps like those in Fig. 5 was increased, to provide more dramatic height differences within the height maps. We create an opacity map, which encodes the stroke’s footprint, to correspond to each height map, by coloring all the points within the stroke white and setting the remaining background transparent. The opacity maps are used to compute a stroke overlap count. The final height field value for a given image position  $(x, y)$  is the sum of the stroke’s height map value at  $(x, y)$  plus a weight representing the number of strokes overlapping at  $(x, y)$ .

To implement fast paint texturing, we draw our stroke list three times. Strokes are rendered differently each time, to collect different information about the scene. The first pass records the colors. The next two passes record the overlap count and the stroke height map information that is used to create the height field. The scene is rendered pixel by pixel, using the color buffer, the height field, and a Phong shading model to calculate each pixel’s color:

1. *Color Pass*: Draw each stroke colored and textured with an opacity map and save the RGB components in the *color* buffer.
2. *Overlap Pass*: Draw each stroke white and textured with the same opacity maps as those used in color pass. The strokes are alpha blended with transparency  $0 < \alpha \ll 1$  to count the number of overlapping strokes at each pixel. Given source and destination blending factors,  $\alpha$  and  $1 - \alpha$ , and source and destination colors,  $(R_s, G_s, B_s, A_s)$  and  $(R_d, G_d, B_d, A_d)$ , respectively, we obtain:

$$R_i = R_s \alpha + R_d (1 - \alpha) \quad (1)$$

$$G_i = G_s \alpha + G_d (1 - \alpha) \quad (2)$$

$$B_i = B_s \alpha + B_d (1 - \alpha) \quad (3)$$

$$A_i = \alpha^2 + A_{id} (1 - \alpha) \quad (4)$$

Starting with a black background sets the initial  $(R_d, G_d, B_d, A_d) = (0, 0, 0, 1)$  at every pixel  $p_i$ . Since each stroke’s color is set to  $(R_s, G_s, B_s, A_s) = (1, 1, 1, \alpha)$ , we can compactly define  $R_{i,k}$ , the red component of pixel  $p_i$

where  $k$  strokes have been drawn, as:

$$R_{i,k} = \begin{cases} 0, & \text{if } k = 0 \\ \alpha + R_{i,k-1}(1 - \alpha), & \text{otherwise} \end{cases} \quad (5)$$

In this way, the final value of any of the color components of a pixel enables us to count the number of strokes that overlap there. We choose to save the red component of the color buffer in the *overlap\_weight* buffer.

3. *Height Pass*: Draw each stroke white and textured with the height map paired with the opacity map used in color and opacity passes. Each pixel  $p_i$  in the resulting color buffer contains the grey value (i.e. the height value) of the last stroke drawn at  $p_i$ . We save the red component of the color buffer in the *height* buffer.
4. Compute the height value at each pixel  $p_i$  as:

$$h(p_i) = \text{overlap\_weight}_i + \text{height}_i \quad (6)$$

Once heights are established, we calculate per-pixel normals using the directional derivatives of the height field.

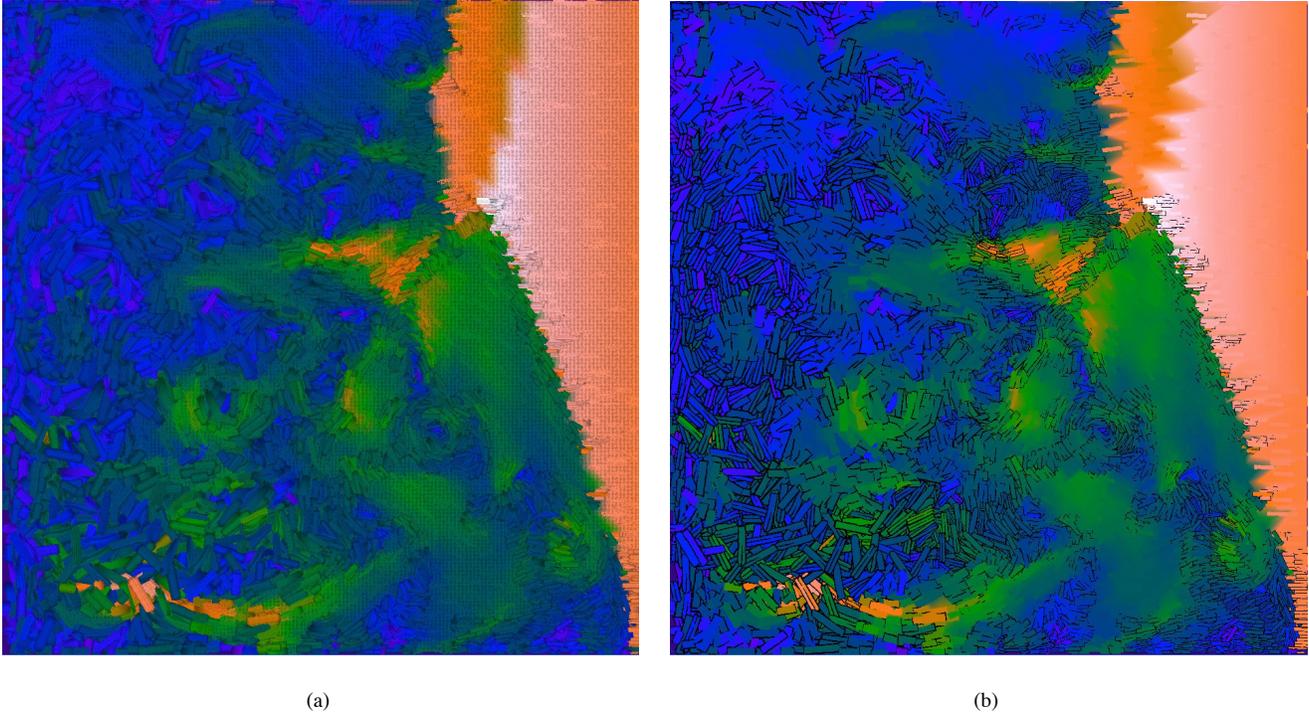
5. For each pixel  $p_i$  we combine the stroke’s diffuse color components in *color<sub>i</sub>* and the surface normal using the Phong model to obtain an RGB color value at  $p_i$ . These value are displayed in the final rendered image (Fig. 9).

**Varied Stroke Aspect Ratio.** This visual feature varies the ratio of stroke length  $l$  to stroke width  $w$ . For a majority of the strokes,  $l$  and  $w$  are selected based on the data attribute  $A_i$  mapped to size.  $A_i$ ’s value determines  $w$ , with  $l = cw$  for some positive constant  $c$ . For a minority of the strokes, however, we let a different attribute  $A_j$ ,  $j \neq i$  control the aspect ratio. If  $A_j$ ’s value is below a preset threshold  $t \gg 0$ ,  $w$  is set based on  $A_i$ ’s value, and  $l = cw$  with  $c = 4$ . If instead  $a_j < t + \frac{\text{MAX}(A_j) - t}{2}$ ,  $w$  is set to a constant value and  $l = c_1 w$  with  $c_1 < 1$ , producing strokes that are wider than they are long. The paint texture reinforces the stroke direction, so that it is not misinterpreted as a stroke with  $l = \frac{1}{c_1} w$ . If neither condition holds,  $w$  is set to a small constant value and  $l = c_2 w$  with  $c_2 \gg c$  to yield strokes that are much longer than they are wide.

**Contrast Strokes.** In Van Gogh wheat field paintings, burgundy strokes contrast in color and direction with a field of golden strokes. We also used the color and orientation of a minority of the strokes to contrast with the majority. A contrast property, either luminance, luminance and orientation, or color, is chosen. If attribute  $A_i$  is mapped to color and  $A_j$  is mapped to orientation,  $A_k$  is mapped to contrast ( $k \neq i \neq j$ ). Stroke color  $C$  and orientation  $\theta$  are selected based on  $A_i$  and  $A_j$ . 10% of the strokes with an  $A_k$  value above a preset threshold are selected to act as contrast strokes. If the contrast property is luminance, the stroke’s luminance is shifted by a constant. If the contrast property is luminance and orientation, the luminance is shifted and the stroke’s orientation is reset to  $-\theta$ . If the contrast property is color, color values in a small interval around  $C$  are assigned to the stroke.

## 6 Practical Applications

We used two real-world data sets to test the practical capabilities of our visualization algorithms: flow from a simulated supernova collapse, and historical weather conditions calculated for positive elevations throughout the world.



**Figure 10:** Visualization of supernova data set, with mapping  $\Delta x$  and  $\Delta y \rightarrow$  orientation, magnitude  $\rightarrow$  color, density  $\rightarrow$  size. a) IC visualization. b) ID visualization.

## 6.1 Supernova Application

A supernova is a massive explosion that occurs at the end of a star's lifetime. Researchers in the Physics Department at NC State University are studying these phenomena. Their current interests include simulating how a supernova collapses. The resulting supernova data sets have time steps and describe a three-dimensional volume of data. Astrophysicists often look at two-dimensional slices of the data at a particular time step to analyze a snapshot of the flow. The slices we are visualizing contain a  $500 \times 500$  regular grid of sample points composed of the attributes  $\Delta x$  and  $\Delta y$  (direction of the velocity vector), *magnitude* (of the velocity vector), *pressure*, and *density*.

We created visualizations of this data using each of our new techniques. To preprocess the data, we first created a mesh by triangulating the sample points and storing *magnitude*, *pressure*, and *density* at each vertex so the simplified mesh would be dense in areas where one or more of the attribute values change rapidly, and sparse in areas where they are relatively constant. Next we calculated the Voronoi diagram of the simplified mesh. This Voronoi diagram was used to partition the data for the IC and ID visualizations.

The IC visualization in Fig. 10a maps *magnitude* to a perceptually balanced color ramp (dark blue to bright pink for low to high).  $\Delta x$  and  $\Delta y$  determine orientation. *Density* is mapped to stroke size (small to large for low to high). *pressure* is not visualized, however, we know that *pressure*, *magnitude*, and *density* are relatively constant where the underpainting is visible. The bright pink area on the right represents a shock wave. The strokes form circular patterns where there are vortices. In Fig. 10b, the data is visualized with the ID technique and the same data-feature mapping. The low and high detail region IC textures in Fig. 10a contrast each other as sharp spatial boundaries. In ID visualization, stroke styles gradually transform across levels of detail.

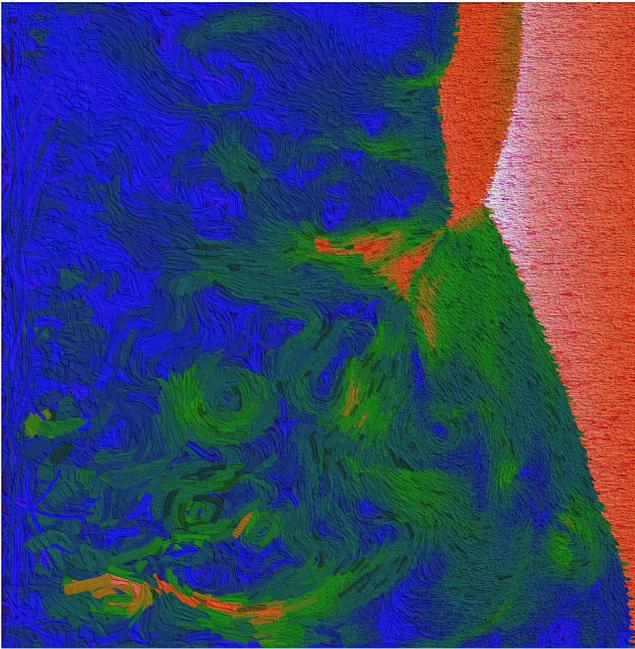
To create a VC visualization, we again mapped *magnitude* to color,  $\Delta x$  and  $\Delta y$  to orientation, and *density* to stroke size. Contrast is used to reinforce magnitude (10% of the strokes have reduced luminance where magnitude  $> 20\%$ ). Pressure is mapped to aspect ratio ( $70\% \leq \textit{pressure} < 85\% \rightarrow$  wide stroke and  $\textit{pressure} \geq 85\% \rightarrow$  long stroke). The resulting visualization is shown in Fig. 11. Wide and long strokes appear on the left where pressure is the highest. The shock wave has contrast strokes, reinforcing the high magnitudes that occur in this region.

Dr. John Blondin, an astrophysicist from the physics department, provided the data and gave us anecdotal feedback on our visualizations. He found the use of glyphs to show multiple attributes in all three techniques assists in the investigation of attribute interactions. He thought visual complexity was the most appealing technique for this application, since the continuity of the glyphs reflects the continuous flow. For example, on the right side of Fig. 11 where the shock wave impact occurs, he could see how the particles were coming together and causing vortices. The sense of flow was reduced in the corresponding regions in Figs. 10a and b. His preference for VC is interesting because he also claimed to find this visualization the most aesthetically pleasing.

## 6.2 Weather Application

Weather data is a second convenient data source, since weather data is plentiful and weather attributes, such as temperature, precipitation, and wind speed, are commonly referenced in day-to-day settings. The data used here, collected by the International Panel on Climate Change between the years 1961 and 1990, consists of 30 year mean monthly weather conditions sampled on a  $\frac{1}{2}^\circ$  latitude by  $\frac{1}{2}^\circ$  longitude regular grid of positive elevations throughout the world.

Figs. 13a and b and 14 show visualizations of this weather data



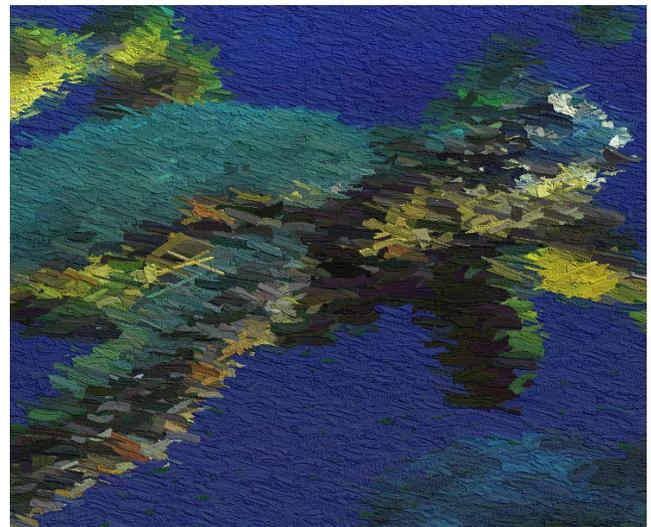
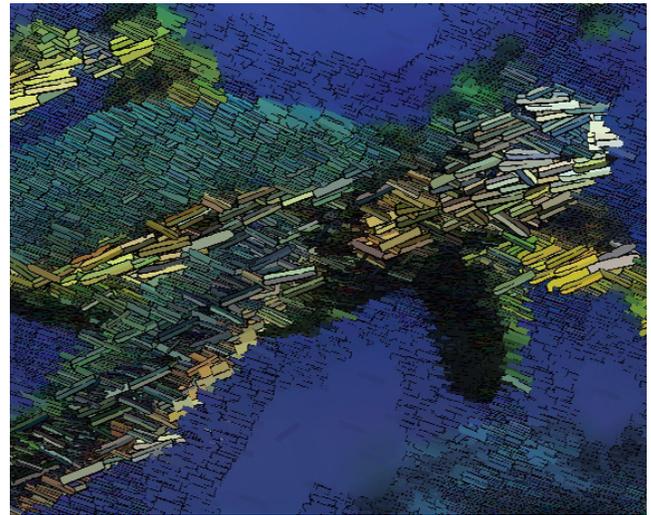
**Figure 11:** VC visualization of supernova data set, with mapping  $\Delta x$  and  $\Delta y \rightarrow$  orientation, magnitude  $\rightarrow$  color, density  $\rightarrow$  size, pressure  $\rightarrow$  aspect ratio; magnitude also controls contrast strokes.

using our three techniques. In Fig. 13a, average African weather conditions in the month of May are visualized. Here *cloud cover*, *mean temperature*, *pressure*, and *diurnal temperature range* are visualized with hue, luminance, size, and orientation. The underlying Voronoi diagram was created with our usual mesh reduction technique, with *temperature*, *wind speed*, and *pressure* stored at each vertex. The luminance is generally higher in northern Africa, indicating the higher temperature there. The canvas shows in interior regions that have relatively constant *temperature*, *wind speed*, and *pressure*.

Mean January South American weather patterns are visualized with our indication and detail technique in Fig. 13b. Here *wind speed*, *minimum temperature*, *wet day frequency*, and *precipitation* are visualized with hue, luminance, size, and orientation. Mesh simplification was performed with *wet day frequency*, *pressure*, and *diurnal temperature range*. The low detail region in the interior of Brazil shows highlight these attributes have low variability. In the Andes, we can see hue variations and a dark luminance, indicating variable wind speeds and low minimum temperatures relative to the rest of the continent.

In Fig. 14, mean weather conditions for January in southwest Canada are visualized with our visual complexity technique. *Mean temperature*, *pressure*, *wet day frequency*, *wind speed*, and *precipitation* are visualized with color, size, orientation, contrast, and aspect ratio, respectively. The temperature is increasing from north to south with high wind speeds in south central regions where contrast strokes are drawn with reduced luminance and opposing orientation. Wide and long strokes appear in the southwest where precipitation is high.

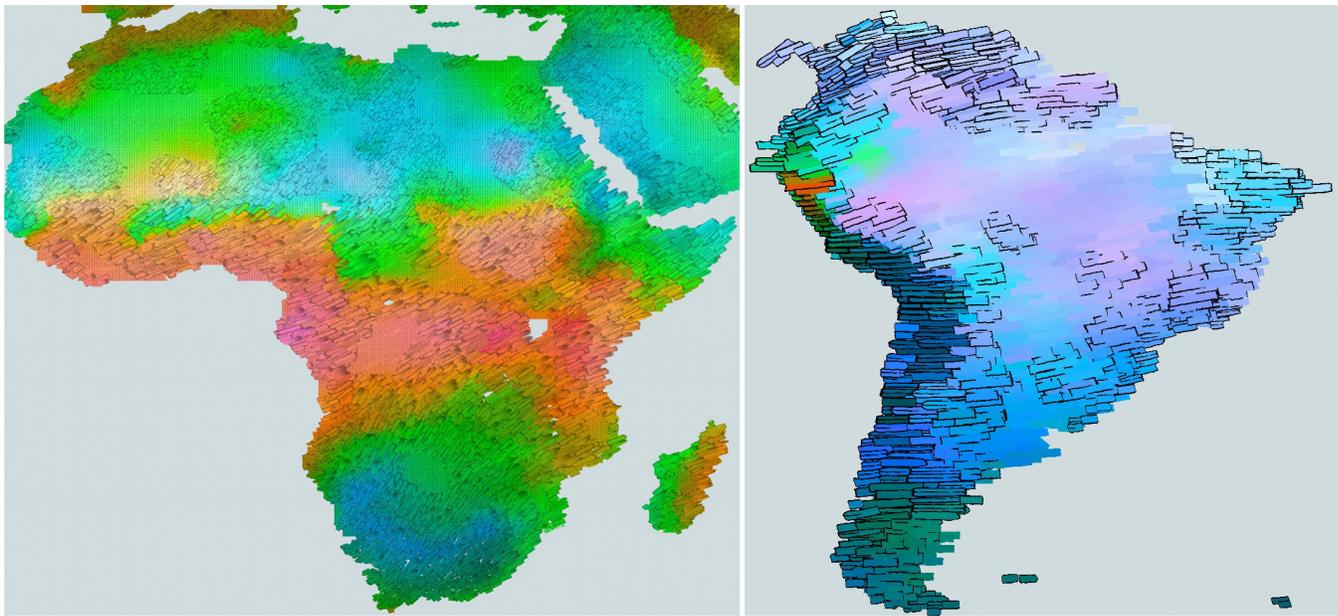
Our techniques are not limited to scientific data visualizations. They can also be applied to other types of input, for example, photographs interpreted as a regular grid of RGB data samples. Fig. 12 shows two examples of an input photograph of a sea turtle, rendered using the ID and VC techniques.



**Figure 12:** A photograph of a sea turtle rendered using: Top the ID algorithm. Bottom the VC algorithm.

## 7 Conclusions

This paper presents new artistic styles for visualizing large scientific data. The styles are designed to both orient viewers with the deliberate use of salient visual features and to engage viewers with aesthetic appeal. We designed these styles by studying psychophysical models of aesthetics and combining them with expressive NPR techniques. We used the new styles to visualize real world data and received encouraging feedback from our domain experts. In fact, our collaborators have started to use our visualizations in their presentations and publications. We are now conducting and analyzing controlled experiments to measure how viewers rate aesthetic, emotional, and visual composition properties of our NPR visualizations relative to real paintings by Abstractionist and Impressionist Masters. Preliminary results have been promising, particularly for IC and VC. This suggests that our techniques can positively impact aesthetic judgment. Given the apparent increase in aesthetic merit produced by our visualizations, the next step in our research will study whether we are increasing engagement, again through the use of controlled experiments.



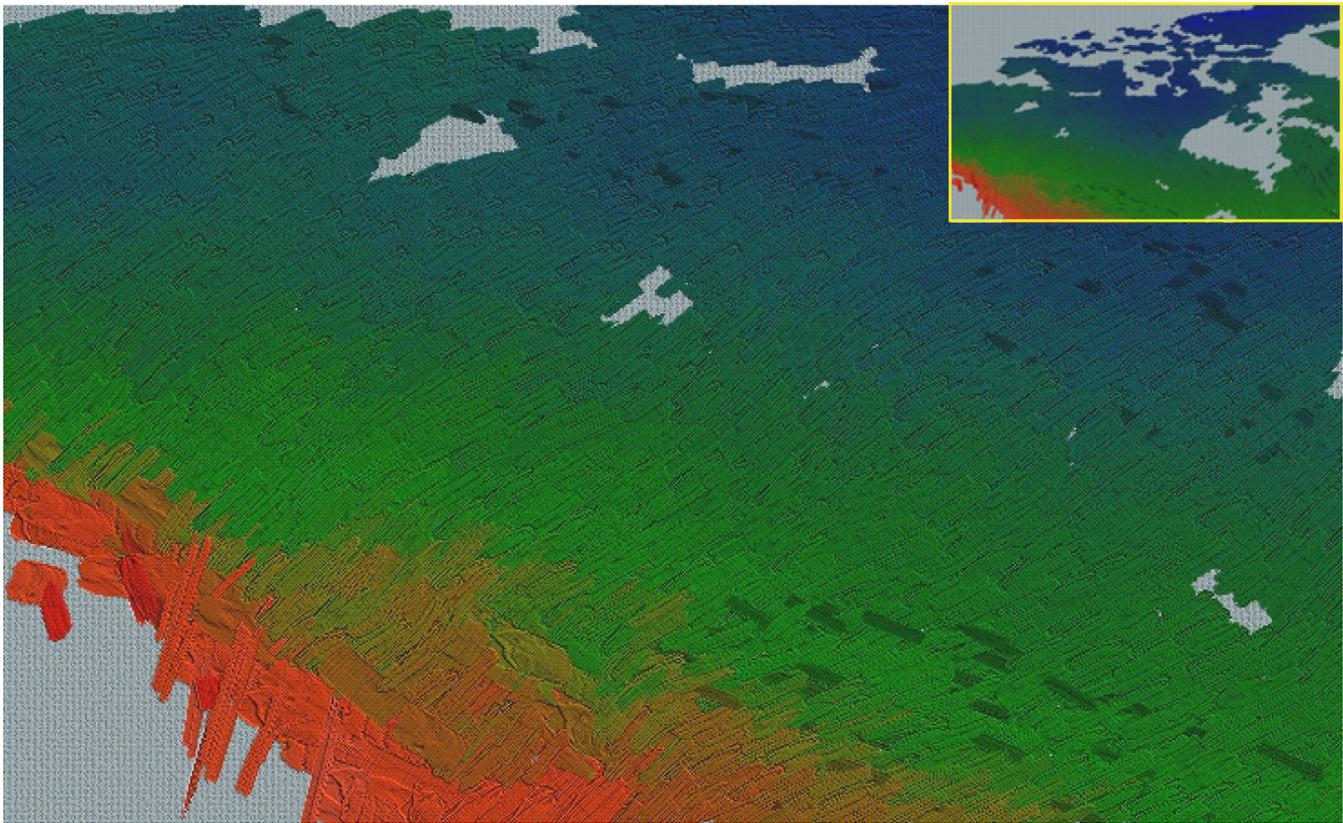
(a)

(b)

**Figure 13:** a) IC visualization of mean weather conditions for May over Africa, cloud cover → hue (blue to red for low to high), temperature → luminance (dark to light for low to high), pressure → size (small to large for low to high), diurnal range → orientation (flat to upright for low to high). b) ID visualization for October over South American, wind speed → hue, minimum temperature → luminance, radiation → size, precipitation → orientation.

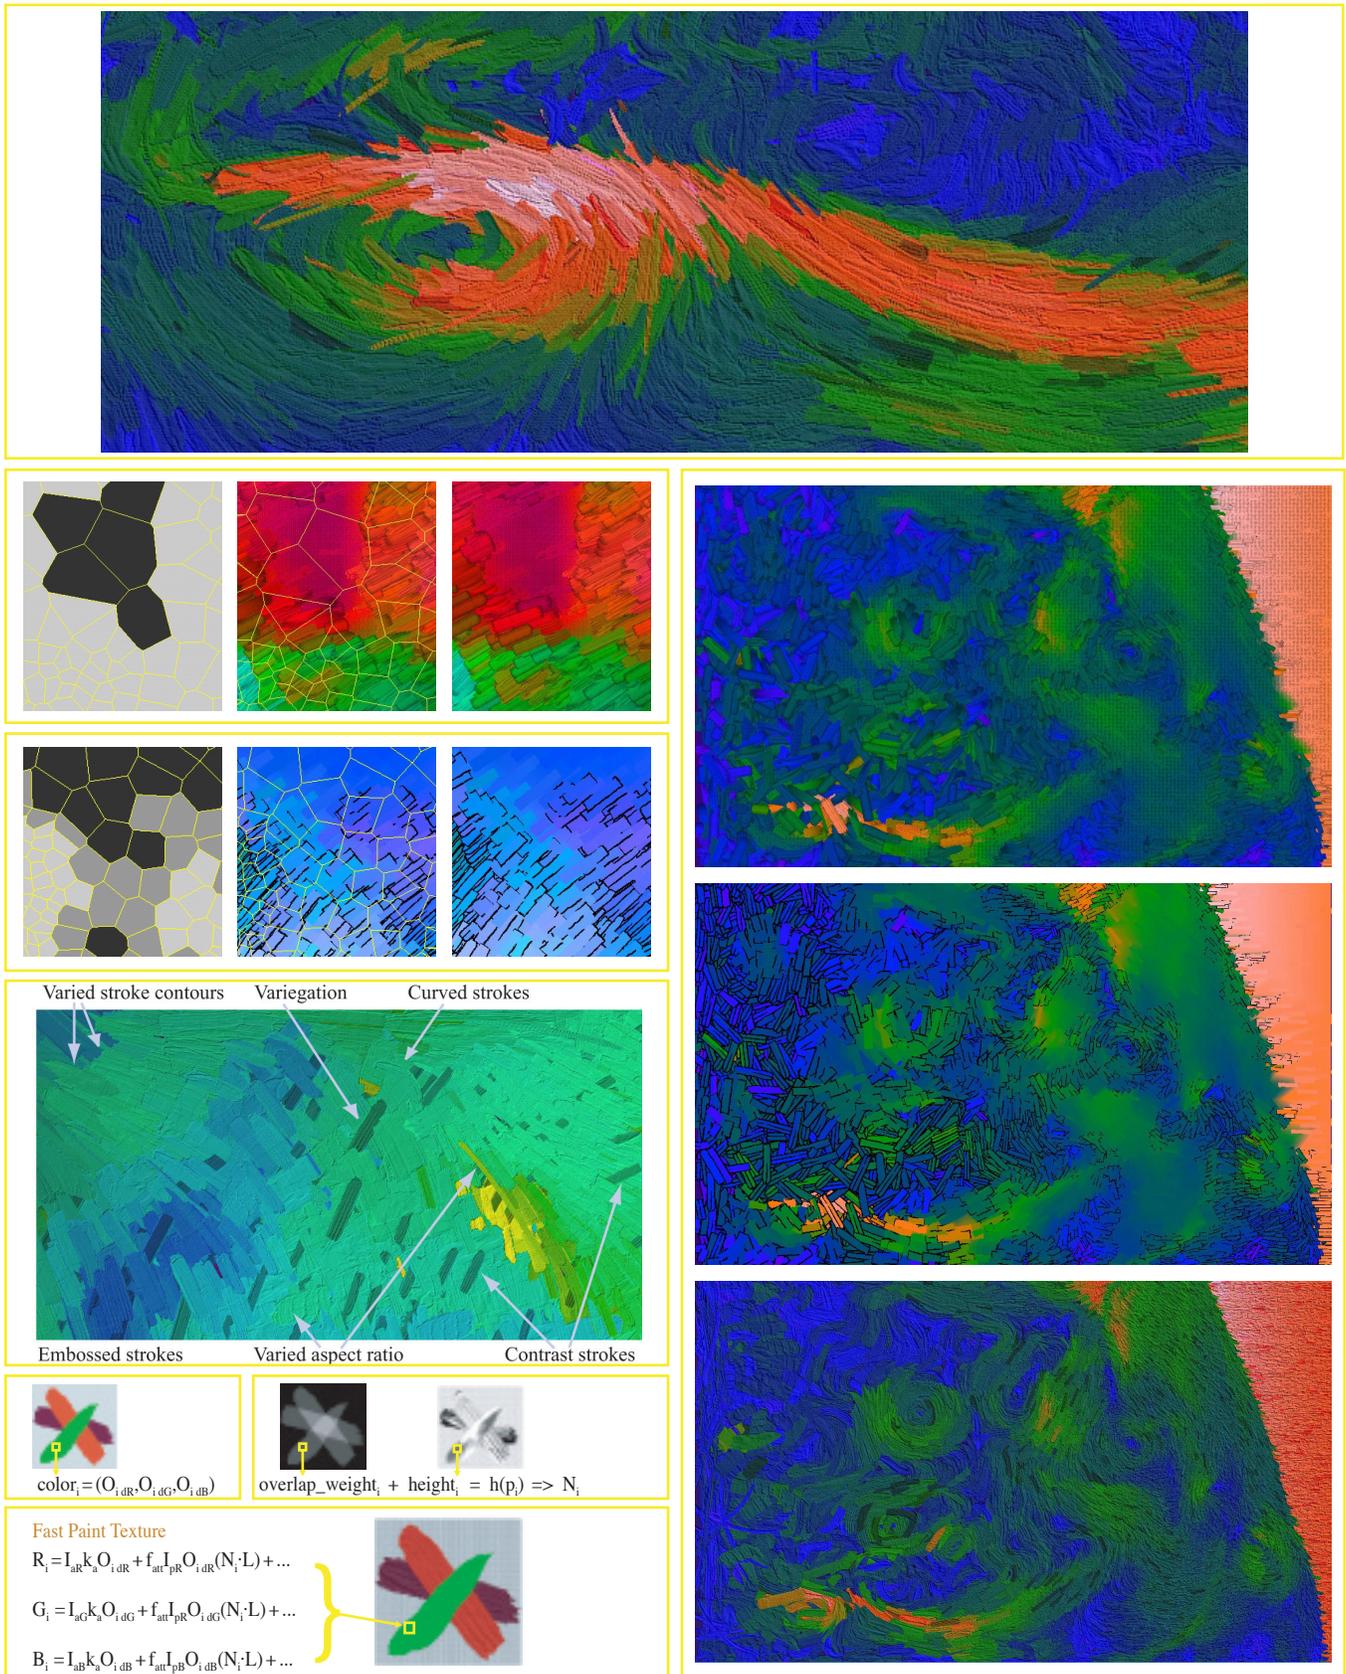
## References

- BALTISSEN, R., AND OSTERMANN, B. 1998. Are the dimensions underlying aesthetic and affective judgment the same? *Empirical Studies of the Arts* 16, 2, 97–113.
- BARBER, C., DOBKIN, D., AND HUHDANPAA, H. 1993. The quickhull algorithm for convex hull. Tech. Rep. TR GCG53, The Geometry Center, University of Minnesota, Minneapolis, MN.
- BARBER, C., DOBKIN, D., AND HUHDANPAA, H. 1996. The quickhull algorithm for convex hulls. *ACM Trans. Math. Softw.* 22, 4, 469–483.
- BERLYNE, D. 1971. *Aesthetics and Psychobiology*. Appleton-Century-Crofts, New York, NY.
- BIRKHOFF, G. 1932. *Aesthetic Measure*. Harvard University Press, Cambridge, MA.
- BROWN, KEVIN, Q. 1979. Voronoi diagrams from convex hulls. *Information Processing Letters* 9, 5, 223–228.
- DAVIS, R. 1936. An evaluation and test of birkhoff's aesthetic measure and formula. *Journal of General Psychology* 15, 231–240.
- GARLAND, M., AND HECKBERT, P. S. 1997. Surface simplification using quadric error metrics. In *SIGGRAPH 97 Conference Proceedings*, T. Whitted, Ed., 209–216.
- GARLAND, M., AND HECKBERT, P. S. 1998. Simplifying surfaces with color and texture using quadric error metrics. In *Proceedings Visualization '98*, 263–269.
- GARLAND, M. 1999. *The Design, Use, and Required Facilities of an Interactive Visual Computer Simulation Language to Explore Production Planning Problems*. Ph.D. thesis, Carnegie Mellon University, Pittsburg, PA.
- GRIGORYAN, G., AND RHEINGANS, P. 2002. Probabilistic surfaces: Point based primitives to show uncertainty. In *Proceedings Visualization 2002*, 147–153.
- HAEBERLI, P. 1990. Paint by numbers: Abstract image representations. *Computer Graphics (SIGGRAPH 90 Conference Proceedings)* 24, 4, 207–214.
- HEALEY, C. G., TATEOSIAN, L., ENNS, J. T., AND REMPLE, M. 2004. Perceptually based brush strokes for nonphotorealistic visualization. *ACM Trans. Graph.* 23, 1, 64–96.
- HERTZMANN, A. 1998. Painterly rendering with curved brush strokes of multiple sizes. In *SIGGRAPH 98 Conference Proceedings*, M. Cohen, Ed., 453–460.
- HERTZMANN, A. 2002. Fast texture maps. In *Proceedings NPAR 2002 Symposium on Non-Photorealistic Animation and Rendering*, 91–96.
- HOLMAN, D., VERTEGAAL, R., SOHN, C., AND CHENG, D. 2004. Attentive display: paintings as attentive user interfaces. In *CHI '04: Extended abstracts of the 2004 conference on Human factors and computing systems*, ACM Press, 1127–1130.
- HSU, S. C., AND LEE, I. H. H. 1994. Drawing and animation using skeletal strokes. In *SIGGRAPH 94 Conference Proceedings*, A. Glassner, Ed., 109–118.
- KIRBY, R. M., MARMANIS, H., AND LAIDLAW, D. H. 1999. Visualizing multivalued data from 2D incompressible flows using concepts from painting. In *Proceedings Visualization '99*, 333–340.
- 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 Visualization '98*, 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*,



**Figure 14:** VC visualization for January for southwest Canada (with overview in corner), temperature  $\rightarrow$  color, pressure  $\rightarrow$  size, wet day frequency  $\rightarrow$  orientation, wind speed  $\rightarrow$  contrast, precipitation  $\rightarrow$  aspect ratio.

- T. Whitted, Ed., 407–414.
- LU, A., MORRIS, C. J., EBERT, D. S., RHEINGANS, P., AND HANSEN, C. 2002. Non-photorealistic volume rendering using stippling techniques. In *Proceedings Visualization 2002*, 211–218.
- MEIER, B. J. 1996. Painterly rendering for animation. In *SIGGRAPH 96 Conference Proceedings*, H. Rushmeier, Ed., 477–484.
- PHAM, B. 1991. Expressive brush strokes. *Computer Vision, Graphics and Image Processing* 53, 1, 1–6.
- RAMACHANDRAN, V. 2000. The science of art: How the brain responds to beauty. In *Understanding wisdom: Sources, science, and society*, W. S. Brown, Ed. Templeton Foundation Press., Philadelphia, PA.
- RHEINGANS, P., AND EBERT, D. 2001. Volume illustration: Non-photorealistic rendering of volume models. *IEEE Transactions on Visualization and Computer Graphics* 7, 3, 253–264.
- SALISBURY, M., ANDERSON, S. E., BARZEL, R., AND SALESIN, D. H. 1994. Interactive pen-and-ink illustrations. In *SIGGRAPH 94 Conference Proceedings*, A. S. Glassner, Ed., 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*, T. Whitted, Ed., 401–406.
- SECOND, A. 2002. Weighted voronoi stippling. In *Proceedings NPAR 2002 Symposium on Non-Photorealistic Animation and Rendering*, 37–43.
- SHIRAIISHI, M., AND YAMAGUCHI, Y. 2000. An algorithm for automatic painterly rendering based on local source image approximation. In *Proceedings NPAR 2000 Symposium on Non-Photorealistic Animation and Rendering*, 53–58.
- STOMPPEL, A., LUM, E., AND MA, K.-L. 2002. Feature-enhanced visualization of multidimensional, multivariate volume data using non-photorealistic rendering techniques. In *10th Pacific Conference on Computer Graphics and Applications*, 1–8.
- STRASSMANN, S. 1986. Hairy brushes. *Computer Graphics (SIGGRAPH 86 Proceedings)* 20, 4, 185–194.
- TATEOSIAN, L. G. 2002. *Nonphotorealistic visualization of multidimensional datasets*. Master’s thesis, North Carolina State University.
- WALTER, J. D., AND HEALEY, C. G. 2001. Attribute preserving dataset simplification. In *VIS ’01: Proceedings of the conference on Visualization ’01*, IEEE Computer Society, 113–120.
- WILSON, B., AND MA, K.-L. 2004. Rendering complexity in computer-generated pen-and-ink illustrations. In *NPAR ’04: Proceedings of the 3rd international symposium on Non-photorealistic animation and rendering*, ACM Press, 129–137.
- WINKENBACH, G., AND SALESIN, D. H. 1994. Computer-generated pen-and-ink illustration. In *SIGGRAPH 94 Conference Proceedings*, A. Glassner, Ed., 91–100.
- WINKENBACH, G., AND SALESIN, D. H. 1996. Rendering parametric surfaces in pen-and-ink. In *SIGGRAPH 96 Conference Proceedings*, H. Rushmeier, Ed., 469–476.
- WOODING, D. 2002. Eye movements of large populations: Ii. deriving regions of interest, coverage, and similarity using fixation maps. *Behavior Research Methods, Instruments, and Computers* 34, 4, 518–528.



**Figure 15:** Top, supernova closeup (Fig. 1); left column, top to bottom: IC process (Fig 4), ID process (Fig 6), VC components (Fig. 7), texture procedure (Fig. 9); right column, top to bottom: IC supernova visualization, ID supernova visualization, VC supernova visualization (Figs. 10, 11)