

Visualization of Multivariate Data Using Preattentive Processing

by

CHRISTOPHER G. HEALEY
B.Math, University of Waterloo, 1990

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF SCIENCE

IN THE FACULTY OF GRADUATE STUDIES
DEPARTMENT OF COMPUTER SCIENCE

We accept this thesis as conforming to the required standard



THE UNIVERSITY OF BRITISH COLUMBIA

October, 1992

© Christopher G. Healey, 1992

Abstract

A new method for designing multivariate data visualization tools is presented. Multivariate data visualization involves representation of data elements with multiple dimensions in a low dimensional environment, such as a computer screen or printed media. Our tools are designed to allow users to perform simple tasks like estimation, target detection, and detection of data boundaries rapidly and accurately. These techniques could be used for large datasets where more traditional techniques do not work well or for time-sensitive applications that require rapid understanding and informative data displays.

Our design technique is based on principles arising in an area of cognitive psychology called preattentive processing. Preattentive processing studies visual features that are “preattentively” detected by the human visual system. Viewers do not have to focus their attention on particular regions of an image to determine whether elements with certain features are present or absent. Examples of preattentive features include colour, orientation, intensity, size, shape, curvature, and line length. Because this ability is part of the low-level human visual system, detection is performed very rapidly, almost certainly using a large degree of parallelism. In this thesis we investigate the hypothesis that these features can be used to effectively represent multivariate data elements. Visualization tools that use this technique will allow users to perform rapid and accurate visual processing of their data displays.

We chose to investigate two known preattentive features, colour and orientation. The particular question investigated is whether rapid and accurate estimation is possible using these preattentive features. Experiments that simulated displays using our preattentive visualization tool were run. The experiments used data similar to that which occurred in a set of salmon migration studies. This choice was made to investigate the likelihood of our techniques being relevant to real-world problems. Analysis of the results of the experiments showed that rapid and accurate estimation is possible with both colour and orientation. A second question, whether interaction occurs between the two features, was answered negatively. Additional information about exposure durations and feature and data interaction were also discovered.

Contents

Abstract	ii
Table of Contents	iii
List of Tables	v
List of Figures	vi
Acknowledgements	viii
1 Introduction	1
2 Related Work	5
Standard Visualization Tools	9
Volume Visualization	13
Flow Visualization	14
Multivariate Data Visualization	15
Visual Interactive Simulation	19
Future Directions	20
3 Preattentive Processing	22
Feature Integration Theory	23
Texton Theory	27
Similarity Theory	28
Three-Dimensional Icons	31
Interference Experiments	33
Iconographic Displays	35
4 Salmon Migration Simulations	39
Standard Visualization Tools	44
5 Psychological Experiments	50
Task Selection	51

Experiment Design	54
Colour Selection	59
Experiment Procedure	66
6 Analysis of Results	69
Estimation Ability	71
Feature Preference	73
Data Type Preference	74
Feature Interference	75
Feature and Data Type Interaction	76
Exposure Duration Experiments	78
7 Conclusions	82
8 Future Work	85
9 Bibliography	87
10 Appendix A	92
Experiment Results and Graphs	92
11 Appendix B	107
Discrimination Experiment Results	107
12 Appendix C	112
Exposure Duration Experiment Results	112

List of Tables

5.1	Experiment Overview	58
5.2	Monitor-Displayable Colours	64
6.1	<i>t</i> -test of Experiment Trial Feature Error Rates	73
6.2	<i>t</i> -test of Control and Experiment Trial Error Rates	75
6.3	<i>F</i> -test of Feature and Data Type Influence on Error Rates	76
6.4	Exposure Duration Estimation Error Rates	79
6.5	<i>t</i> -test of Control and Experiment Trial Exposure Error Rates	81
10.1	Combined Landfall/Colour Results	93
10.2	Combined Landfall/Orientation Results	93
10.3	Combined Stream/Colour Results	94
10.4	Combined Stream/Orientation Results	94
11.1	Subject 1 Discrimination Results	109
11.2	Subject 2 Discrimination Results	109
11.3	Subject 3 Discrimination Results	110
11.4	Subject 4 Discrimination Results	110
11.5	Combined Subject Discrimination Results	111
12.1	Subject 1 Exposure Duration Results	113
12.2	Subject 2 Exposure Duration Results	113
12.3	Subject 3 Exposure Duration Results	113
12.4	Subject 4 Exposure Duration Results	114
12.5	Subject 5 Exposure Duration Results	114
12.6	Combined Subject Exposure Duration Results	114

List of Figures

2.1	Coherency Visualization	17
3.1	Target Detection	23
3.2	Boundary Detection	24
3.3	Feature Map From Early Vision	26
3.4	Textons	27
3.5	N-N Similarity	29
3.6	Three-Dimensional Icons	32
3.7	Hue and Brightness Segregation	33
3.8	Form and Hue Segregation	34
3.9	“Stick-Men” Icons	36
3.10	Chernoff Faces	37
4.1	British Columbia Coast	40
4.2	OSCURS Output	42
4.3	Latitude of Landfall Visualization Tool	45
4.4	Swim Speed Visualization Tool	46
4.5	Date of Landfall Visualization Tool	47
5.1	Rectangle Orientations	54
5.2	Example Experiment Display	56
5.3	Experiment Design Overview	58
5.4	Munsell Colour Space	60
5.5	Discrimination Graph	65
6.1	Experiment Subsection Average Error Graph	72
6.2	Exposure Duration Graph	80
10.1	Landfall/Colour Control Block 1	95
10.2	Landfall/Colour Control Block 2	96

10.3	Landfall/Colour Experiment Block	97
10.4	Landfall/Orientation Control Block 1	98
10.5	Landfall/Orientation Control Block 2	99
10.6	Landfall/Orientation Experiment Block	100
10.7	Stream/Colour Control Block 1	101
10.8	Stream/Colour Control Block 2	102
10.9	Stream/Colour Experiment Block	103
10.10	Stream/Orientation Control Block 1	104
10.11	Stream/Orientation Control Block 2	105
10.12	Stream/Orientation Experiment Block	106

Acknowledgements

I would like to thank both my supervisors Dr. Kellogg Booth (Computer Science) and Dr. James Enns (Psychology) for the effort, support, and guidance they provided me throughout my thesis. Dr. Paul LeBlond (Oceanography) also offered valuable suggestions and access to research being conducted in his department.

A number of people in the Department of Oceanography provided me with ideas, direction, and research results which I used during my experiments. Dr. Keith Thomson designed and ran the original salmon migration simulations, and spent many hours deriving data I needed for my experiments. He also helped me choose an experiment task which addressed a question of interest to Oceanography. Dr. James Ingraham wrote the OSCURS simulation software which was used to run the salmon migration simulations. He spent time modifying his software to provide me with data I needed for my experiments.

Researchers in the Department of Psychology offered a number of insights and suggestions. Dr. Lana Trick helped design the original framework for the psychological experiments. Dr. Ronald Rensink wrote the software which was used to run the experiments. He also provided many helpful suggestions and references in the area of preattentive processing.

Many other people took time and effort to improve my work. Thanks to Mr. Chris Romanzin for reading the final draft of my thesis and proposing a number of important changes, and for looking at my original visualization designs.

Perhaps the most important acknowledgement goes to my parents. Without their support and guidance, most of the opportunities I've enjoyed would not have been possible, including a chance to complete the work contained here.

Financial support for this research was provided in part by a graduate scholarship from the Natural Sciences and Engineering Research Council of Canada and by the University of British Columbia.

Chapter 1

Introduction

“A picture shows me at a glance what it takes dozens of pages of a book to expound”

– Fathers and Sons, *Turgenev*

Many different disciplines use computers to store and analyse data. The data are generated in various ways. Computer simulations are being used to model real-world systems. These programs produce large datasets, reflecting the state of the system at a given point in time. Experiments in areas like physics, chemistry, astronomy, and management science often produce a variety of data. These are usually post-processed using a computer. Real time systems such as air traffic control and medical imaging often require a rapid and intuitive method for displaying information as it is generated.

Different methods have been used to convert raw data into a more usable visual format. The best known example is the conversion of numeric data into different types of graphs. Many commercial statistical languages supply primitives for various types of visualization that go beyond simple bar graphs, pie charts, and scatter plots. Recently, specialized visualization software tools such as the Application Visualization System (AVS), apE, VIS-5D, and the Wavefront Data Visualizer have been developed for computer graphics workstations. The field of visual interactive simulation has emerged to research ways of adding useful visualization and

user interaction to simulation programs.

Empirical methods and guidelines are now needed to build more complex visualization systems. Work in vision and psychology has shown how various object properties affect the human visual system. Certain features such as colour, size, shape, orientation and texture seem to “pop out” of a scene immediately. Careful use of these visual properties may be important when trying to design informative visualization tools. Many of these properties are studied in preattentive processing, a phenomenon investigated in cognitive psychology. Suggestions from this area on how to apply these features to visualization have arisen. The work reported in this thesis represents the first two steps in a three-step process to develop effective visualization tools for a particular class of multivariate data. The first step was examining the literature on preattentive processing to find new techniques for visualization. The second step was designing and experimentally testing prototype visualization tools based on these preattentive features. The final step will involve building robust visualization tools for actual applications by incorporating the experience gained from the prototypes.

The Department of Oceanography at UBC is currently using simulation programs to determine causal effects of ocean currents on sockeye salmon migration patterns. This investigation could benefit from a number of different types of visualization. One common problem is the need to plot multidimensional data in some intuitive form in a two-dimensional environment, such as a computer screen or printed report. Plots may need to be animated in time, to study temporal trends and relationships.

In order to help researchers in Oceanography analyse their simulation data, we designed a set of “standard” visualization tools. These tools were written in an ad-hoc manner, without explicitly using results from research in preattentive processing. Three tools were designed: one to compare latitude of landfall for two different years, one to compare salmon swim speeds for two different years, and one to compare landfall dates for two different years. Strong evidence was found to support oceanography’s hypothesis that ocean currents affect salmon migration patterns. Results from this work are included in two papers to be published in the literature

[Tho92a][Tho92b]. The researchers in Oceanography are currently conducting the second and third phases of their experiment. Data from these simulations will be analysed using the same three visualization tools.

The researchers made a number of suggestions after using our visualization tools. In particular, they wanted to combine multiple displays into a single presentation format. To address this request, we designed a multivariate data visualization tool that explicitly uses results and visual features from preattentive processing. We wanted a tool that worked well with a real-world task and real-world data. Researchers in Oceanography helped us choose a task that was of interest to them. We used data from the salmon migration simulations when we tested the visualization design.

We decided to study the use of colour and orientation to represent multiple data values in a single display. Psychological experiments were designed to test these features. A display that contained 174 rectangles was shown to a subject for 450 milliseconds. Rectangles were coloured either blue or red. Rectangles were oriented either at 0° or at 60° (Figure 5.2). Half of the subjects were asked to estimate the percentage of blue rectangles in the display. The other half were asked to estimate the percentage of rectangles oriented at 60° . Each display simulated a visualization tool presenting salmon migration data. Subjects were performing the task proposed by oceanography when they estimated the number of rectangles with a given preattentive feature. Results from these experiments were analysed to answer the following questions:

- Can estimation be performed reliably within the exposure duration we chose?
- Do subjects perform better with a particular feature (i.e., is one feature better for the estimation task)?
- Do features interfere with one another? Does encoding a data value irrelevant to the estimation task with a secondary preattentive feature affect the subject's estimation ability?

We computed a number of statistics, including t -test and analysis of variance F -test values,

to analyse the data generated during the experiments. Results from these tests indicated the following conclusions:

- rapid and accurate estimation can be performed using both colour and orientation within a 450 millisecond exposure duration
- there is no evidence of a subject preference for either colour or orientation during the estimation task
- there is evidence of a subject preference for the underlying data being displayed during the estimation task
- there is no evidence that orientation interferes with a subject's ability to perform colour estimation
- there is no evidence that colour interferes with a subject's ability to perform orientation estimation
- there is no evidence of interaction between the primary preattentive feature and the underlying data being displayed

Chapter 2 of this thesis describes research related to scientific visualization. The four visualization methods, volume visualization, flow visualization, multivariate data visualization, and visual interactive simulation are examined. Chapter 3 discusses material related to preattentive processing. Three opposing theories that hypothesize how the human visual system performs preattentive processing are explained. Chapter 4 describes Oceanography's salmon migration simulations and the visualization techniques used to examine their data. The task selection and design process for our psychological experiments are presented in Chapter 5. Chapter 6 analyses data from the experiments, statistically testing for a number of hypotheses. Conclusions from the analysis are presented in Chapter 7, and directions for future work are discussed in Chapter 8.

Chapter 2

Related Work

*“Felix qui potuit rerum cognoscere causas
[Lucky is he who could understand the causes of things]”*

– Virgil (70-19 B.C.)

Scientific visualization as a discipline within computer graphics is a relatively recent development. The first reference to “scientific visualization” per se occurred sometime in the late 1980s. Panels and workshops in a variety of different disciplines are now addressing scientific visualization and its relationship to their work [Wol88][Tre89][Bec91b]. The area is expanding into a number of subfields that use computer graphics to solve various types of problems. Examples of these techniques include volume visualization, medical imaging, flow visualization, and multivariate data visualization.

The term “visualization” is often used to refer to the presentation of information to a viewer. Common presentation methods include tables, graphs, pictures, and sound. However, some may question where the actual visualization occurs. Is the presentation of data “visualization”, or does “visualization” occur in the visual and auditory processing of the presentation done by the user? Vande Wetters notes visualization is “...an intelligent, cognitive process performed by human beings” and that visualization tools “...merely facilitate mental insight by the

researchers who use them”¹. Much of the published work to date seems to blur or ignore this distinction. Our notion of “visualization” is the presentation of data in a format specifically designed to make use of the human visual system, especially low-level human vision. In this way we take into account both the representation of data and the cognitive processing performed by the viewer.

Visualization can be defined in a mathematical context to more formally illustrate the path from data to information presentation. The most general setting looks at visualization of arbitrary relations. We will be content to concentrate on functional relationships. We represent these symbolically as

$$f : D \Rightarrow R \tag{2.1}$$

where (informally) the domain D is the set of “independent” values, and the range R is the set of “dependent” results. Assuming D is a subset of \mathbb{R}^m and R is a subset of \mathbb{R}^n then $f : D \Rightarrow R$ can be considered an $(m + n)$ -dimensional dataset. Visualization is a similar mapping of data to some presentation format.

$$g : data \Rightarrow presentation \tag{2.2}$$

Researchers usually want to determine or verify the mapping f . The visualization mapping g is designed to display data in a way that makes it easy to understand or validate f . This can be done in a variety of ways. g can present data elements in a more useable or intuitive manner. It can be used as a filter to extract and juxtapose some subset of the data. It can highlight specific data elements or relationships between different elements. An important aspect of g is the choice of features used to represent data elements. Common features include spatial location within a display, colour, size, shape, orientation, and sound. The set *data* often includes values not only from D and R , but from other sources as well. Known results can be juxtaposed with experimental results to ensure f simulates some real-world system in a logical or expected way. A simulated time axis can be added to a dataset to allow frame-by-frame animation.

¹Visualization Tools: An Introduction. *Pixel*, 1(2), 1990, pg. 32

Visualization tools are utilities that perform the mapping g . They take as input the set *data*, and produce as output a presentation of the data elements in some specified format. Most tools require the user to define g explicitly. Others automate this process by attempting to choose an appropriate mapping based on the information the user wants to extract from the data. This requires knowledge about the type of presentation that is most appropriate for the given task. A number of papers have been written that discuss visualization for specific types of problems and environments. Relatively few papers discuss rules or guidelines for the design and use of general visualization tools. Researchers are now searching for visualization techniques that can be applied to a variety of visualization problems.

Visualization of data and real-world phenomena is an area of study with a long history. Tufte speaks of some 500 years of information representation, starting with the perfection of perspective drawings during the Italian Renaissance [Tuf90]. Books on graphic design show charts and graphs dating back some 150 years to the mid 1800s [Her74] [Tuf83]. More recently, computer simulations have prompted the design of a variety of computer-based visualization techniques. Researchers have found that appropriate visual displays can significantly increase the usefulness of a simulation program. Tools have been written that allow simulation languages to create graphical displays. Programs that post-process data from computer and mathematical models also create graphical displays. These programs can visualize data in a variety of complex and useful ways.

The current stage of research is concerned with the intelligent design of visualization tools. Scientists are now turning to computer graphics, psychology, and visual arts to understand how the human visual system analyses images. This has led to the use of visual properties to make displays more intuitive. These visual properties take advantage of the fundamental workings of the human visual system itself. Some of the most interesting work has used the psychological phenomena known as preattentive processing. This allows users to quickly detect interesting trends or properties in data through the use of abstract features such as colour, texture, shape, and orientation.

Not everyone believes that graphical displays are better than text descriptions. DeSanctis reviewed research that compared graphical and textual displays [DeS84]. He found that most experiments tested one of the following areas:

- interpretation accuracy, to ensure data is correctly interpreted
- problem comprehension, to see how quickly data is understood
- task performance, to see how presentation of data affects solving the task at hand
- decision quality, to see how presentation of data affects the quality of decisions made
- speed of comprehension, to see how quickly data can be summarized
- decision speed, to see how quickly a decision can be reached
- memory for information, to see how presentation of data affects ability to recall specific facts
- viewer preference, to see which type of data presentation users prefer

DeSanctis notes some obvious results, for example that tables are usually better than graphs for obtaining specific data values, but graphs are usually better than tables for identifying trends. He also states some controversial proposals, for example that colour does not enhance comprehension or task performance. In fact, DeSanctis claims that contradictory results exist for all the experiment criteria. In some cases, researchers found graphics were better than text. In other cases, text was favoured over graphics. DeSanctis feels these contradictions imply graphics do not necessarily offer advantages for any of the experimental criteria.

One should remember that DeSanctis focused on the comparison of graphs versus textual tables only. He did not research possible advantages of more complex visual displays. The contradictions he found suggest that graphics should be applied only in certain situations. If used correctly, they can offer an improvement over textual displays. It certainly seems unlikely

that graphical displays never offer an improvement over textual displays in any situation. In fact, DeSanctis suggests that guidelines be developed that tell designers when graphics may be useful.

Standard Visualization Tools

Visualization in scientific computing is attempting to address a number of related problems. The National Science Foundation panel on scientific visualization calls these the “domain of visualization” [McC87], by which they mean the application areas to which visualization is being applied, the tools being brought to bear, and the research problems arising as a result. Three key areas being researched are meaningful representation of data, immediate visual feedback and interaction, and management and analysis of large datasets.

One obvious use of visualization is the presentation of data in a useful, intuitive, or meaningful manner. This is usually done by attaching “features” such as colour, spatial location, and sound to each data element. Features are chosen to show properties within and relationships among data elements. An example of this technique is the presentation of data from computer and mathematical simulations. Originally, visualization of this data consisted of so called “business graphics”. This included printed tables of formulas or numbers and graphs produced using groups of characters to represent bars or lines. Simulation programs that generated the data provided no commands to display results in a concise or coherent manner [Gor69][Fra77][Bra83]. A set of final results was printed when the simulation ended. Printing intermediate values to ensure the program was working correctly often produced a seemingly endless stream of data. Commercial software manufacturers have attempted to solve this problem by providing separate programs or libraries of routines that perform different types of visualization. Some packages take data produced by a simulation program and display it in a meaningful way. Others add commands to the language, allowing visualization to be included as part of the simulation program.

A problem with traditional simulation systems is the lack of immediate visual feedback and interaction. Users want to be able to view their results in real time, as they are being produced. This allows experimentation when interesting phenomena are discovered during simulation execution. Users want to be able to “steer” the simulation, to direct its path to follow interesting trends as the data is generated. This class of software tool is called visual interactive simulation. Interaction will increase productivity if the visualization leads to a good choice. A number of researchers have created visual interactive simulation tools that allow interactive analysis [Mel85][OKe86][Set88]. They provide various empirical and anecdotal results that show visual interactive simulation as an improvement over existing simulation models [Kau81][Bel87].

The requirements for visual interactive simulation are similar to another important class of problem, the visualization of output from real-time applications. Systems like air traffic control require rapid and informative visualization of multivariate data. These displays are often shared by different operators, who visually acquire different data from different parts of the display at the same time. The visualization technique must allow a variety of tasks to be performed rapidly and accurately on dynamically changing subsets of the overall display. Medical imaging systems such as CT, MRI, and ultrasound are another type of application that could benefit from real-time visualization techniques. This method of exploratory analysis presents data to the user in real time. The user analyses the data and decides how to proceed. An informative visualization technique that allows rapid and accurate visual analysis would decrease the amount of time needed to complete the analysis task. This is important, because these types of systems often cannot be time-shared.

Recent developments in computing have created a number of high volume data sources. Examples include supercomputers, satellites, geophysical monitoring stations, and medical scanners. Much of this data is stored without ever being analysed, due in part to the amount of time required to apply traditional analysis techniques. Research in visualization tries to address both the management and presentation of this type of data. The underlying database must allow efficient query and retrieval, filtering, and transformation of data elements. A high speed communication system may be needed to move data from a repository to a workstation

where visualization is performed. Different visualization tools can use a single database and common access methods to obtain the data they require. This data is then presented to the user in some predefined format, depending on the tool being used.

A number of dedicated visualization tools have been written to solve some of the problems mentioned above. They allow post-processing of experiment and simulation data, and are usually run on graphics workstations. The programs offer a number of very powerful visualization techniques such as two- and three-dimensional modelling and rendering, volume rendering, temporal and spatial animation, and even stereo displays. Most of the tools also allow user interaction both before and during the visualization process.

The Application Visualization System (AVS) is a visualization tool built from several data-manipulation modules [Van90a]. It runs as a set of individual UNIX processes combined and controlled by a network flow manager. The AVS modules are grouped into four basic types:

- data sources, that convert data from a given input format into an AVS specific data type (bytes, integers, reals, strings, fields, colour maps, geometric objects, volumes, pixel maps)
- data filters, that convert an input data type to a new output data type. Filters can be used for tasks like cropping and transposing
- data mappers, that convert an input field into a geometric output. Standard mappers include alpha-blending, isosurface rendering, and field flow visualization
- renderers, that allow manipulation and rendering of AVS objects. Standard renderers include a geometry editor, an image editor, and a volume visualization module

Users can place modules in a configuration of their choice to convert, filter, map, and render data in different ways. A network manager binds and controls communication between each module. Users can also write their own modules, to extend the standard AVS package.

apE is a set of visualization tools written at Ohio State University [Van90b]. Like AVS,

apE allows users to place individual modules in a flowchart-like configuration. Input data is converted to a proprietary data format called flux. A network flow manager directs data from a set of input files through each module. Rendering modules at the end of the flow display results in different ways. apE includes its own set of modules to perform data conversion, filtering, transformation, editing, and rendering. Standard display modules include contour and colour maps, isosurfaces, volume visualization, and rendering of geometric objects.

The VIS-5D visualization system was written by the University of Wisconsin Space Science and Engineering Center. It was originally designed to help earth scientists interactively analyse and animate large datasets [Hib90]. VIS-5D expects data to be formatted as a five-dimensional grid. Three dimensions are used to represent spatial location, one to represent time, and one to represent multiple physical variables. Users can interactively change the presentation of the data by controlling the following factors:

- the viewpoint in three dimensions
- the combination of variables being simultaneously displayed
- choice of visual feature for any variable being displayed
- time dynamics
- the spatial extents of the region being shown
- calculation of new variables from existing ones
- an object's transparency

Although AVS, apE, and VIS-5D give users access to previously unavailable visualization techniques, they do have drawbacks. Most excel at a specific subset of tasks required for general visualization. There is no simple way to integrate the packages together to form a single powerful unit. The tools are usually used to post-process already existing data. Their ability to produce data in a simulation-like manner seems to be limited. Finally, many people

do not have access to the equipment or expertise required to use the software. These people need a set of tools that allow intuitive visualization using simple graphic primitives.

Volume Visualization

Volume visualization is a rendering technique that attempts to display data representing volumetric objects or models. Volume visualization allows users to address two related problems:

- a need to visualize three-dimensional objects in order to obtain a better understanding about some aspect of the object
- a need to visualize not only the surface but also the interior of the object

A number of traditional three-dimensional rendering techniques such as scan-line rendering or ray tracing can solve the first requirement. The second requirement is more difficult to satisfy. Early graphics systems employed “cutting planes” to solve the problem. More recently, various methods have been proposed to allow a user to see “inside” the objects they are visualizing.

Yagel, Kaufman, and Zhang have developed a technique that allows a user to render and manipulate objects in a three-dimensional environment [Yag91]. Their system presents a set of tools that are both familiar and intuitive to the user. Volumetric objects are discretized into volume elements (voxels) using a scan-conversion algorithm. The voxel-based objects are then rendered using a modified ray-tracing algorithm called 3D raster ray tracing. Users can manipulate their objects using the following tools:

- a user can place mirrors in the scene to obtain multiple views of an object or objects
- a user can tell the system to cast shadows to help visualize the spatial location and orientation of objects
- a user can change the specularity of an object to help visualize the object’s surface features

- a user can perform a variety of constructive solid geometry operations on an object (e.g., dissection, clipping, filtering, thresholding). One obvious use of these operations would be to clip away part of the surface, to see what was inside the object

Another way of looking inside an object is to make the outer layers of the object semi-transparent. This allows us to “see” through these layers to whatever lies inside the model. Drebin used this technique to visualize medical images [Dre88]. The body’s outer surfaces (tissue and fat) are rendered as though they are transparent. The inner surfaces (bone) are solid. The resulting image is a skeleton surrounded by a semi-transparent “skin”. Users can thus obtain visual information about both the surface and the interior of the body being displayed.

Flow Visualization

Flow visualization is concerned with the simulation, visualization, and analysis of atmospheric and fluid flows. Wind tunnel simulations study how air flows around and through various objects. Fluid mechanics uses visualization tools to analyse phenomena such as vortices, turbulence, and eddies in a fluid flow field.

Sethian has designed a flow visualization tool that runs interactively in real time on a parallel processor computer, the Connection Machine CM-2 [Set88]. Perhaps the most important feature of Sethian’s visualization tool is its real-time interactive ability. Users can watch the progress of a simulation and immediately respond to interesting events. This often involves changing system parameters to explore unexpected results as they evolve. A more complete description of such “visual interactive simulation” techniques is given later in this chapter.

Sethian’s visualization tool allows a user to design the shape of a flow tunnel and place various geometric objects inside it. Coloured particles (smoke and dye) can be interactively injected to follow the fluid flow. Objects that provide a constant stream of particles from a fixed spatial location are also available. These are used to analyse the direction of the flow over

time. Sethian provides a number of examples where unexpected flow fields are easily observed using his visualization software.

A group at the NASA Ames Research Center has been using virtual reality as part of the design of their wind tunnel simulation software [Bry91]. Users interact with their “virtual wind tunnel” using a data glove and a head-mounted display. They can place tufts and coloured particles into their world to help analyse the flow field. Tufts are small vanes that are anchored at a fixed position in the flow. They show changes in the flow field at that spatial location. Coloured particles move through the flow field over time. They form particle paths and streamlines, which help users understand and analyse the shape of the flow field.

Multivariate Data Visualization

Multivariate data visualization deals with the problem of representing high-dimensional data in a low-dimensional environment. High-dimensional data is a set of data elements, each of which encapsulates a relatively large number of data values. An example we will see later in this thesis is a set of data elements that represent migrating salmon. Each data element, or salmon, has a number of data values associated with it, like current geographical location, future location of landfall, swim speed, swim orientation, and type of salmon. Researchers need a way to visually represent this information in a two- or three-dimensional environment. Multivariate data visualization attempts to solve this problem by addressing two related questions. First, is it useful or even reasonable to visually display the data elements with all their associated data values? There may be a limit to the amount of information we can process at one time. Second, if we do want to “see” all the data values, how can we display the information in a low-dimensional environment, such as a computer screen or printed media? One possibility is to display the data so the user can distinguish between different data values. Another possibility is a display that shows trends in the data, rather than individual values.


Treinisch has been developing a flexible multidimensional visualization system [Tre91]. He

is designing an environment that helps scientists analyse the data they produce. This goal has led to a number of system requirements. Treinish feels visualization techniques should take advantage of the power of the human visual system. The system should be usable by the scientist directly and should not require a visualization specialist. A discipline-independent system should handle arbitrary types of data. In order to meet these requirements, Treinish proposes a construct called the visualization pipeline as follows:


- the raw data is stored in a data repository in some abstract, data-independent format (Treinish suggests using the Common Data Format). Users can select the data they want to visualize from the repository
- the selected data can be passed through a number of filters, in order to convert it into the desired format. These may include scaling and projecting the data, as well as conversion to a specific format or coordinate system
- if the data is continuous, it may need to be sampled or gridded to a user-selected resolution
- finally, the data is visualized or rendered using one of a number of visualization tools provided with the system (e.g., two- and three-dimensional plots, contour plots, surfaces, or field flows)

Treinish stresses the importance of the data repository and the supporting data management facilities. He feels this is key to allowing discipline-independent visualization. Data selection and manipulation is separated from data visualization. This gives the user enough flexibility to work with arbitrary data sets and a variety of visualization techniques.

Ware and Beatty have designed a method that uses colour to represent multidimensional data elements [War85][War88]. Data with up to five dimensions can be processed by their system. Elements are represented as spatially located coloured squares. Individual data values are assigned to the following properties of the square representing the data element: x-coordinate, y-coordinate, red intensity, green intensity, and blue intensity. An obvious extension to six dimensions would be to plot elements in a simulated three-dimensional environment.



(a)



(b)

Figure 2.1: Examples of Ware and Beatty’s coherency visualization technique. This 5-dimensional dataset is made up of four coherent groups of elements: (a) data represented by four “clouds” of grey squares; (b) data represented by four “clouds” of coloured squares

Ware and Beatty's visualization tool is used for coherency testing (cluster analysis). Coherency occurs when data elements can be separated into groups of elements with common properties. The data values of the data elements in a given group are similar. Coherent data is represented visually as spatial groups of similarly coloured squares. The data in Figure 2.1a contains either three or four coherent groups. It is difficult to determine whether squares in the center of the figure belong to one or two groups. The use of colour in Figure 2.1b shows very clearly that the data contains four coherent groups. These are represented visually by the four coloured "clouds" of squares. The importance of colour is illustrated by the four overlapping data elements in the center of the figure, two of which are yellow and two of which are purple. Colour allows us to classify these elements into their respective groups. Without colour, we would probably guess all four elements belonged to the purple group. This method allows users to quickly determine whether or not a dataset is coherent. One possible drawback to this technique is the fact that it cannot support data elements with more than five or six dimensions.

Pickett and Grinstein have been using results from cognitive psychology as part of the design of their visualization tools [Pic88][Gri89][Lev90]. Pickett and Grinstein are also developing tools for coherency testing. Their tools display structure in the data as a set of textures and boundaries. Groups of data elements with similar values appear as a spatial group in the display with a unique texture. Different groups are identified by their differing textures.

Pickett and Grinstein's method has a number of important features. The use of texture is based on an area of research in cognitive psychology called preattentive processing. The display mechanism is specifically designed to take advantage of the human low-level vision system. This method is also able to display data with more than five dimensions. This is an improvement over colour contour plots and Ware and Beatty's colour square method. Recently, Pickett and Grinstein have extended their method to use other visual primitives, such as colour [Lev91] and sound. A more detailed explanation of their visualization tools is presented in the "Iconographic Displays" section in the next chapter of this thesis.

Visual Interactive Simulation

Visual interactive simulation (VIS) was a term first conceived by Hurrion in 1976. VIS is concerned with two important problems: visualization of simulation data and user interaction with a running simulation [Bel87]. After building a number of models for various companies in the United Kingdom, Hurrion offered a set of anecdotal observations supporting the use of VIS [Hur80]. He suggested VIS made users more attentive, partly because pictures were more interesting than text and partly because users felt they had more control over the system. Visualization of intermediate results sometimes gave rise to interesting situations the user had never envisioned. This often led to a new line of research and experimentation.

A number of papers have been written showing practical examples of VIS applications. Kaufman showed an example of converting a traditional simulation program into an interactive program that used graphics [Kau81]. Researchers in the United Kingdom have developed a set of library routines that allow VIS on an Apple microcomputer [OKe86]. After applying their software to various practical simulation problems, they concluded that a machine with relatively simple graphic primitives and computational power offered an acceptable platform for VIS software.

Scientists at AT&T Bell Labs have developed a general VIS package, the Performance Analysis Workstation [Mel85]. The workstation allows design and testing of queueing networks. All the tasks on the workstation are graphical. Users design the network topology by positioning resources on the screen and placing links between them. They can allocate screen space for various graphs or informational messages. The simulation can be started, interrupted, modified, and monitored using commands available through pop-up menus. The workstation visually displays activity in the network, including movement and queueing of requests as the simulation executes.

Workgroups in Bell Labs have been using the Performance Analysis Workstation for a variety of applications. Results from the workstation mirror the observations of Hurrion, specifically

that users enjoy using the workstation, and that interesting phenomena often arise while the simulation is running.

Future Directions

Various panels and workgroups that consulted scientists studying visualization seem to confirm opinions suggested in previous work. These people provided some interesting insight on the current state of scientific visualization [McC87][Bro88][Tre89]. They feel visualization is made up of a number of different disciplines, including physical science, computer graphics and computer science, psychology, and visual arts. The recent need for some form of intelligent data representation has stimulated the formation of interdisciplinary groups. These groups are pooling their expertise in an attempt to come up with guidelines and solutions to the visualization problem.

Another interesting point was the strong support shown for visual interactive simulation. The panel members agreed that up to this point visualization has been a post-production task. Scientists visualize data that was obtained some time in the past. People now want to produce visual displays while the data is being generated. Through user interaction, they want to be able to “steer” the path of computation in response to the visualization they see while the simulation runs. This is exactly the concept Bell and O’Keefe suggested in their research paper [Bel87]. Visual interaction simulation methods can also be used to visualize data from real-time systems such as air traffic control and medical imaging systems.

Finally, researchers recognize the fact that not everyone has the sophisticated computer equipment needed to generate complex visual images. This creates two additional visualization requirements. First, there should be some simple way to transport data from where it is generated to where it will be visualized. Various groups are working to produce advanced networking tools that will allow supercomputers and workstations to communicate efficiently. Second, there is an understanding that visualization does not need to be complicated in order

to be useful; some data can benefit from advanced techniques such as stereo display or volume rendering, but most data can be visualized using simple wireframe objects or colour-coded plots. This type of simple visualization can be done on personal workstations or even personal computers.

In order to address these requirements, we decided to use the built-in processing of the human visual system to assist with visualization. In particular, we studied results from an area of cognitive psychology called preattentive processing. Preattentive processing describes a set of simple visual features that are detected in parallel by the low-level human visual system. Background information on preattentive processing is presented in the next chapter. We hypothesize that the use of preattentive features in a visualization tool will allow users to perform rapid and accurate visual tasks such as grouping of similar data elements, detection of elements with a unique characteristic, and estimation of the number of elements with a given value or values. We tested this hypothesis using psychological experiments that simulated a preattentive visualization tool. This is described in two chapters later in this thesis.

Chapter 3

Preattentive Processing

“One man does not see everything in a glance”

– Phœnissœ, *Euripides*

Researchers in psychology and vision have been working to explain how the human visual system analyses images. One interesting result has been the discovery of visual properties that are “preattentively” processed. These properties are detected immediately by the visual system. Viewers do not have to focus their attention on an image to determine whether elements with the given property are present or absent.

An example of preattentive processing is detecting a filled circle in a group of empty circles (Figure 3.1). The target object contains a preattentive property, “filled”, that the distractor objects do not. A viewer can quickly glance at the image to determine whether the target is present or absent.

Properties that are preattentively processed can be used to highlight important image characteristics. Experiments in psychology have used preattentive properties to assist in performing the following visual tasks:

- target detection, where users attempt to rapidly and accurately detect the presence or absence of a “target” element that uses a unique preattentive property within a field of




Figure 3.1: Examples of target detection: (a) target can be preattentively detected because it contains the unique feature “filled”; (b) filled circle target cannot be preattentively detected because it contains no preattentive feature unique from its distractors

distractor elements (Figure 3.1)

- boundary detection, where users attempt to rapidly and accurately detect a texture boundary between two groups of elements, where all of the elements in each group have a common preattentive property (Figure 3.2)
- counting, where users attempt to count or estimate the number of elements in a display that use a unique preattentive property

Feature Integration Theory

Triesman has provided some exciting insight into preattentive processing by researching two important problems [Tri85]. First, she has tried to determine which properties are detected preattentively. She calls these properties “preattentive features”. Second, she has formulated a hypothesis about how the human visual system performs preattentive processing.

Triesman ran experiments using target and boundary detection to classify preattentive features. For target detection, subjects had to determine whether the target was present or absent in a field of distractors. Boundary detection involved placing a group of objects that




Figure 3.2: Examples of boundary detection: (a) the horizontal boundary between two groups is preattentively processed because each group contains a unique feature; (b) the vertical boundary is not apparent, because both groups use the same features (filled versus empty and square versus circle)

used a unique feature within a set of distractors to see if the boundary could be preattentively detected.

Researchers test for preattentive target detection by varying the number of distractors in a scene. If search time is relatively constant and below some chosen threshold, independent of the number of distractors, the search is said to be preattentive. Similarly, for boundary detection, if users can classify the boundary within some fixed time, the feature used to define the boundary is said to be preattentive. A common threshold time is 250 milliseconds, because this allows subjects “one look” at the scene. The human visual system cannot decide to change where the eye is looking within this time frame.

Objects that are made up of a conjunction of unique features cannot be detected preattentively. A conjunction occurs when the target object is made up of two or more features, each of which is contained in the distractor objects. Figure 3.1b shows an example of a conjunction target. The target is made up of two features, filled and circle. Both of these features occur in the distractor objects (filled squares and empty circles). Thus, the target cannot be preattentively detected.

Triesman has compiled a list of features that can be preattentively detected [Tri85][Tri88]. These features include line length, orientation, contrast, hue, curvature, and closure. It is important to note that some of these features are asymmetric. For example, a sloped line in a sea of vertical lines can be preattentively detected. However, a vertical line in a sea of sloped lines cannot be preattentively detected. Another important consideration is the effect of different types of background distractors on the feature target. These kinds of factors must be taken into consideration when trying to design systems that rely on preattentive processing.

Triesman breaks human early vision into a set of feature maps and a master map of locations in an effort to explain preattentive processing. Each feature map registers activity in response to a given feature. Triesman proposes a manageable number of feature maps, including one for each of the human vision colour primaries red, yellow, and blue as well as separate maps for orientation, shape, texture, and other preattentive features.

When the human visual system first sees an image, all the features are encoded in parallel into their respective maps. One can check to see if there is activity in a given map, and perhaps get some indication of the amount of activity. The individual feature maps give no information about location, spatial arrangement, or relationships to activity in other maps.

The master map of locations holds information about intensity or hue discontinuities at given spatial locations. Focused attention acts through the master map. By examining a given location, one automatically gets information about all the features present at that location. This is provided through a set of links to individual feature maps (Figure 3.3).

This framework provides a general hypothesis that explains how preattentive processing occurs. If the target has a unique feature, one can simply access the given feature map to see if any activity is occurring there. Feature maps are encoded in parallel, so feature detection is almost instantaneous. A conjunction target cannot be detected by accessing an individual feature map. Activity there may be caused by the target, or by distractors that share the given preattentive feature. In order to locate the target, one must search serially through the master map of locations, looking for an object with the correct features. This requires focused

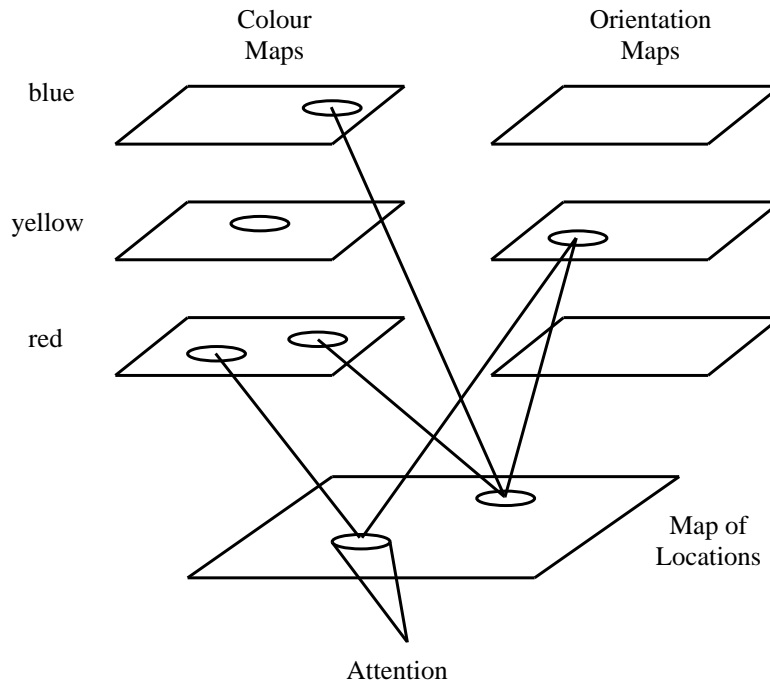


Figure 3.3: Framework for early vision that explains preattentive processing. Individual maps can be accessed to detect feature activity. Focused attention acts through a serial scan of the master map of locations.

attention and a relatively large amount of time.

In later work, Triesman has expanded her strict dichotomy of features being detected either in parallel or in serial [Tri88][Tri91]. She now believes that parallel and serial represent two ends of a spectrum. “More” and “less” are also encoded on this spectrum, not just “present” and “absent”. The amount of differentiation between the target and the distractors for a given feature will affect search time. For example, a long vertical line can be detected immediately among a group of short vertical lines. As the length of the target shrinks, the search time increases, because the target is harder to distinguish from the distractors. At some point, the target line becomes shorter than the distractors. If the length of the target continues to




Figure 3.4: Example of similar textons: (a) two textons that appear different in isolation; (b) the same two textons cannot be distinguished in a randomly oriented texture environment

decrease, search time decreases, because the degree of similarity between the target and the distractors is decreasing.

Texton Theory

Texture segregation involves preattentively locating groups of similar objects and the boundaries that separate them. Triesman used texture segregation during her experiments with boundary detection. Figure 3.2a is an example of a horizontal texture boundary with empty shapes on the top and filled shapes on the bottom. Figure 3.2b is an example of a vertical texture boundary with filled circles and empty squares on the left, and empty circles and filled squares on the right.

Julèsz has also investigated texture perception and its relationship to preattentive process-

ing [Jul81][Jul83][Jul84]. He has proposed his own hypothesis on how preattentive processing occurs. Julèsz believes that the early vision system detects a group of features called textons. Textons can be classified into three general categories:

1. Elongated blobs (e.g., line segments, rectangles, ellipses) with specific properties such as hue, orientation, and width
2. Terminators (ends of line segments)
3. Crossings of line segments

Julèsz believes that only differences in textons or their density can be detected preattentively. No positional information about neighboring textons is available without focused attention. Like Triesman, Julèsz believes preattentive processing occurs in parallel and focused attention occurs in serial.

Figure 3.4 shows an example of an image that supports the texton hypothesis. Although the two objects look very different in isolation, they are actually the same texton. Both are blobs with the same height and width. Both are made up of the same set of line segments and each has two terminators. When oriented randomly in an image, one cannot preattentively detect the texture boundary between the two groups of objects.

Similarity Theory

Some researchers do not support the dichotomy of serial and parallel search modes. Initial work in this area was done by Quinlan and Humphreys [Qui87]. They investigated conjunction searches by focusing on two factors. First, search time may depend on the number of items of information required to identify the target. Second, search time may depend on how easily a target can be distinguished from its distractors, regardless of the presence of unique preattentive features. Note that in later work Triesman addressed this second factor [Tri88]. Quinlan and




Figure 3.5: Example of N-N similarity affecting search efficiency: (a) high N-N similarity allows easy detection of target shaped like the letter L; (b) low N-N similarity increases difficulty of detecting target shaped like the letter L

Humphreys found that Triesman's feature integration theory was unable to explain all the results obtained from the above experiments.

Duncan and Humphreys have now proposed their own explanation of preattentive processing, which assumes search ability varies continuously depending on the type of task and display conditions [Dun89][Mül90]. According to their theory, search time depends on two important criteria, T-N similarity and N-N similarity. T-N similarity is the amount of similarity between the targets and nontargets (distractors). N-N similarity is the amount of similarity within the nontargets themselves. These two factors affect search time as follows:

- as T-N similarity increases, search efficiency decreases and search time increases
- as N-N similarity decreases, search efficiency decreases and search time increases
- T-N similarity and N-N similarity are related (Figure 3.5). Decreasing N-N similarity has little effect if T-N similarity is low. Increasing T-N similarity has little effect if N-N

similarity is high

Triesman's feature integration theory has difficulty explaining the results of Figure 3.5. In both cases, the distractors seem to use exactly the same features as the target, namely oriented, connected lines of a fixed length. Yet experimental results show displays similar to Figure 3.5a produce an average search time increase of 4.5 milliseconds per additional distractor, while displays similar to Figure 3.5b produce an average search time increase of 54.5 milliseconds per additional distractor.

In order to explain the above and other search phenomena, Duncan and Humphreys propose a three-step theory of visual selection.

1. The visual field is segmented into structural units. Individual structural units share some common property (e.g., spatial proximity, hue, shape, motion). Each structural unit may again be segmented into smaller units. This produces a hierarchical representation of the visual field. Within the hierarchy, each structural unit is described by a set of properties (e.g., spatial location, hue, texture, size). This segmentation process occurs in parallel.
2. Because access to visual short-term memory is limited, Duncan and Humphreys assume that there exists a limited resource that is allocated among structural units. Because vision is being directed to search for particular information, a template of the information being sought is available. Each structural unit is compared to this template. The better the match, the more resources allocated to the given structural unit relative to other units with a poorer match.

Because units are grouped in a hierarchy, a poor match between the template and a structural unit allows efficient rejection of other units that are strongly grouped to the rejected unit.

3. Structural units with a relatively large number of resources have the highest probability of access to the visual short-term memory. Thus, structural units that most closely match

the template of information being sought are presented to the visual short-term memory first. Search speed is a function of the speed of resource allocation and the amount of competition for access to the visual short-term memory.


Given these three steps, we can see how T-N and N-N similarity affect search efficiency. Increased T-N similarity means more structural units match the template, so competition for visual short-term memory access increases. Decreased N-N similarity means we cannot efficiently reject large numbers of strongly grouped structural units, so resource allocation time and search time increases.

Three-Dimensional Icons


To date, most of the features identified as preattentive have been relatively simple properties. Examples include hue, orientation, line length, and size. Enns and Rensink have identified a class of three-dimensional elements that can be detected preattentively [Enn90a][Enn90b]. They have shown that the three-dimensionality is what makes the elements “pop-out” of the visual scene. This is important, because it suggests that more complex high-level concepts may be processed preattentively by the low-level vision system.

Figure 3.6 shows an example of these three-dimensional icons. The elements in Figure 3.6a are made up of three planes. The planes are arranged to form an element that looks like a three-dimensional cube. Subjects can preattentively detect the group of cubes with a three-dimensional orientation that differs from the distractors. The elements in Figure 3.6b are made up of the same three planes. However, the planes are arranged to produce an element with no apparent three-dimensionality. Subjects cannot preattentively detect the group of elements that have been rotated 180 degrees. Apparently, information about three-dimensional orientation is preattentively processed.

Enns and Rensink have also shown how lighting and shadows provide three-dimensional information that is preattentively processed [Enn90c][Enn90d]. Spheres are drawn with shadows



(a)



(b)

Figure 3.6: Three-dimensional icons: (a) when the cubes appear “three-dimensional”, the group with a different orientation is preattentively detected; (b) when three-dimensional cues are removed, the unique group cannot be preattentively detected




Figure 3.7: Hue and brightness segregation: (a) hue texture boundary with varying brightness across the array; (b) brightness texture boundary with varying hue across the array

so they appear to be lit either from above or from below. Subjects can preattentively detect the group of spheres that appear to be lit differently than the distractors.

Interference Experiments

Callaghan has done research to see how similarity within feature groups affects texture segregation [Cal90]. She found that varying certain irrelevant features within a group can interfere with boundary detection. Her initial experiments dealt with location of a horizontal or vertical texture boundary [Cal84]. Subjects were presented with a six by six array of elements. The texture boundary was formed by either differences in hue or differences in brightness. For hue segregation, the brightness in both groups varied randomly between two values. For brightness segregation, hue varied randomly between two values (Figure 3.7). Subjects had to determine whether the texture boundary was vertical or horizontal. Control experiments were run to see how quickly subjects could detect simple hue and brightness boundaries. The control arrays had a uniform brightness during hue segregation, and a uniform hue during brightness




Figure 3.8: Form and hue segregation: (a) hue boundary is preattentively detected, even though form varies in both groups; (b) hue interferes with detection of form boundary

segregation.

Callaghan found that non-uniform brightness will interfere with hue segregation. It took subjects longer to determine where the texture boundary occurred, relative to the control array. However, a non-uniform hue did not interfere with brightness segregation. A brightness texture boundary can be detected in a fixed amount of time, regardless of whether the hue fluctuates or not. This asymmetry was verified through further experimentation [Cal90].

Callaghan's more recent work has shown a similar asymmetry between form and hue [Cal89]. As before, subjects were asked to locate a horizontal or vertical texture boundary in a six by six array. During the experiment, the arrays were segregated by either hue or form. For hue segregation, form varied randomly within the array (circle or square). For form segregation, hue varied randomly. Results showed that variation of hue interfered with form segregation. However, variation of form did not interfere with hue segregation (Figure 3.8).

These interference asymmetries suggest some preattentive features may be "more important" than others. The visual system reports information on one type of feature over and above

others that may also be present. Callaghan's experiments suggest that brightness overrides hue information and that hue overrides shape information.



Iconographic Displays

Pickett and Grinstein have been working to develop a method of displaying multidimensional data in a two-dimensional environment [Pic88]. The two most common display mediums, computer screens and printed documents, are both two-dimensional.

Initially, work has focused on spatially or temporally coherent data sets. This type of data generally contains clusters of data elements with similar values. Previously, data with up to three dimensions was plotted in colour. Each data dimension controlled the intensity of one of the three primary colours red, green, and blue. Coherent areas within the data set occurred where colour values were similar. Relationships between data elements were shown as spatial changes in colour. Pickett decided to use texture as a medium that could show relationships among higher dimensional data. The texture segregations would preattentively display areas with similar data elements. Researchers could then quickly decide whether further analysis was required.

Pickett developed icons that could be used to represent each data element. An icon consists of a body segment plus a number of limbs (Figure 3.9). Each value in the data element controls the angle of one of the limbs. The icons in Figure 3.9 can support five-dimensional data elements. The four limbs support the first four dimensions. The final dimension controls the orientation of the icon's body in the image.

Once the data-icon mapping is defined, an icon can be produced for each data element. These icons are then displayed on a two-dimensional grid in some logical fashion. The result is an image that contains various textures that can be preattentively detected. Groups of data elements with similar values produce similar icons. These icons, when displayed as a group, form a texture pattern in the image. The boundary of this pattern can be seen, because icons



(d)

Figure 3.9: Examples of “stick-men” icons: all three icons have four limbs and one body segment (shown in bold), so they can support five-dimensional data elements; (d) an iconographic display of 5-D weather satellite data from the west end of Lake Ontario



Figure 3.10: Examples of Chernoff faces: notice that the various facial characteristics, such as nose length, eyes, mouth, jowls, etc., are controlled by various data values in each face

outside the given group have a different form and produce a different texture pattern.

The key to this technique is designing icons that, when displayed, produce a good texture boundary between groups of similar data elements. To this end, Pickett and Grinstein have been working on an icon toolkit, to allow users to produce and test a variety of icons with their data sets [Gri89]. They have also added an audio component to the icons. Running the mouse across the image will produce a set of tones. Like the icons, the tones are mapped to the values in each data element. It is believed these tones will allow researchers to detect interesting relationships in their data.

Other researchers have suggested using various types of “icons” to plot individual data elements. One of the more unique suggestions has been to use faces (Figure 3.10) with different expressions to represent multidimensional data [Che73][Bru78].

Each data value in a multidimensional data element controls an individual facial characteristic. Examples of these characteristics include the nose, eyes, eyebrows, mouth, and jowls. Chernoff claims he can support data with up to eighteen dimensions. He also claims groupings in coherent data will be drawn as groups of icons with similar facial expressions. This technique seems to be more suited to summarizing multidimensional data elements, rather than segmenting them. Still, it shows that researchers are exploring a wide variety of ideas.

We will look at two preattentive features, hue and orientation, and will investigate their use for a common task, estimation. Psychological experiments that simulate a visualization tool for estimation using these features will be described in Chapter 5.

Chapter 4

Salmon Migration Simulations

*“Di qua, di là, di giù, di su li mena
[Hither, tither, upward, downward they are driven]”*

– Inferno, *Dante*

A salmon is a well-known type of fish that lives, among other areas, on the western Canadian coast. Salmon are split into six species: Atlantic, Sockeye, Coho, Chinook, Chum, and Pink. The life history of a salmon consists of four stages: birth, freshwater growth stage, ocean growth stage, and migration and spawning [Pea92]. Salmon are born as fry in freshwater rivers and streams. After birth, the fry spend time feeding and maturing before swimming downstream to the open ocean. The amount of time spent can vary widely from a few months to several years, depending on the type of salmon. Upon reaching the ocean, the salmon moves to its “open ocean habitat”, where it spends most of its ocean life feeding and growing. Ocean habitats for the various species of salmon are not well defined. Sockeye salmon are thought to feed in the Subarctic Domain, an area of the Pacific Ocean north of 40° latitude stretching from the coast of Alaska to the Bering Sea. After a period of one to six years, salmon begin their migration run. This consists of an open ocean stage back to the British Columbia coast and a coastal stage back to a freshwater stream to spawn. Salmon almost always spawn in the stream where they were born. Scientists now know salmon find their stream of birth using smell when they




Figure 4.1: The British Columbia coast, showing Vancouver Island, the Juan de Fuca Strait, the Johnstone Strait, and the Fraser River. Arrows represent the migration patterns of returning salmon


reach the coast. The direction finding methods used to navigate from the open ocean habitat to the coast are still being researched. Once the salmon reaches its stream of birth, it spawns and then dies. Typical spawning produces between 2,000 and 5,000 eggs.

The ocean phase of sockeye salmon migration is not well understood [LeB90]. It is recognized that it is rapid, well directed, and well timed. Previous work has examined the effect of climate and ocean conditions during migration to see how they affect the point of landfall for Fraser River sockeye. The entrance to the Fraser River is located on the southwest coast of British Columbia, near Vancouver. When the Gulf of Alaska is warm, sockeye make landfall at the north end of Vancouver Island and approach the Fraser River primarily via a northern


route through the Johnstone Strait. When the Gulf of Alaska is cold, sockeye are distributed further south, make landfall on the west coast of Vancouver Island, and approach the Fraser River primarily via a southern route through the Juan De Fuca Strait (Figure 4.1). The percentage of sockeye returning around the northern end is called the Johnstone Strait Diversion (hereafter JSD). The JSD has a distinct interannual variability and is measured by records of salmon caught by fishermen each year. Research is being conducted to determine the factors that drive this variability.

Recent work in plotting ocean currents has provided scientists with a possible explanation for sockeye migration patterns. It has been speculated that the interannual variability of ocean surface circulation has impact on where the sockeye make landfall. A multi-institutional investigation has been initiated to examine the influences of currents, temperature and salinity on open ocean return migrations of sockeye salmon. Researchers are using the OSCURS (Ocean Surface Circulation Simulation) model [Ing88][Ing89]. Possible return migration paths of the sockeye are plotted by placing 174 simulated salmon at fixed locations in the open ocean (Figure 4.2). The simulated salmon use a common set of “rules” to find their way back to the British Columbia coast. Three separate experiments have been proposed. Salmon from each experiment react to different sets of environmental stimuli. These experiments are designed to test the following:

1. The influence of currents on compass-oriented migrations. Compass-oriented salmon take a single “look” before their migration run to determine the direction to the coast. They use a biological “compass” to swim in this fixed direction during their migration run, regardless of external forces (e.g., currents) that may shift their migration path
2. The influence of currents on bicoordinate-orientated migrations. Bicoordinate-oriented salmon make multiple “looks” during their migration run. Each “look” allows them to adjust for external forces by updating the direction they need to swim to reach the coast
3. The influence of currents, temperature and salinity on compass- and bicoordinate-orientated migrations



(a)



(b)

Figure 4.2: Examples of output from OSCURS program: (a) dots represent starting positions of 174 simulated salmon; (b) trailer beginning at each salmon's starting position tracks its path to the British Columbia coast

The simulations address the following questions regarding the ocean phase of sockeye return migrations:

- What is the range of sockeye swimming speeds (as opposed to migration rates)?
- What direction-finding mechanisms are used by sockeye (i.e., compass-orientation, plotting and/or bicoordinate-orientation)?
- How does the interannual variability of open ocean current, temperature and salinity fields impact the observed interannual variability of the point of landfall and timing of sockeye returns?

The model simulations of the first phase are complete. They focused on two years with significantly different JSDs, 1982 (22% JSD) and 1983 (80% JSD). Simulations of passive salmon for each of the two years confirmed that the magnitudes of the surface currents are sufficient to influence the migrating sockeye. The average net north-south current drifts were 2.4 and 5.1 kilometres per day, respectively. Over a migration period of sixty days, the difference in these drifts could account for sockeye landfall in 1983 of up to 150 kilometres further north than 1982. A shift in landfall of this distance might explain the large difference in JSDs.

Fifty-four simulations of compass-orientated sockeye were run. Each simulation “released” 174 sockeye, distributed over the Northeast Pacific sub-arctic habitat (Figure 4.2a). The following model parameters were used:

- 1982 and 1983 surface currents
- sockeye swim speeds of 20.8 cm/s (18 km/d), 34.7 cm/s (30 km/d) and 55.6 cm/s (48 km/d)
- sockeye swim directions of 90°, 112.5°, and 135°
- “release” times (simulated start dates) of May 1, June 1, and July 1

These studies illustrate the problems discussed in Chapter 2, namely display of multidimensional data in a two-dimensional environment, time animated displays, and displays that allow comparison of plots to quickly detect differences.

Standard Visualization Tools

We started our analysis of Oceanography's simulation data by building a set of "standard" visualization tools. These tools were written in an ad-hoc manner, without using results from research in preattentive processing. This helped us understand the Oceanography simulations and the type of data analysis researchers required. We also wanted to see what kind of visualization tools we could build using standard design techniques. The design process involved consulting with the end users and writing tools to analyse a single specific problem. Three different programs were written. The first compared latitude of landfall for two different years. The second compared salmon swim speeds for two different years. The third compared date of landfall for two different years.

An example of output from the latitude of landfall visualization tool is shown in Figure 4.3. The top and center images show latitude of landfall values for the two years under study, 1982 and 1983, respectively. The researchers in oceanography subdivided the range of possible landfall values into three regions: from 45° to 51° latitude, from 51° to 54° latitude, and from 54° to 60° latitude. Salmon making landfall below 45° or above 60° are not displayed. Each region is represented by a unique grey scale range. A salmon's latitude of landfall is shown by drawing a grey square at the salmon's starting position. The intensity of the square identifies both the region and the exact latitude where the salmon made landfall.

The bottom image shows the difference in latitude of landfall between 1982 and 1983. The range of possible values is subdivided into two regions. A salmon that landed farther north in 1983 versus 1982 is shown by drawing a light grey square at its starting position. A salmon that landed farther south in 1983 versus 1982 is shown by drawing a dark grey square at its starting

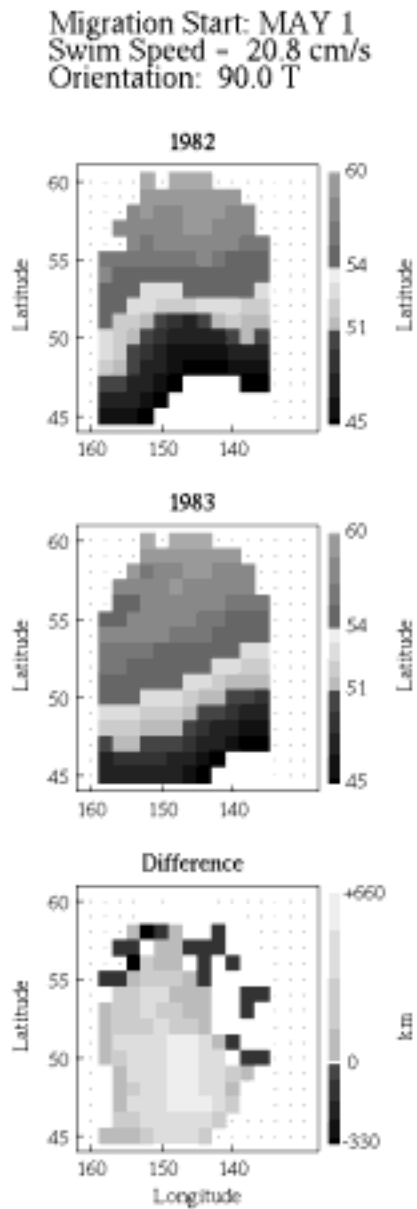


Figure 4.3: Example output from the landfall visualization tool. Top image shows latitude of landfall for 1982. Center image shows latitude of landfall for 1983. Bottom image shows difference in latitude of landfall between 1983 and 1982




Figure 4.4: Example output from the swim speed visualization tool. Top image shows difference in swim speed and migration speed for 1982. Center image show difference in swim speed and migration speed for 1983. Bottom image show difference in migration speed between 1983 and 1982

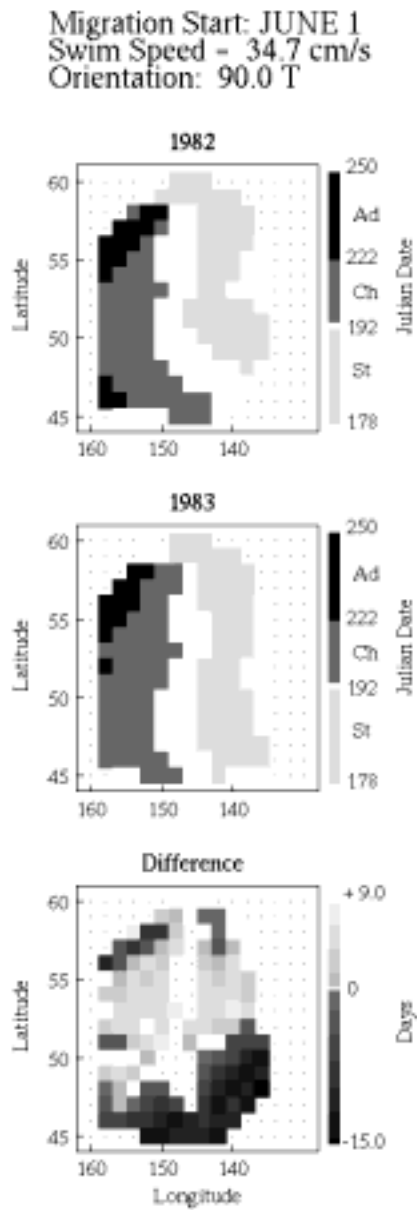


Figure 4.5: Example output from the date of landfall visualization tool. Top image shows stock classification for salmon in 1982, based on date of landfall. Center image shows stock classification for salmon in 1983. Bottom image shows difference in date of landfall between 1983 and 1982

position. Salmon with the same latitude of landfall in both 1982 and 1983 are represented by an empty square. Again, salmon that made landfall outside the range 45° to 60° in either of the two years are not displayed. The data in Figure 4.3 support the hypothesis that ocean currents affect salmon migration patterns. The majority of the squares in the bottom image are light grey, indicating the majority of salmon in 1983 landed farther north than they did in 1982. These results are similar to the real data, which shows a 1982 JSD of 22% and a 1983 JSD of 80%.

The other two programs visualize data in a similar way. The swim speed visualization tool compares a salmon's swim speed and its actual migration speed (Figure 4.4). Swim speed is a constant speed used by the salmon to swim towards the British Columbia coast. Migration speed is the speed the salmon makes after taking into account the effect of the current. The current may assist or resist the salmon's forward progress. The top and center images compare swim speed and migration speed for 1982 and 1983 respectively. A salmon with a migration speed faster than its swim speed is coloured light grey. A salmon with a migration speed slower than its swim speed is coloured dark grey. The bottom image compares the difference in migration speed between 1982 and 1983. Salmon that are faster are coloured light grey, while salmon that are slower are coloured dark grey.

The date of landfall visualization tool classifies salmon as belonging to one of three different stocks, Stikine, Chilko, or Adams (Figure 4.5). "Stock" refers to a specific area of the Fraser River where the salmon are born and where they migrate to spawn. Researchers know the time periods during which salmon from a given stock arrive at the mouth of the Fraser river. The visualization tools can assign a salmon to a given stock using its date of landfall. Salmon that arrived during the Stikine migration period are coloured light grey. Salmon that arrived during the Chilko migration period are coloured dark grey. Salmon that arrived during the Adams migration period are coloured black. Salmon that arrived during none of the above migration periods are represented by an empty square. The bottom image shows the difference in date of landfall between 1983 and 1982. Salmon that landed earlier in 1983 versus 1982 are coloured light grey. Salmon that landed later are coloured dark grey.

The visualization tools could be more useful if they had some way to “superimpose” all three frames of data into a single display. This would allow the researchers to search more easily for dependencies among the different variables. To design this type of tool, one must answer the question “How should the variables be combined and displayed simultaneously?” An ad-hoc assignment of features such as colour, shape, orientation, and intensity to individual data values may not result in a useful visualization tool. In fact, as discussed in the next chapter, it may result in a tool that inhibits the user’s ability to extract the desired information. Results from research in preattentive processing can be used to design a “preattentive” visualization tool. Such tools intelligently represent multivariate data using multiple preattentive features.

Chapter 5

Psychological Experiments

*“Errors, like straws, upon the surface float
He who would search for pearls must dive below”*

– John Dryden (1631-1700)

Through experimentation, we hope to determine whether or not research in preattentive processing can help design more useful and intuitive scientific visualization tools, as described in Chapter 3. DeSanctis gave a number of criteria that measure the usefulness of a graphical display: interpretation accuracy, problem comprehension, task performance, decision quality, speed of comprehension, decision speed, memory for information, and viewer preference. We use estimation error and reaction time to measure accuracy for three of these criteria. In our experiments, a “more useful” visualization tool will be one with an improved interpretation accuracy, task performance, or decision speed.

The experiments used data from the salmon migration simulations. They were limited to a specific set of preattentive features and a specific type of task. Due to the general nature of our data and task, we believe these experiments can be extended to a broad range of visualization problems. In general, we expected to address the following issues:

- are preattentive features “better” than standard visualization methods?

- in what situations are preattentive features useful or not useful?
- are preattentive visualization tools “intuitive” to the user, or do they require a large amount of initial explanation?
- does encoding separate data values with different preattentive features cause an interference effect for the given task due to feature preference?
- do our observations match what we would expect, given results from previous experiments in preattentive processing?

Task Selection

We consulted with researchers in Oceanography to identify a suitable task to use during the psychological experiments. We wanted to choose a task that was simple, but that still posed a question of interest to Oceanography. Part of this involved using results from the salmon migration simulations during the experiments. This allowed us to see how our visualization techniques performed with “real-world” data. We decided to ask subjects to estimate the number of simulated salmon that made landfall north of some fixed latitude. Relative landfall was encoded on a two-dimensional map of the open ocean at the spatial position the salmon started its migration run. A preattentive feature was used to represent position of landfall. For example, during one experiment salmon that landed north of the given latitude were coloured blue, while salmon that landed south were coloured red. Subjects were asked to estimate the percentage of blue elements in each display.

We wanted to see how preattentive features interfere with one another. Callaghan’s experiments showed a “feature preference” hierarchy for her texture boundary detection task [Cal84][Cal89]. We wanted to see if irrelevant preattentive features interfere with a subject’s estimation ability, similar to the way irrelevant preattentive features interfere with boundary detection ability. We decided to use the stream function for our “interference” variable. The

stream function is a scalar value that represents current direction and speed at a given position in the ocean, for a given instant in time. It is the potential function for the velocity field. Given the stream function $\phi(x, y)$, the x and y components of the current vector can be calculated as follows.

$$\begin{aligned} V_x &= \frac{\partial \phi}{\partial x} \\ V_y &= -\frac{\partial \phi}{\partial y} \end{aligned} \tag{5.1}$$

A stream function value was encoded at each spatial position where a salmon started a migration run. A preattentive feature was used to represent the stream function. Stream function is an unnecessary piece of information, because our estimation task asks only about latitude of landfall. We needed to ensure that stream function acted only as an interference value, independent of latitude of landfall. In fact, we suspect that stream function and landfall are correlated. Subjects might “learn” about this dependence during the experiment (which is exactly what researchers in Oceanography had to do). It would then be possible to use either landfall or stream function information to complete the estimation task. We would be unable to determine how different choices for our primary preattentive feature affect estimation ability. Subjects would simply gather information about the data value encoded with the “preferred” feature. If stream function and latitude of landfall are dependent, information about either one might be used to give a landfall estimation.

The latitude of landfall values from the migration simulations were modified to give us an even distribution of “northern salmon” (this is explained further below). These changes also provided the desired independence. Any dependence that might have existed was lost when we modified the landfall values. Note that when Oceanography studies the actual simulation data, they will be looking for a hypothesized correlation between latitude of landfall and the stream function field.

Open ocean current data was available for the years 1946 to 1990. The salmon migration

simulation was run for each of these years. The latitude position where each salmon made landfall was recorded. This allowed us to calculate an average latitude of landfall for a salmon, based on its starting position. We compared these average values to individual simulation results. For a given year, this allowed us to classify a salmon as landing farther north or south versus the average landfall data.

Each simulation ran for a predetermined number of simulated days. Stream function information for each day was recorded. This allowed us to calculate a set of average stream function values for each year from 1946 to 1990. As mentioned above, these values were encoded at each spatial position where a salmon started a migration run. In summary, our proposed task was as follows:

- our primary data value was whether a salmon lands north or south of the average latitude of landfall for the salmon's starting position
- the primary data value was encoded at the spatial position the salmon started its migration run, using preattentive feature A
- our secondary data value was the average stream function value at every spatial location where a salmon started a migration run
- the secondary data value was encoded using preattentive feature B
- subjects were asked to estimate the percentage of elements in the display with a given preattentive feature (i.e., the number of salmon landing north of the average latitude of landfall) to the nearest 10%
- estimation error (percentage of over- or under-estimation) and reaction time were used as a measure of subject performance



Figure 5.1: Examples of the rectangle orientations used during the experiment: (a) 0° rotation; (b) 60° rotation

Experiment Design

Through experimentation, we hoped to answer two specific sets of questions about preattentive features and their use in visualization tools:

- is it possible for subjects to provide a reasonable estimation of the relative number of elements in a display with a given preattentive feature? What features under what conditions allow this?
- how does encoding an “irrelevant” data dimension with a secondary preattentive feature interfere with a subject’s estimation ability? Which features interfere with one another and which do not?

We decided to examine two features, hue and orientation. These have been shown to be preattentive in various experiments by Julész [Jul83] and Triesman [Tri88]. Two unique line orientations were used: 0° rotation and 60° rotation (Figure 5.1). Two different hues were used, H_1 and H_2 . These hues were chosen from the Munsell colour space, and they satisfied the following properties:

1. The hues were isoluminant, that is, the perceived brightness of both hues were equal

2. The perceived difference between hues H_1 and H_2 were equal to the perceived difference between a rectangle rotated 0° and one rotated 60°

An explanation of how we determined these properties will be given later in this chapter. Our design allowed us to use oriented, coloured rectangles to represent latitude of landfall and stream function values at the starting position of each salmon.

Latitude of landfall and stream function values for each year were split into two subgroups before the experiment began. Latitude of landfall was divided into a subgroup of values south of the average latitude of landfall position, and a subgroup of values equal to or north of the average latitude of landfall position. The stream function was divided into two equally sized subgroups. All the values in the first subgroup were less than or equal to all the values in the second subgroup. We could then use either of our two preattentive features to encode landfall and stream function values.

The experiment was divided into four subsections or “blocks”, B_1 , B_2 , B_3 , and B_4 . The primary and secondary data value varied within each block, as did the primary and secondary preattentive feature. This gave us the following:

1. Primary data value was latitude of landfall, represented by hue. Secondary data value was stream function, represented by orientation (Figure 5.2)
2. Primary data value was latitude of landfall, represented by orientation. Secondary data value was stream function, represented by hue
3. Primary data value was stream function, represented by hue. Secondary data value was latitude of landfall, represented by orientation
4. Primary data value was stream function, represented by orientation. Secondary data value was latitude of landfall, represented by hue

Each block was further divided into two control subsections and one experiment subsection (Figure 5.3). The control subsections used a fixed value for the secondary feature. This allowed




Figure 5.2: Example of a display from block B₁, latitude of landfall represented by hue, stream function represented by orientation

us to see how a real interference feature affected estimation error and reaction time. As an example, block B_1 was subdivided as follows:

- Control subsection 1 consisted of 36 trials. Landfall was encoded using hue. Stream function was ignored, and every rectangle was oriented 0°
- Control subsection 2 consisted of 36 trials. Landfall was encoded using hue. Stream function was ignored, and every rectangle was oriented 60°
- The experiment subsection consisted of 72 trials. Landfall was encoded using hue. Stream function was encoded using orientation

The 72 trials used in control subsections were almost the same as the 72 trials used in the experiment subsection. The only difference was whether stream function was ignored or used to determine orientation. In fact, there were only 36 unique trials. Each trial was shown four times during the experiment to obtain the desired number of trials. Each of the 36 unique trials contained a certain percentage of target elements. We wanted an equal number of trials for each percentage value. For example, in block B_1 , the target element was a salmon that lands north of the average landfall position. We wanted 4 trials in which 5-15% of the salmon were “northern”, 4 trials in which 15-25% of the salmon were “northern”, and so on up to 85-95%. These percentage values are called intervals, and are numbered from 1 (5-15% northern) to 9 (85-95% northern). The original simulation data was modified to give us the desired distribution.

Blocks B_2 , B_3 , and B_4 were subdivided in a similar manner. Each subject completed either blocks B_1 and B_3 (testing hue), or blocks B_2 and B_4 (testing orientation), for a total of 288 trials. The 144 trials within each block were presented in a random order to the subject.

The only difference between blocks B_1 and B_3 and blocks B_2 and B_4 was the primary data value. Blocks B_1 and B_2 used landfall as the primary data value. Blocks B_3 and B_4 used stream function as the primary data value. Stream function and latitude of landfall were subdivided




Figure 5.3: Overview of the experiment design

Block Name	Block Type	Primary Value	Primary Feature	Secondary Value	Secondary Feature	Total Trials
Sub 1-1	control	landfall	hue	—	0° orientation	36
Sub 1-2	control	landfall	hue	—	60° orientation	36
Sub 1-3	experiment	landfall	hue	stream	orientation	72
Sub 2-1	control	landfall	orientation	—	red	36
Sub 2-2	control	landfall	orientation	—	blue	36
Sub 2-3	experiment	landfall	orientation	stream	hue	72
Sub 3-1	control	stream	hue	—	0° orientation	36
Sub 3-2	control	stream	hue	—	60° orientation	36
Sub 3-3	experiment	stream	hue	landfall	orientation	72
Sub 4-1	control	stream	orientation	—	red	36
Sub 4-2	control	stream	orientation	—	blue	36
Sub 4-3	experiment	stream	orientation	landfall	hue	72

Table 5.1: Information for each subsection in the experiment, including name, type, primary data value, primary preattentive feature, secondary data value, secondary preattentive feature, and total number of trials

for blocks B₃ and B₄ as follows. An average stream function value was calculated using data for the years 1946 to 1990. A particular stream function value could then be classified as lower than, equal to, or greater than the average value. Latitude of landfall was split into two equal groups for each year 1946 to 1990. The primary and secondary data values were subdivided in a way identical to blocks B₁ and B₂.

The above variation was designed to give us two different types of spatial patterns. The “question” asked in blocks B₃ and B₄ is: what percentage of stream function values are equal to or greater than the average value? Although this is not a question of interest for Oceanography, it was asked to help validate results from the psychological experiments. Landfall values tend to subdivide into two separate “groups” of elements. Stream function values tend to subdivide into a concentric ring-like pattern. We wanted to compare estimation error and reaction time between blocks B₁ and B₃ and blocks B₂ and B₄. A difference would suggest estimation ability depends, at least in part, on the type of spatial pattern presented to the subject.

Colour Selection

The hues used during the psychological experiments were chosen from the Munsell colour space. This colour space was originally proposed by Albert H. Munsell in 1898. It was later revised by the Optical Society of America in 1943 to more closely approximate Munsell’s desire for a functional and perceptually balanced colour system. A colour from the Munsell colour space is specified using the three “dimensions”, hue, chroma, and value.

In Munsell space, hue refers to some uniquely identifiable colour, or as Munsell suggested, “the quality by which we distinguish one colour from another, as a red from a yellow, a green, a blue, or a purple”¹. Hue is represented by a circular band divided into ten sections. Munsell named these sections 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 for short). Each section

¹*Munsell: A Grammar of Colour*. New York, New York: Van Nostrand Rienhold Company, 1969, pg. 18




Figure 5.4: Munsell colour space, showing it's three dimensions hue, value, and chroma

can be further divided into ten subsections if finer divisions of hue are needed. A number preceding the hue name is used to define the subsection (for example, 5R or 7BG).

Value refers to a colour's lightness or darkness. Munsell defined value as "the quality by which we distinguish a light colour from a dark one"². Value is divided into nine sections, numbered 1 through 9. Dark colours have a low value, while lighter colours have a higher value.

Chroma defines a colour's strength or weakness. Chroma is measured in numbered steps starting at 1. Weak colours have low chroma values. Strong colours have high chroma values. Greys are colours with a chroma value of zero. The maximum possible chroma value depends on the hue and value being used.

²Ibid, pg. 20

A visual representation of the Munsell colour space is shown in Figure 5.4³. The circular band represents hue. The pole running through the center of the colour space represents value. Colours with increasing chroma radiate outwards from the value pole. A Munsell colour is specified by writing “hue value/chroma”. For example, R 6/6 would be a relatively strong red. BG 9/2 would be a weak cyan.

We chose the Munsell space because it provides a number of desirable properties. Munsell originally designed his colour space to be used by artists. One feature he tried to incorporate into his system was perceptual “balance”. Hues directly opposite one another will be balanced, provided their value and chroma are equal. Thus, BG 5/5 is perceptually balanced with R 5/5, and Y 2/3 is balanced with PB 2/3. Opposite hues with different values and chromas can also be balanced by varying the amount of each colour used within a given area. Given two Munsell colours $H_1 V_1/C_1$ and $H_2 V_2/C_2$, we need $V_2 C_2$ parts of hue H_1 and $V_1 C_1$ parts of hue H_2 . For example, colours R 5/10 and BG 5/5 can be balanced by using BG 5/5 in two-thirds of the area, and R 5/10 in one-third of the area. As we would expect, the stronger chroma and higher value take up less of the total area than the weaker chroma and lower value.

A second and perhaps more important property is that Munsell colours with the same value are isoluminant. Thus, colours R 5/5, G 5/6, B 5/3, and any other colours with value 5 are all perceived as having equal luminance. This property was provided when the Munsell colour table was revised in 1943.

Munsell colours must be converted into RGB triples to be displayed on a computer monitor. The first step in this process is calculating a CIE to RGB conversion matrix for the monitor being used. Given the CIE chromaticity values $(x_r, y_r), (x_g, y_g), (x_b, y_b)$ for the monitor phosphors, and the luminance of the monitor’s maximum-brightness red, green, and blue (Y_r, Y_g, Y_b) we can compute the following.

³Ibid, pg. 23

$$\begin{aligned}
 z_r &= 1 - x_r - y_r \\
 z_g &= 1 - x_g - y_g \\
 z_b &= 1 - x_b - y_b
 \end{aligned}
 \tag{5.2}$$

$$\begin{aligned}
 C_r &= \frac{Y_r}{y_r} \\
 C_g &= \frac{Y_g}{y_g} \\
 C_b &= \frac{Y_b}{y_b}
 \end{aligned}
 \tag{5.3}$$

A colour is usually specified in CIE colour space with a triple (x, y, Y) . These values are used to obtain X and Z as shown in Equation 5.4 below. The values X , Y and Z are then inserted into Equation 5.5 to obtain the monitor R, G, B values for the given colour [Fol90]. This process assumes the intensity steps produced by the video hardware are linear.

$$\begin{aligned}
 X &= \frac{x}{y} Y \\
 Z &= \frac{1 - x - y}{y} Y
 \end{aligned}
 \tag{5.4}$$

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} x_r C_r & x_g C_g & x_b C_b \\ y_r C_r & y_g C_g & y_b C_b \\ z_r C_r & z_g C_g & z_b C_b \end{bmatrix}^{-1} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}
 \tag{5.5}$$

Our experiments were run on a Apple Macintosh II microcomputer using a software package written by Rensink and Enns [Enn91]. This software was specifically designed to run preattentive psychology experiments. The microcomputer was equipped with an Apple RGB 13-inch colour display and a Mac II High-Resolution Video Card. It was capable of displaying 256 colours simultaneously. The CIE chromaticities for the monitor phosphors were supplied by the manufacturer. A spot photometer was used to measure the luminance values for the mon-

itor's maximum-intensity red, green, and blue. This gave us the following values, which were used to obtain the conversion matrix shown in Equation 5.8.

$$\begin{aligned}(x_r, y_r, z_r) &= (0.625, 0.340, 0.035) \\(x_g, y_g, z_g) &= (0.280, 0.595, 0.125) \\(x_b, y_b, z_b) &= (0.155, 0.070, 0.775)\end{aligned}\tag{5.6}$$

$$\begin{aligned}Y_r &= 5.5 \\Y_g &= 16.6 \\Y_b &= 2.8\end{aligned}\tag{5.7}$$

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 0.1313 & -0.0574 & -0.0211 \\ -0.0439 & 0.0806 & 0.0015 \\ 0.0025 & -0.0080 & 0.0325 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}\tag{5.8}$$

Tables available in Wyszecki and Stiles give CIE chromaticity values for a large number of Munsell colours [Wys82]. Once we have the CIE values for a Munsell colour, we can use Equation 5.5 to convert them into RGB triples. Not all the Munsell colours can be displayed by the monitor. Many of them fall outside the monitor's colour gamut. In particular, the number of displayable green, blue-green, and blue hues is quite limited. Table 5.2 shows the number of different chroma values available for each of the primary Munsell hue and value combinations.

We know how to convert Munsell colours to monitor RGB colour space. We must still choose two Munsell colours for our experiment. Recall that we required the following properties of our two hues.

1. The hues will be isoluminant, that is, the perceived brightness of both hues will be equal
2. The perceived difference between hues H_1 and H_2 will be equal to the perceived difference between a rectangle rotated 0° and one rotated 60°

Hue	Value								
	1	2	3	4	5	6	7	8	9
5 R	2/5	4/7	5/8	6/9	8/10	9/9	7/7	5/5	3/3
5 YR	1/2	2/3	3/4	4/6	5/7	5/9	6/10	6/7	3/3
5 Y	1/1	1/2	2/3	2/4	3/6	3/7	4/8	4/9	5/10
5 GY	1/2	2/3	2/4	3/5	3/6	4/7	4/8	5/10	5/9
5 G	1/4	2/8	3/11	4/13	4/14	5/14	6/13	8/11	6/6
5 BG	1/3	2/6	2/9	3/10	3/11	4/10	5/10	5/8	5/5
5 B	1/3	2/5	2/6	3/7	4/8	4/8	5/7	4/4	2/2
5 PB	2/5	3/8	4/9	6/10	7/9	7/7	5/5	3/3	1/1
5 P	3/11	5/14	7/16	9/16	11/14	10/10	7/7	5/5	2/2
5 RP	3/7	4/9	5/10	7/11	8/12	9/11	9/9	6/6	3/3

Table 5.2: Number of displayable colours for all primary Munsell hue and value combinations (Apple RGB 13-inch colour display), in the format “displayable chroma/total chroma”

Requirement 1 was satisfied by ensuring both hues had the same value in Munsell space. We chose Munsell value 7, because that slice through Munsell space provided the largest range of displayable colours for a variety of different hues.

Requirement 2 was satisfied by running a set of preliminary experiments. We started with a simple target detection task. Users were asked to detect the presence or absence of a rectangle rotated 60° in a field of distractor rectangles rotated 0° . Both the target and distractor rectangles were coloured 5R7/8. The experiment consisted of 30 trials. 15 of the 30 trials were randomly chosen to contain the target. Each trial consisted of 36 rectangles (including the target, if present), drawn in random positions on the screen. The average reaction time for detection was computed from the trials in which the user responded correctly.

After the first experiment, the target and distractors were changed. The target was a rectangle coloured 10RP7/8. The distractors were rectangles coloured 5R7/8. Note that the target is a single counter-clockwise “hue step” from the distractors in Munsell space. Both

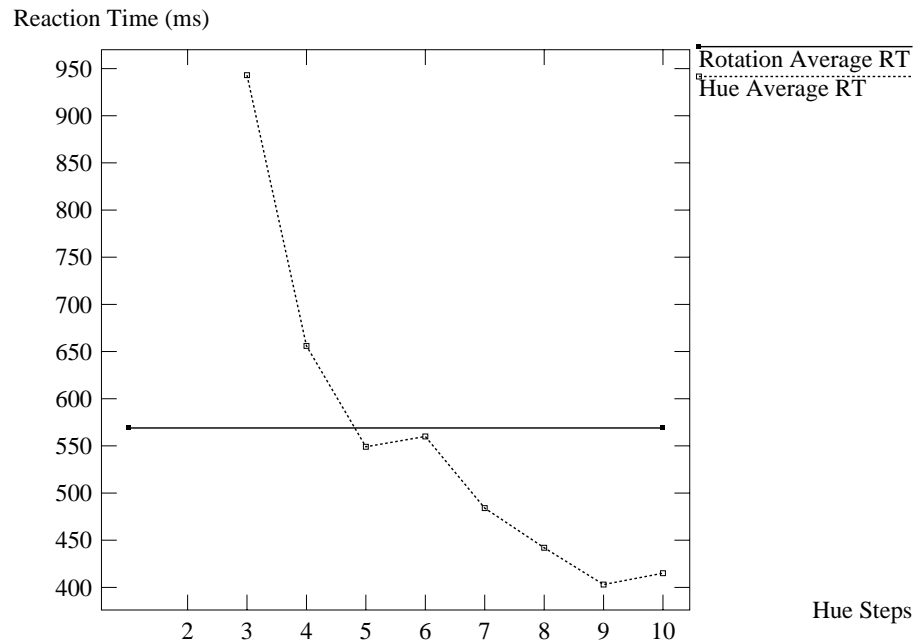


Figure 5.5: Results from hue and rotation difference experiments. The solid line shows reaction time for rotation discrimination. Points along the dashed line show reaction time for discriminating the given Munsell hue

the target and distractor rectangles were rotated 0° . The experiment consisted of 30 trials. 15 of the 30 trials were randomly chosen to contain the target. Each trial consisted of 36 rectangles (including the target, if present), drawn in random positions on the screen. The average reaction time for detection was computed from the trials in which the user responded correctly.

The hues used for the target and distractors during the second experiment were very similar. Because of this, the average reaction time for the second experiment was higher than the average reaction time for the first experiment. Additional experiments were run as follows.

- the target was moved another counter-clockwise “hue step” away from the distractors (i.e., 5RP 7/8, 10P 7/8, and so on)

- the second experiment was re-run, and average reaction time was computed
- this process continued until an average reaction time equal to or below the average reaction time of the first experiment was obtained

This process provided two isoluminant hues H_1 and H_2 with a perceived difference equal to that of a 60° rotation, where perceived difference was measured by reaction time in the target detection experiment.

Figure 5.5 shows a graph of the combined results of all subjects. The solid horizontal line shows the average reaction time for detecting the rotated rectangle, 569 milliseconds. The dashed line shows the average reaction times for detecting each coloured rectangle, starting with 5P 7/8, which is three hue steps from the distractors, and ending with 5BG 7/8, which is ten hue steps from the distractors. The exact reaction times are provided in Table 11.5. At five hue steps, or 10PB 7/8, the average reaction time for hue detection was less than the average reaction time for target detection. We chose to use the hues 5R 7/8 and 5PB 7/8 during the experiment. These hues are 6 hue steps from one another. We assumed the equality in perceived difference would transfer from the target detection task to the estimation task.

Experiment Procedure

Twelve subjects were used during the experiment, nine males and three females, all students or staff at the University of British Columbia. All subjects had normal or corrected vision and none were known to be colour blind. Experiments were run in the Department of Psychology's vision laboratory, using a Macintosh II microcomputer equipped with a 13-inch RGB monitor and video hardware capable of displaying 256 colours simultaneously. In general, it took subjects between 50 and 80 minutes to finish the entire experiment. Subjects were paid an honorarium of \$10 for their participation in the experiment.

Each subject completed either blocks B_1 and B_3 (blocks using hue as the primary feature)

or blocks B_2 and B_4 (blocks using orientation as the primary feature) for a total of 288 trials. At the beginning of the experiment, subjects were shown a sample display frame. The experiment procedure and task were explained to the subject. For example, subjects completing the hue blocks were told:

During the experiment, you will be shown a number of displays or “trials” similar to this one. Notice the trial is made up of rectangles, some of which are coloured blue, some of which are coloured red. Your task will be to estimate, to the nearest 10%, the number of rectangles that are coloured blue. These trials will be shown in the following way. First, the screen will blank for a short period of time. Just before the trial is shown, a focus circle will be displayed, so you can prepare yourself. The trial will be shown for a fixed period of time. The screen will blank, and the computer will wait for you to type in your estimation. If your response is correct, a plus sign will be shown. If your response is incorrect, a minus sign will be shown. The computer will then proceed to the next trial. You can take as much time as you like to make your estimation. However, in general, the longer you wait before making your estimation, the less accurate your answer will be

Subjects were then shown how to enter their estimation. This was done by typing a digit on the keyboard between 1 and 9, which corresponded to the percentage of rectangles they estimated contained the target feature: 5-15%, 15-25%, and so on up to 85-95%. Subjects were told no trial would contain less than 5% or more than 95% of the target rectangles.

Subjects started with a set of practice trials called “Fixed Practice”. This consisted of nine trials, one for each of the nine possible intervals. In one trial 10% of the rectangles were targets, in another 20% were targets, and so on up to 90%. The practice trials were designed to calibrate the subjects and to give them an idea of the speed of the trials and the experiment. Trials were displayed one after another to the subject in the manner described above. If subjects estimated correctly, they moved immediately to the next trial. If they estimated incorrectly, the trial was redisplayed, and they were told the correct answer.

Next, subjects completed a second set of practice trials called “Random Practice”. This set

of trials consisted of 18 trials, two for each of the nine possible intervals. Trials were displayed in a random order to the subject. This experiment was designed to run exactly like a real experiment block. Trials in which the subject estimated incorrectly were not redisplayed and subjects were not told the correct answer, although they were told whether their estimation was right or wrong.

Finally, subjects completed the two experiment blocks B_1 and B_3 or B_2 and B_4 . Each block consisted of 72 control trials and 72 experiment trials. The 144 trials from each block were presented to the subject in a random order. Subjects were provided with an opportunity to rest after every 48 trials. Data from all four phases was saved for later analysis.

Chapter 6

Analysis of Results

*“King, father, royal Dane, O! answer me
Let me not burst in ignorance; but tell . . .”*

– Hamlet, *Shakespeare*

A variety of information was obtained during the psychological experiments. Four data files were produced for each subject: a fixed practice file, a random practice file, a file for the first experiment block, and a file for the second experiment block. Data for each trial included a record of the subject’s estimation, the correct estimation, response time, and trial type (control or experiment) for the experiment blocks. Estimation error is the absolute value of the difference between the subject’s estimation and the correct estimation for a given trial.

Some trials showed large errors, usually when the correct estimation was in the lower 3 or upper 3 intervals. We hypothesize this was caused by subjects accidentally estimating the wrong feature, either the number of red rectangles instead of blue rectangles or the number of rectangles rotated 0° instead of 60° . This was supported to some extent by subjects who mentioned they had made this type of mistake. We decided to modify the data by complementing subject estimations that displayed an unusually large error. Estimations were corrected as follows:

- for intervals 1 (5-15% targets), 2 (15-25% targets), and 3 (25-35% targets), estimations greater than 5 were “flipped” about 5, that is, 6 became 4, 7 became 3, 8 became 2, and 9 became 1
- for intervals 7 (65-75% targets), 8 (75-85% targets), and 9 (85-95% targets), estimations less than 5 were “flipped” about 5, that is, 4 became 6, 3 became 7, 2 became 8, and 1 became 9

In total, less than 1% (21 of 2,304) of the trials were modified. This data formed the basis for the analysis described in this chapter. We performed a number of t -tests and analysis of variance (ANOVA) F -tests while investigating our data. A t -test tests two independent samples X_1 and X_2 to see if their means are equal within some confidence interval $1 - \alpha$. Equality of means is called the null hypothesis $H_o : \mu_1 = \mu_2$. α represents the probability of rejecting H_o when it is in fact true. For example, an $\alpha = 0.10$ gives a 10% probability of error. The t value for a given α and number of degrees of freedom v is denoted $(1-\frac{\alpha}{2})t_v$.

An ANOVA F -test compares three or more independent samples X_1, \dots, X_n to see if their means are equal within some confidence interval $1 - \alpha$. Equality of means is called the null hypothesis $H_o : \mu_1 = \dots = \mu_n$. α represents the probability of rejecting H_o when it is in fact true. The F value for a given α and number of degrees of freedom v is denoted $(1-\alpha)F_v$. The analysis that follows addresses the following questions:

- can rapid and accurate estimation be performed using each of hue and orientation?
- is there evidence of a subject preference for either hue or orientation for the estimation task?
- is there evidence of a subject preference for the type of data being displayed during the estimation task?
- is there evidence that orientation interferes with a subject’s ability to perform hue estimation?

- is there evidence that hue interferes with the subject's ability to perform orientation estimation?

A t -test assumes X_1 and X_2 are independent and normally distributed with means μ_1 and μ_2 and the same variance σ^2 . Similarly, an F -test assumes all X_i are independent and normally distributed with means μ_i and the same variance σ^2 . However, research by Glass and others [Gla84] has shown that the violation of normality has little effect on t -test robustness. Also, if the size of the samples is equal, that is, if $n_1 = n_2$, violation of homogeneity-of-variance has little effect on t -test robustness. These assumptions extend to ANOVA F -test results. Because our data meets these criteria, we did not test for normality or equal variance during our analysis.

Estimation Ability

The first question we addressed was whether subjects were able to perform accurate estimation in a 450 millisecond exposure duration. Figures 10.1–10.12 in Appendix A show graphs of combined subject data for the control and experiment subsections of blocks B₁, B₂, B₃, and B₄. Each graph plots average subject response \bar{V} for each interval, standard deviation of subject response $\sigma(V)$ for each interval, and standard deviation of subject estimation error $\sigma(e)$ for each interval. Tables 10.1–10.4 give exact values for these measurements, as well as average subject estimation error for each interval \bar{e} .

The results show that accurate estimation was possible during the experiment for all four blocks. In the experiment subsections the total \bar{e} ranged from a low of 0.54 in block B₂ to a high of 0.65 in block B₁. $\sigma(e)$ was below 1.0 in all four blocks. This indicates that subject responses were clustered close to the correct estimate. Results from the two control subsections show similar trends.

Figure 6.1 graphs \bar{e} across all nine intervals for the experiment subsection of blocks B₁, B₂, B₃, and B₄. \bar{e} for each block appears symmetric, with a maximum around interval 5 or 6,




Figure 6.1: Graph of average error values for the experiment subsection of all four blocks B_1 , B_2 , B_3 , and B_4

and a minimum at interval 1 or 9. It appears subjects had more trouble estimating trials from intervals 4, 5, and 6, and less trouble estimating trials from intervals 1, 2, 8, and 9. However, this is a well known phenomena called the “end effect”. Because subjects have more freedom of choice at intervals 4, 5, and 6, their observed estimation error is higher than at intervals 1 and 9, where there is only one direction of freedom available. This trend is also due in part to the fact that we “flipped” subject estimations from intervals 1, 2, 3, 7, 8, and 9 that displayed an unusually large error.

Feature Preference

A point of interest was whether a subject's estimation ability differed depending on the feature being estimated. A t -test was computed to see if estimation error mean was equal across primary features for both the control and experiment subsections. Blocks B_1 and B_3 were combined to form B_c . This super-block contained all trials where hue was the primary preattentive feature. Blocks B_2 and B_4 were combined to form B_o . This super-block contained all trials where orientation was the primary preattentive feature. The data displayed during B_c 's trials is the same data displayed during B_o 's trials. The only differences are the features used to encode the primary and secondary data values. The t -test compared means of the control and experiment subsection's estimation error.

Subsection	n_1	n_2	v	t
Control 1	432	432	862	0.36
Control 2	432	432	862	1.43
Experiment	864	864	1726	0.45

(a)

Subsection	n_1	n_2	v	t
Control 1	432	432	862	2.06
Control 2	432	432	862	1.73
Experiment	864	864	1726	1.84

(b)

Table 6.1: t -test results for estimation error rates from: (a) hue and orientation trials; (b) landfall and stream function trials

Table 6.1a shows the subsection, the number of hue trials n_1 , the number of orientation trials n_2 , the degrees of freedom v , and the t -value t . The control t -values are less than $0.975t_{862} = 1.962$ and the experiment t -value is less than $0.975t_{1726} = 1.960$. Therefore, there appears to be no feature preference for the estimation task. Any difference in means is probably due to sampling

error, and not the choice of hue or orientation as a primary preattentive feature. We did not expect to observe a feature preference, because we calibrated the perceived difference between our two hues and our two orientations to be equal before the experiment.

Data Type Preference

It is possible that the spatial distribution of the data affects a subject's estimation ability. We used two different data sources during the experiment, latitude of landfall and stream function values. A difference in estimation error mean across data types would indicate estimation ability depends, at least in part, on the spatial distribution of the data being displayed. Blocks B_1 and B_2 were combined to form B_l . This super-block contained all trials where latitude of landfall was the primary data value. Blocks B_3 and B_4 were combined to form B_s . This super-block contained all trials where stream function was the primary data value. Both B_l and B_s have 864 trials where hue was the primary preattentive feature, and 864 trials where orientation was the primary preattentive feature. The main difference between the two super-blocks is the underlying data being displayed. The t -test compared means of the control and experiment subsection's estimation error.

Table 6.1b shows the subsection, the number of landfall trials n_1 , the number of stream function trials n_2 , the degrees of freedom v , and the t -values t . Control subsection 1's t -value is greater than $_{0.975}t_{862} = 1.962$. This suggests that data type did have an effect on estimation error in control subsection 1. Control subsection 2's t -value is less than 1.962, but it does fall between $_{0.90}t_{862} = 1.283 < p < _{0.975}t_{862}$. Similarly, the experiment subsection's t -value falls between $_{0.90}t_{1726} = 1.282 < p < _{0.975}t_{1726} = 1.960$. The t -test results indicate the possibility of a data type influence on estimation error. With an α value of 0.10, we would not be able to reject the null hypothesis, that data type does not affect estimation error, in any of the subsections.

Feature Interference

One question of interest was whether encoding an irrelevant data value with a secondary preattentive feature affected a subject's estimation ability. We began by checking to see if orientation interfered with a subject's ability to estimate using hue. t -tests were computed to compare estimation error mean across control and experiment subsections for blocks B_1 and B_3 , the blocks that used hue as their primary preattentive feature.

Subsection	n_1	n_2	v	t
B_1	432	432	862	0.03
B_3	432	432	862	0.21

(a)

Subsection	n_1	n_2	v	t
B_2	432	432	862	0.23
B_4	432	432	862	1.15

(b)

Table 6.2: t -test results for estimation error rates from: (a) control and experiment hue trials; (b) control and experiment orientation trials

Table 6.2a shows the block, the number of control trials n_1 , the number of experiment trials n_2 , the degrees of freedom v , and the t -value t . The t -values for both blocks are less than $0.975t_{862} = 1.962$. Therefore, there appears to be no interference due to encoding of an irrelevant data value using orientation. Any difference in means is probably due to sampling error.

We continued to investigate interference by checking to see if hue interfered with a subject's ability to estimate using orientation. t -tests were computed to compare estimation error mean across control and experiment subsections for blocks B_2 and B_4 , the blocks that used orientation as their primary preattentive feature.

Table 6.2b shows the block, the number of control trials n_1 , the number of experiment

trials n_2 , the degrees of freedom v , and the t -value t . The t -values for both blocks are less than $0.975t_{862} = 1.962$. Therefore, there appears to be no interference due to encoding of an irrelevant data value using hue. Any difference in means is probably due to sampling error.

Feature and Data Type Interaction

Previous analysis has shown no subject preference for primary preattentive feature, but a possible subject preference for the spatial distribution of primary data values. We wanted to see if feature and data type interacted with one another. Suppose subjects prefer one type of spatial distribution over another. Is the improvement in estimation ability different depending on our primary preattentive feature? That is, does estimation using hue receive a larger improvement from the preferred spatial distribution, compared to estimation using orientation? Perhaps orientation receives a larger improvement versus hue. A 2-factor ANOVA F -test was computed to check for feature and data type interaction. A 2-factor ANOVA allows comparison of means and interaction of two independent data values. We compared the two independent variables feature and data type, each of which has two levels, hue and orientation and landfall and stream function, respectively. These four groups of data correspond to blocks B_1 , B_2 , B_3 , and B_4 .

Subsection	J	v_b	v_w	F_f	F_d	F_i
Control 1	2	8	860	0.22	7.03	0.46
Control 2	2	8	860	3.65	5.34	0.01
Experiment	2	8	1724	0.35	5.83	0.89

Table 6.3: F -test results of estimation error rates across feature and data type

Table 6.3 shows the subsection, number of levels J , between level degrees of freedom v_b , within level degrees of freedom v_w , and F -values F_f , F_d , and F_i . F_f is the F -value for estimation error mean across feature, F_d is the F -value for estimation error mean across data type, and

F_i is the F -value for feature and data type interaction. The F_f values for all subsections are less than ${}_{0.95}F_{860} = 3.85$, and the F_d values for all subsections are greater than ${}_{0.95}F_{860}$. This confirms our previous analysis, namely feature type does not affect estimation error, but data type may affect estimation error.

The interaction F -values F_i are all less than ${}_{0.95}F_{1724} = 3.84$. Therefore, no interaction seems to be present. Neither feature gains a greater improvement in estimation error when using a more preferred data type.

To conclude, we have addressed and statistically answered all the questions posed at the beginning of the chapter. It has been demonstrated that the data are consistent with the following conclusions:

- rapid and accurate estimation can be performed using either hue or orientation
- there is no evidence of a subject preference for either hue or orientation during the estimation task for the particular hue and orientation values used
- there is evidence of a subject preference for the underlying data being displayed during the estimation task
- there is no evidence that orientation interferes with a subject's ability to perform hue estimation
- there is no evidence that hue interferes with a subject's ability to perform orientation estimation
- there is no evidence of interaction between primary preattentive feature and the underlying data being displayed

These conclusions apply to data displayed for an exposure duration of 450 milliseconds. Because we already have robust estimation ability, large increases in exposure duration probably would not provide significant improvement in estimation ability. A more interesting direction of

investigation would be to decrease the exposure duration. This would allow us to examine two important questions. First, at what exposure duration are subjects no longer able to perform robust estimation? At some exposure duration below 450 milliseconds we expect subjects to be unable to give accurate estimations. Second, do any interference effects begin to appear at lower exposure durations? For example, we found that hue did not interfere with estimation of orientation at a 450 milliseconds exposure duration. It may be that an interference effect does exist, but 450 milliseconds gives subjects enough time to overcome this effect. If this is true, the interference should appear at lower exposure durations. Feature preference may also be dependent on exposure duration. We began investigating these possibilities by running a set of informal experiments described below.

Exposure Duration Experiments

We conducted a set of post-experiments to obtain information on how exposure duration affects the estimation task. 90 control and 90 experiment trials from block B_1 were used during the experiment. Exposure duration for each trial varied among five possible values: 15 milliseconds, 45 milliseconds, 105 milliseconds, 195 milliseconds, and 450 milliseconds. Presentation of trials differed somewhat from the previous experiment. A “mask” of randomly oriented grey rectangles was displayed for 105 milliseconds immediately following each trial. This was designed to remove any “after-image” of the trial that may have been present in the subject’s visual short term memory. In summary, trials were presented in the following way:

- a blank screen was displayed for 195 milliseconds
- a focus circle was displayed for 105 milliseconds
- the trial was displayed for its exposure duration (one of 15, 45, 105, 195, or 450 milliseconds)
- a mask of randomly oriented grey rectangles was displayed for 105 milliseconds

- the screen blanked, and subjects were allowed to enter their estimation

Because trials came from block B_1 , our primary data value was latitude of landfall, represented by hue, and our secondary value was stream function, represented by orientation. Subjects estimated the number of blue rectangles in each trial. As before, an equal number of trials for each interval was used. There were 10 control and 10 experiment trials where 5-15% of the rectangles were blue, 10 control and 10 experiment trials were 15-25% of the rectangles were blue, and so on up to 85-95% for a total of 180 trials. Trials at each interval were split evenly among the five exposure durations. For example, there were 2 control and 2 experiment trials with an exposure duration of 15 milliseconds at each interval (5-15%, 15-25% and so on up to 85-95% for a total of 36 trials). Trials were presented to the subject in a random order so the various exposure durations were intermixed.

Exposure	Control 1			Control 2			Experiment		
	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$
15 ms	45	1.93	2.74	45	2.15	2.73	90	1.52	2.07
45 ms	45	1.06	1.46	45	1.44	1.94	90	1.13	1.51
105 ms	45	0.91	1.51	45	0.80	1.16	90	0.75	1.17
195 ms	45	0.71	1.44	45	0.71	1.06	90	0.76	1.10
450 ms	45	0.60	1.01	45	0.80	1.12	90	0.57	0.91

Table 6.4: Average estimation error and standard deviation of estimation error for control and experiment trials at each exposure duration

Analysis of data from the previous experiment showed estimation was accurate at every interval. Because of this, we combined trials with a given exposure duration into a single block of data. For example, trials that were displayed for 105 milliseconds formed a single group of 2 control and 2 experiment trials from each interval for a total of 18 control and 18 experiment trials. We plotted average estimation error versus exposure duration to see if estimation ability was affected by display time. Figure 6.2 shows the graph of average estimation error versus exposure duration for experiment trials. Table 6.4 shows the number of samples

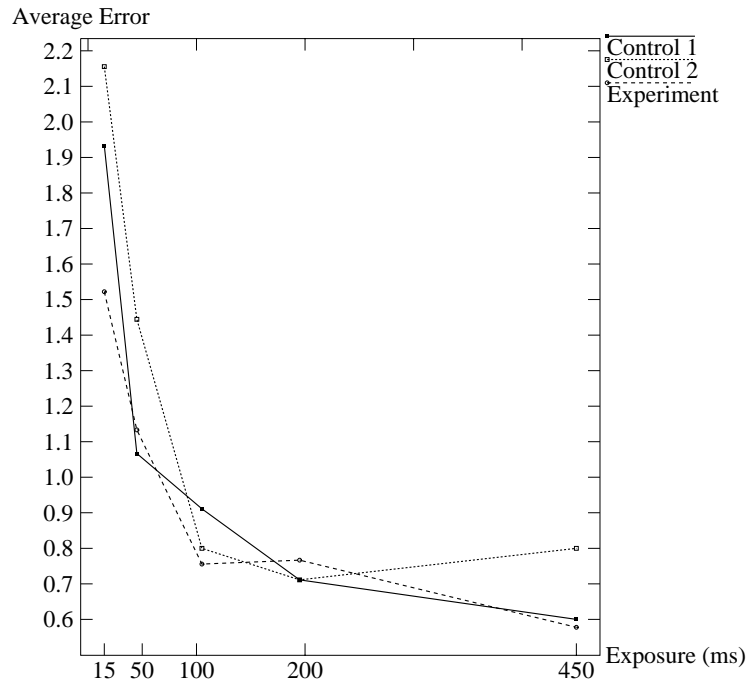


Figure 6.2: Graph of average error across exposure duration for combined results from exposure duration experiment

n , average estimation error \bar{e} , and standard deviation of estimation error $\sigma(e)$ for the control and experiment trials.

Average estimation error and standard deviation of error seem to be reasonably stable, even down to 105 milliseconds. Below that duration both values increased rapidly to a maximum of 2.14 and 2.87 respectively. This indicates the minimum exposure duration for robust hue estimation lies somewhere between 45 and 105 milliseconds.

We concluded our analysis by checking to see if orientation interfered with hue estimation at any of the exposure durations. t -tests were computed to compare estimation error means across control and experiment subsections for all five exposure durations. Table 6.5 shows the exposure duration, the number of control trials n_1 , the number of experiment trials n_2 , the degrees of

Exposure	n_1	n_2	v	t
15 ms	90	90	178	1.45
45 ms	90	90	178	0.51
105 ms	90	90	178	0.53
195 ms	90	90	178	0.31
450 ms	90	90	178	0.83

Table 6.5: t -test results of estimation error rates from control and experiment trials for all five exposure durations

freedom v , and the t -value t . The t -values for all durations are less than $_{0.975}t_{178} = 1.972$, except for the 15 millisecond exposure duration, where $_{0.90}t_{178} = 1.286 < p < _{0.95}t_{178} = 1.653$. This suggests that orientation is not interfering with hue estimation at any of the exposure durations tested. Additional experiments should test for hue interference, because Callaghan reported asymmetric interference patterns during her research.

Chapter 7

Conclusions

“O most lame and impotent conclusion!”

– Othello, *Shakespeare*

Preattentive features can be used to design simple and efficient visualization tools. In our studies, we simulated a visualization tool based on preattentive features that allowed rapid and accurate estimation to be performed within 450 milliseconds. “Real world” data from Oceanography’s salmon migration simulations was used during the experiments. Half the subjects were asked to estimate the percentage of rectangles in each display coloured blue, to the nearest 10%. The other half were asked to estimate the number of rectangles rotated 60° , to the nearest 10%. Average subject estimation error across all intervals was never more than 1.25. Standard deviation of error across all intervals was never more than 1.03. This shows subject estimations clustered around the correct answer for all intervals.

There was an “end-effect” present during the estimation task. Subjects made the largest estimation errors when 35-65% of the rectangles were targets (intervals 4, 5, and 6). Subjects made the smallest errors when 5-15% or 85-95% of the rectangles were targets (intervals 1 and 9).

Subjects showed no feature preference. They estimated targets that used hue and targets that used orientation equally well. This is no doubt due in part to the fact that we calibrated the perceived difference between our two hues and our two orientations to be equal for a related target detection task. There was evidence of a data type preference. Subjects performed better estimation when the primary data value was stream function. This indicates a subject's estimation ability depends, at least in part, on the spatial distribution of the data values being estimated. There was no evidence of interaction between feature and data type. Neither feature gained a greater improvement in estimation error when the preferred data type was used.

Probably the most important new result is that no feature interference was observed. Encoding an irrelevant data value with orientation did not affect hue estimation ability. Encoding an irrelevant data value with hue did not affect orientation estimation ability. This suggests that the two preattentive features can be used together in a single display. Users might use hue and orientation to encode data elements in two different ways. They would then be able to perform different types of preattentive data analysis on a single display.

Further experiments that varied exposure duration showed robust estimation is possible below 450 milliseconds. Estimation seemed reasonably accurate down to an exposure duration of 105 milliseconds. Below that, estimation ability deteriorated rapidly. There was no evidence that orientation interfered with estimation using hue at any of the five exposure durations tested.

Some of our results could be predicted from research in preattentive processing. Others were more surprising. We expected estimation using hue and orientation to be an extension of the results of Julész [Jul83] and Triesman [Tri88]. They used hue and orientation in their work to perform rapid and accurate target and boundary detection. Our data analysis shows this turned out to be the case. On the other hand, we also expected to observe an interference effect during our experiment. This would have been an extension of Callaghan's work, where she showed a "feature hierarchy" of brightness-hue-shape [Cal84][Cal89].

As our experiments suggest, results from preattentive processing can form a basis for the design of visualization tools, but the results cannot be applied without some additional investigation to ensure the desired effects hold in the particular visualization environment to which they are applied.

Chapter 8

Future Work

*“Issun saki wa yami
[No man knows his future]”*

– Japanese Proverb

Our experiments and related analysis leave a number of interesting avenues for future work. One obvious extension is to test for relationships among additional features. Intensity, size, and shape are three features often used for data visualization. Information on feature preference and interference would provide a more general and more useful set of guidelines on the use of these features in the design of visualization software.

Many visualization tasks require more than two data values to be encoded at each spatial location. Additional work could examine how to encode higher-dimensional elements in a low-dimensional environment. One obvious possibility is using three or more features in a single display. This type of visualization tool could exhibit new and unexpected types of interference. There may also be a limit to the amount of information a subject can extract and process at one time.

Our analysis showed no evidence of feature interference. Colour did not interfere with orientation estimation, and orientation did not interfere with hue estimation. It is possible that interference is sensitive to exposure duration. Subjects may have been able to overcome

any interference that did exist within our 450 millisecond exposure duration. With experiments designed to test for interference at different exposure durations, we could determine whether interference exists and the exposure durations at which it occurs. This type of experiment could also be used to search for possible feature preferences that may occur during shorter exposure durations.

Some subjects noted after the experiment they felt it would have been easier to estimate the number of red rectangles, or the number of rectangles rotated 0° . There may be a subject preference within feature for estimating. That is, red may be easier to estimate than blue, 0° rotation may be easier to estimate than 60° rotation, or vice-versa. It would be simple to set up a set of experiments similar to the ones we ran to test for these phenomena.

We explicitly chose two hues whose perceived difference from one another was equal to the perceived difference between two rectangles oriented 0° and 60° . A choice of features perceptually different from one another might cause a subject feature preference during the estimation task. For example, we could choose two isoluminant hues perceptually as far apart from one another as possible. A set of experiments could be run to see if estimation using hue was more rapid or accurate than estimation using orientation.

Estimation was shown to depend on exposure duration. Estimation error remained stable down to 105 milliseconds, then increased rapidly as the exposure duration fell. It seems unlikely that increases in exposure duration from 450 milliseconds will produce noticeably better estimation. Searching for feature interference and preference at exposure durations between 45 and 105 milliseconds would allow us to determine whether these phenomena occur at the boundary of estimation ability.

The data values used in our experiment were derived from salmon migration studies in Oceanography. More comprehensive studies based on actual tasks performed by researchers are needed before conclusive evidence will exist for using preattentive features. Other types of data should be investigated as well if general visualization tools are to be based on preattentive processing.

Bibliography

- [Asi85] Asimov, D. (1985). The Grand Tour: A Tool for Viewing Multidimensional Data. *SIAM Journal on Scientific and Statistical Computing*, **6**(1), 128–143.
- [Bec91a] Becker, R.A. and W.S. Cleveland (1991). Viewing Multivariate Scattered Data. *Pixel*, **2**(2), 36–41.
- [Bec91b] Becker, R.A. and W.S. Cleveland (1991). Take a Broader View of Scientific Visualization. *Pixel*, **2**(2), 42–44.
- [Bel87] Bell, P.C. and R.M. O’Keefe (1987). Visual Interactive Simulation—History, Recent Developments, and Major Issues. *Simulation*, **49**(3), 109–116.
- [Bir69] Birren, F. (ed.) (1969). *Munsell: A Grammar of Color*. New York, New York: Van Nostrand Reinhold Company.
- [Bra83] Bratley, P., B.L. Fox, and L.F. Schrage (1983). *A Guide to Simulation*. New York, New York: Springer-Verlag.
- [Bro88] Brown, M., D. Greenberg, M. Keeler, A.R. Smith, and L. Yaeger (1988). The Visualization Roundtable. *Computers in Physics*, **2**(3), 16–26.
- [Bru78] Bruckner, L.A. (1978). On Chernoff Faces. P.C.C. Wang (ed.), *Graphical Representation of Multivariate Data*, 93–121. New York, New York: Academic Press.
- [Bry91] Bryson, S. and C. Levit (1991). The Virtual Windtunnel: An Environment for the Exploration of Three-Dimensional Unsteady Flows. *Proceedings Visualization ’91*, 17–24. San Diego, United States.
- [Cal84] Callaghan, T.C. (1984). Dimensional Interaction of Hue and Brightness in Preattentive Field Segregation. *Perception & Psychophysics*, **36**(1), 25–34.
- [Cal89] Callaghan, T.C. (1989). Interference and Domination in Texture Segregation: Hue, Geometric Form, and Line Orientation. *Perception & Psychophysics*, **46**(4), 299–311.
- [Cal90] Callaghan, T.C. (1990). Interference and Dominance in Texture Segregation. D. Brogan (ed.), *Visual Search*, 81–87. New York, New York: Taylor & Francis.

- [Che73] Chernoff, H. (1973). The Use of Faces to Represent Points in k-Dimensional Space Graphically. *Journal of the American Statistical Association*, **68**(342), 361–367.
- [DeF91] DeFanti, T.A. and M.D. Brown (1991). Visualization in Scientific Computing. M.C. Yovits (ed.), *Advances in Computers*, **33**, 247–305. New York, New York: Academic Press.
- [DeS84] DeSanctis, G. (1984). Computer Graphics As Decision Aids: Directions for Research. *Decision Sciences*, **15**, 463–487.
- [Dre88] Drebin, R.A., L. Carpenter, and P. Hanrahan (1988). Volume Rendering. *Computer Graphics*, **22**(4), 65–74.
- [Dun89] Duncan, J. and G.W. Humphreys (1989). Visual Search and Stimulus Similarity. *Psychological Review*, **96**(3), 433–458.
- [Enn90a] Enns, J.T. and R.A. Rensink (1990). Sensitivity to Three-Dimensional Orientation in Visual Search. *Psychology Science*, **1**(5), 323–326.
- [Enn90b] Enns, J.T. and R.A. Rensink (1990). Influence of Scene-Based Properties on Visual Search. *Science*, **247**, 721–723.
- [Enn90c] Enns, J.T. (1990). The Promise of Finding Effective Geometric Codes. *Proceedings Visualization '90*, 389–390. San Francisco, United States.
- [Enn90d] Enns, J.T. (1990). Three-Dimensional Features that Pop Out in Visual Search. D. Brogan (ed.), *Visual Search*, 37–45. New York, New York: Taylor & Francis.
- [Enn91] Enns, J.T. and R.A. Rensink (1991). VSearch Colour: Full-Colour Visual Search Experiments on the Macintosh II. *Behaviour Research Methods, Instruments, & Computers*, **23**(2), 265–272.
- [Fol90] Foley, J.D., A. Van Dam, S.K. Feiner, and J.F. Hughes (1990). *Computer Graphics, Principles and Practice*. Reading, Massachusetts: Addison-Wesley.
- [Fra77] Franta, W.R. (1977). *The Process View of Simulation*. New York, New York: North-Holland.
- [Gla84] Glass, G.V. and K.D. Hopkins (1984). *Statistical Methods in Education and Psychology*. Englewood Cliffs, New Jersey: Prentice-Hall.
- [Gor69] Gordan, G. (1969). *System Simulation*. Englewood Cliffs, New Jersey: Prentice-Hall.
- [Gri89] Grinstein, G., R. Pickett, and M. Williams (1988). EXVIS: An Exploratory Visualization Environment. *Proceedings Graphics Interface '89*, 254–261. London, Canada.

- [Her74] Herdeg, W. (ed.) (1974). *Graphis/Diagrams: The Graphical Visualization of Abstract Data*. New York, New York: Hastings House Publishers.
- [Hib90] Hibbard, B. and D. Santek (1990). The VIS-5D System for Easy Interactive Visualization. *Proceedings Visualization '90*, 28–35. San Francisco, United States.
- [Hur80] Hurrion, R.D. (1980). An Interactive Visual Simulation System for Industrial Management. *European Journal of Operational Research*, **5**, 86–93.
- [Ing88] Ingraham, W.J., Jr. and R.K. Miyahara (1988). Ocean Surface Current Simulations in the North Pacific Ocean and Bering Sea (OSCURS Numerical Model). *NMFS F/NWC-130*. National Oceanic and Atmospheric Association Technical Memo, 155 pp. Seattle, United States.
- [Ing89] Ingraham, W.J., Jr. and R.K. Miyahara (1989). OSCURS Numerical Model to Ocean Surface Current Measurements in the Gulf of Alaska. *NMFS F/NWC-168*. National Oceanic and Atmospheric Association Technical Memo, 67 pp. Seattle, United States.
- [Jul81] Julész, B. (1981). Textons, the Elements of Texture Perception, and their Interactions. *Nature*, **290**, 91–97.
- [Jul83] Julész, B. and J.R. Bergen (1983). Textons, the Fundamental Elements in Preattentive Vision and Perception of Textures. *The Bell System Technical Journal*, **62**(6), 1619–1645.
- [Jul84] Julész, B. (1984). A Brief Outline of the Texton Theory of Human Vision. *Trends in Neuroscience*, **7**(2), 41–45.
- [Kau81] Kaufman, A. and M.Z. Hanani (1981). Converting a Batch Simulation Program to an Interactive Program with Graphics. *Simulation*, **36**(4), 125–131.
- [LeB90] LeBlond, P. (1990). Influences of Currents, Temperature and Salinity on Open Ocean Migrations of Sockeye Salmon in the Northeast Pacific Ocean. *DFO/NSERC Science Subvention Program Submission*. Department of Oceanography, University of British Columbia.
- [Lev90] Levkowitz, H. and R.M. Pickett. Iconographic Integrated Displays of Multiparameter Spatial Distributions. Rogowitz, B.E. and J.P. Allebach (eds.), *Human Vision and Electronic Imaging: Models, Methods, and Applications*, 345–355. Bellingham, Washington: SPIE.
- [Lev91] Levkowitz, H. (1990). Color Icons: Merging Color and Texture Perception for Integrated Visualization of Multiple Parameters. *Proceedings Visualization '91*, 164–170. San Diego, United States.

- [McC87] McCormick, B.H., T.A. DeFanti, and M.D. Brown (1987). Visualization in Scientific Computing—A Synopsis. *IEEE Computer Graphics & Applications*, **7**(7), 61–70.
- [Mel85] Melamed, B. and R.J.T. Morris (1985). Visual Simulation: The Performance Analysis Workstation. *IEEE Computer*, **18**, 87–94.
- [Mül90] Müller, H.J., G.W. Humphreys, P.T. Quinlan, and M.J. Riddoch (1990). Combined-Feature Coding in the Form Domain. D. Brogan (ed.), *Visual Search*, 47–55. New York, New York: Taylor & Francis.
- [OKe86] O’Keefe, R. and R. Davies (1986). A Microcomputer System for Simulation Modelling. *European Journal of Operational Research*, **24**, 23–29.
- [Pea92] Pearcy, W.G. (1992). *Ocean Ecology of North Pacific Salmonids*. Seattle, Washington: University of Washington Press.
- [Pic88] Pickett, R. and G. Grinstein (1988). Iconographic Displays for Visualizing Multidimensional Data. *Proceedings of the 1988 IEEE Conference on Systems, Man, and Cybernetics*, 514–519. Beijing and Shenyang, China.
- [Qui87] Quinlan, P.T. and G.W. Humphreys (1987). Visual Search for Targets Defined by Combinations of Color, Shape, and Size: An Examination of the Task Constraints on Feature and Conjunction Searches. *Perception & Psychophysics*, **41**(5), 455–472.
- [Reu90] Reuter, L.H. (1990). Human Perception and Visualization. *Proceedings Visualization ’90*, 401–406. San Francisco, United States.
- [Set88] Sethian, J.A., J.B. Salem, and A.F. Ghoniem (1988). Interactive Scientific Visualization and Parallel Display Techniques. *Proceedings Supercomputing ’88*, 132–139. Orlando, United States.
- [Tho92a] Thomson, K.A., W.J. Ingraham, M.C. Healey, P.H. LeBlond, C. Groot, and C.G. Healey (1992). The Influence of Ocean Currents on Latitude of Landfall and Migration Speed of Sockeye Salmon Returning to the Fraser River. *Fisheries Oceanography Journal*, **1**(2), 163–179.
- [Tho92b] Thomson, K.A., W.J. Ingraham, M.C. Healey, P.H. LeBlond, C. Groot, and C.G. Healey (1992). The Influence of Ocean Currents on Return Timing of Fraser River Sockeye Salmon. *Fisheries Oceanography Journal*, in press.
- [Tre89] Treinish, L.A., J.D. Foley, W.J. Campbell, R.B. Haber, and R.F. Gurwitz (1989). Effective Software Systems for Scientific Data Visualization. *Computer Graphics*, **23**(5), 111–136.
- [Tre91] Treinish, L.A. and T. Goettsche (1991). Correlative Visualization Techniques for Multidimensional Data. *IBM Journal of Research and Development*, **35**(1/2), 184–204.

- [Tri85] Triesman, A. (1985). Preattentive Processing in Vision. *Computer Vision, Graphics, and Image Processing*, **31**, 156–177.
- [Tri88] Triesman, A. and S. Gormican (1988). Feature Analysis in Early Vision: Evidence from Search Asymmetries. *Psychological Review*, **95**(1), 15–48.
- [Tri91] Triesman, A. (1991). Search, Similarity, and Integration of Features Between and Within Dimensions. *Journal of Experimental Psychology: Human Perception and Performance*, **17**(3), 652–676.
- [Tuf83] Tufte, E.R. (1983). *The Visual Display of Quantitative Information*. Cheshire, Connecticut: Graphics Press.
- [Tuf90] Tufte, E.R. (1990). *Envisioning Information*. Cheshire, Connecticut: Graphics Press.
- [Van90a] Vande Wetering, M. (1990). The Application Visualization System—AVS 2.0. *Pixel*, **1**(3), 30–33.
- [Van90b] Vande Wetering, M. (1990). avE 2.0. *Pixel*, **1**(4), 30–35.
- [War85] Ware, C. and J.C. Beatty (1985). Using Colour As A Tool in Discrete Data Analysis. *Technical Report CS-85-21*. Computer Science Department, University of Waterloo.
- [War88] Ware, C. and J.C. Beatty (1988). Using Colour Dimensions to Display Data Dimensions. *Human Factors*, **30**(2), 127–142.
- [Wol88] Wolff, R.S. (1988). Visualization in the Eye of the Scientist. *Computers in Physics*, **2**(3), 29–35.
- [Wys82] Wyszecki, G. and W.S. Stiles (1982). *Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd Edition*. New York, New York: John Wiley & Sons, Inc.
- [Yag91] Yagel, R., A. Kaufman, and Q. Zhang (1991). Realistic Volume Imaging. *Proceedings Visualization '91*, 226–231. San Diego, United States.

Chapter 10

Appendix A

Experiment Results and Graphs

The following tables and graphs summarize the preattentive estimation experiment results. Subjects were asked to estimate the relative percentage of elements in each display with a given preattentive feature, to the nearest 10%. In blocks B₁ and B₃ subjects were asked to estimate the number of rectangles coloured blue, to the nearest 10%. In blocks B₂ and B₄, subjects were asked to estimate the number of rectangles oriented 60°.

Twelve subjects were used during the experiment. Each subject completed either blocks B₁ and B₃ (blocks using hue as the primary feature) or blocks B₂ and B₄ (blocks using orientation as the primary feature) for a total of 288 trials.

Tables 10.1–10.4 list summary information for all four blocks. Each block is subdivided into control 1 trials, control 2 trials, and experiment trials for each of the nine intervals. The tables show average subject response \bar{V} , standard deviation of subject response $\sigma(V)$, average subject error \bar{e} , and standard deviation of subject error $\sigma(e)$. Figures 10.1–10.12 show graphs of the control and experiment subsections of each table.

Class	Control 1				Control 2				Experiment			
	\bar{V}	$\sigma(V)$	\bar{e}	$\sigma(e)$	\bar{V}	$\sigma(V)$	\bar{e}	$\sigma(e)$	\bar{V}	$\sigma(V)$	\bar{e}	$\sigma(e)$
1	1.25	0.53	0.25	0.53	1.33	0.70	0.33	0.70	1.29	0.68	0.29	0.68
2	1.83	0.82	0.58	0.58	2.04	0.86	0.62	0.58	2.17	0.83	0.46	0.71
3	2.71	0.75	0.46	0.66	2.75	0.85	0.67	0.56	2.79	0.71	0.54	0.50
4	4.17	1.13	0.75	0.85	3.75	1.11	0.83	0.76	3.83	1.49	1.08	1.03
5	5.50	1.32	1.00	0.98	5.08	1.67	1.42	0.83	5.54	1.41	1.25	0.84
6	5.96	1.27	0.96	0.81	6.71	1.23	1.21	0.72	6.31	1.17	0.94	0.76
7	6.83	1.01	0.75	0.68	7.42	0.78	0.67	0.56	7.19	0.73	0.52	0.55
8	8.13	0.80	0.46	0.66	8.33	0.56	0.42	0.50	8.15	0.62	0.40	0.49
9	8.71	0.55	0.29	0.55	8.96	0.20	0.04	0.20	8.65	0.53	0.35	0.53
Total	5.01	2.71	0.61	0.75	5.15	2.84	0.69	0.74	5.10	2.72	0.65	0.77

Table 10.1: Summary of combined landfall/hue experiment results

Class	Control 1				Control 2				Experiment			
	\bar{V}	$\sigma(V)$	\bar{e}	$\sigma(e)$	\bar{V}	$\sigma(V)$	\bar{e}	$\sigma(e)$	\bar{V}	$\sigma(V)$	\bar{e}	$\sigma(e)$
1	1.25	0.44	0.25	0.44	1.21	0.51	0.21	0.51	1.13	0.39	0.13	0.39
2	2.13	0.45	0.21	0.41	2.25	0.68	0.42	0.58	2.15	0.74	0.50	0.58
3	2.83	0.70	0.50	0.51	2.92	0.65	0.42	0.50	3.08	0.77	0.54	0.54
4	4.13	1.08	0.71	0.86	3.17	0.76	0.83	0.76	3.71	0.77	0.58	0.58
5	5.50	1.29	1.00	0.93	5.13	1.15	0.96	0.62	5.50	1.25	1.00	0.90
6	6.29	1.20	0.96	0.75	6.67	0.96	0.92	0.72	6.40	1.35	1.19	0.73
7	7.08	0.97	0.75	0.61	7.38	1.01	0.88	0.61	7.25	1.00	0.83	0.60
8	8.17	0.70	0.50	0.51	7.96	0.86	0.63	0.58	8.08	0.71	0.46	0.54
9	8.83	0.38	0.17	0.38	8.79	0.51	0.21	0.51	8.83	0.43	0.17	0.43
Total	5.13	2.69	0.56	0.68	5.05	2.73	0.61	0.66	5.13	2.72	0.60	0.69

Table 10.2: Summary of combined landfall/orientation experiment results

Class	Control 1				Control 2				Experiment			
	\bar{V}	$\sigma(V)$	\bar{e}	$\sigma(e)$	\bar{V}	$\sigma(V)$	\bar{e}	$\sigma(e)$	\bar{V}	$\sigma(V)$	\bar{e}	$\sigma(e)$
1	1.13	0.34	0.13	0.34	1.21	0.41	0.21	0.41	1.23	0.47	0.23	0.47
2	1.83	0.70	0.50	0.51	2.00	0.59	0.33	0.48	2.17	0.52	0.29	0.46
3	3.17	0.82	0.50	0.66	3.21	0.78	0.63	0.49	3.27	0.87	0.65	0.64
4	3.71	0.86	0.71	0.55	4.54	1.14	0.96	0.81	4.42	1.09	0.92	0.71
5	5.08	0.78	0.58	0.50	5.63	0.97	0.88	0.74	5.02	0.91	0.69	0.59
6	6.04	0.95	0.71	0.62	6.21	0.88	0.83	0.65	6.27	1.16	0.90	0.78
7	6.71	0.55	0.38	0.49	7.00	1.10	0.83	0.70	7.33	0.66	0.50	0.55
8	7.83	0.70	0.42	0.58	7.67	0.82	0.58	0.65	7.69	0.59	0.35	0.56
9	8.75	0.44	0.25	0.44	8.75	0.44	0.25	0.44	8.69	0.47	0.31	0.47
Total	4.92	2.60	0.46	0.55	5.13	2.58	0.59	0.66	5.12	2.56	0.54	0.63

Table 10.3: Summary of combined stream function/hue experiment results

Class	Control 1				Control 2				Experiment			
	\bar{V}	$\sigma(V)$	\bar{e}	$\sigma(e)$	\bar{V}	$\sigma(V)$	\bar{e}	$\sigma(e)$	\bar{V}	$\sigma(V)$	\bar{e}	$\sigma(e)$
1	1.08	0.28	0.08	0.28	1.08	0.28	0.08	0.28	1.08	0.28	0.08	0.28
2	2.17	0.56	0.33	0.48	2.29	0.81	0.46	0.72	2.13	0.61	0.33	0.52
3	3.25	0.90	0.67	0.64	3.13	0.61	0.38	0.49	3.15	0.68	0.40	0.57
4	4.08	0.65	0.42	0.50	4.29	0.95	0.71	0.69	4.21	0.77	0.54	0.58
5	4.50	0.93	0.83	0.64	5.46	0.93	0.71	0.75	4.79	1.20	1.00	0.68
6	5.88	1.15	0.79	0.83	6.13	1.08	0.71	0.81	5.83	1.23	0.92	0.82
7	6.71	0.91	0.63	0.71	7.04	1.04	0.79	0.66	6.79	0.99	0.79	0.62
8	8.04	0.69	0.38	0.58	7.96	0.75	0.46	0.59	7.92	0.87	0.63	0.60
9	8.88	0.34	0.13	0.34	8.83	0.38	0.17	0.38	8.75	0.48	0.25	0.48
Total	4.95	2.60	0.47	0.62	5.13	2.61	0.50	0.65	4.96	2.59	0.55	0.66

Table 10.4: Summary of combined stream function/orientation experiment results




Figure 10.1: Graph of combined results for control 1 subsection of block B₁, primary value landfall, primary feature hue

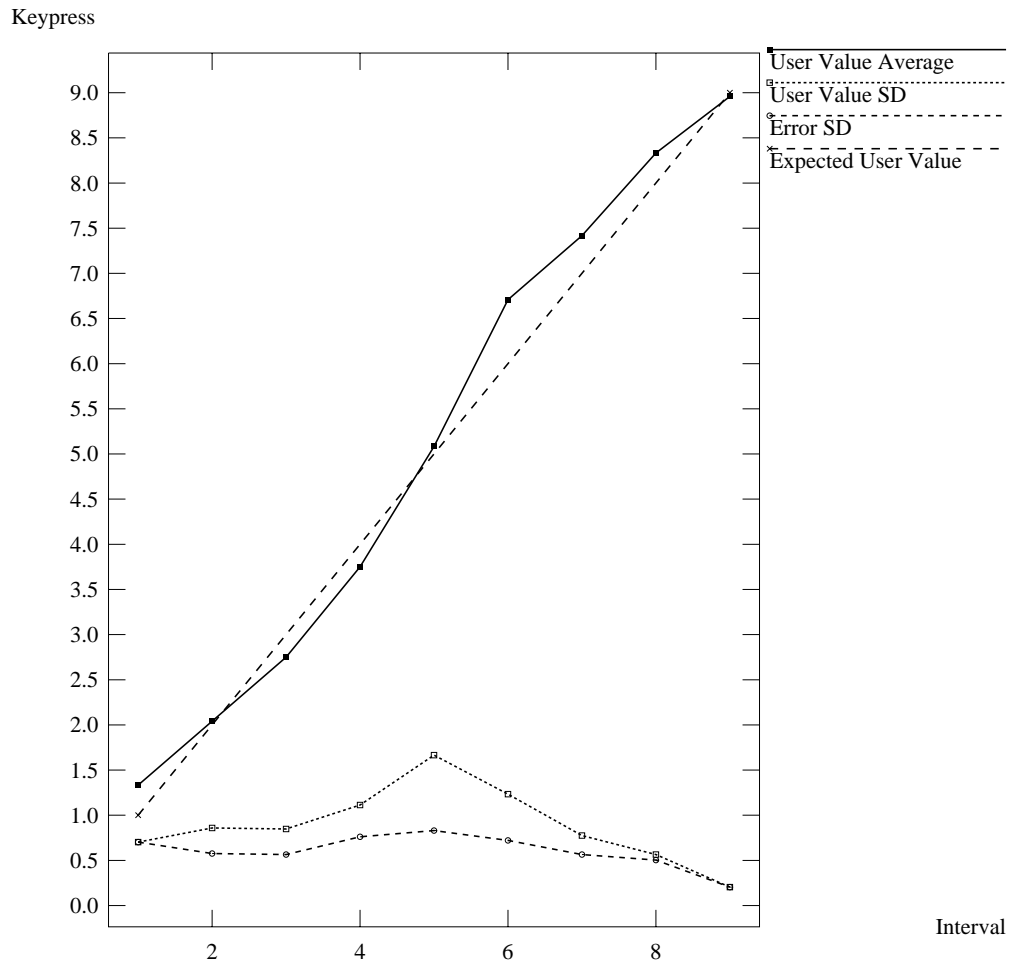


Figure 10.2: Graph of combined results for control 2 subsection of block B₁, primary value landfall, primary feature hue




Figure 10.3: Graph of combined results for experiment subsection of block B₁, primary value landfall, primary feature hue




Figure 10.4: Graph of combined results for control 1 subsection of block B₂, primary value landfall, primary feature orientation




Figure 10.5: Graph of combined results for control 2 subsection of block B₂, primary value landfall, primary feature orientation

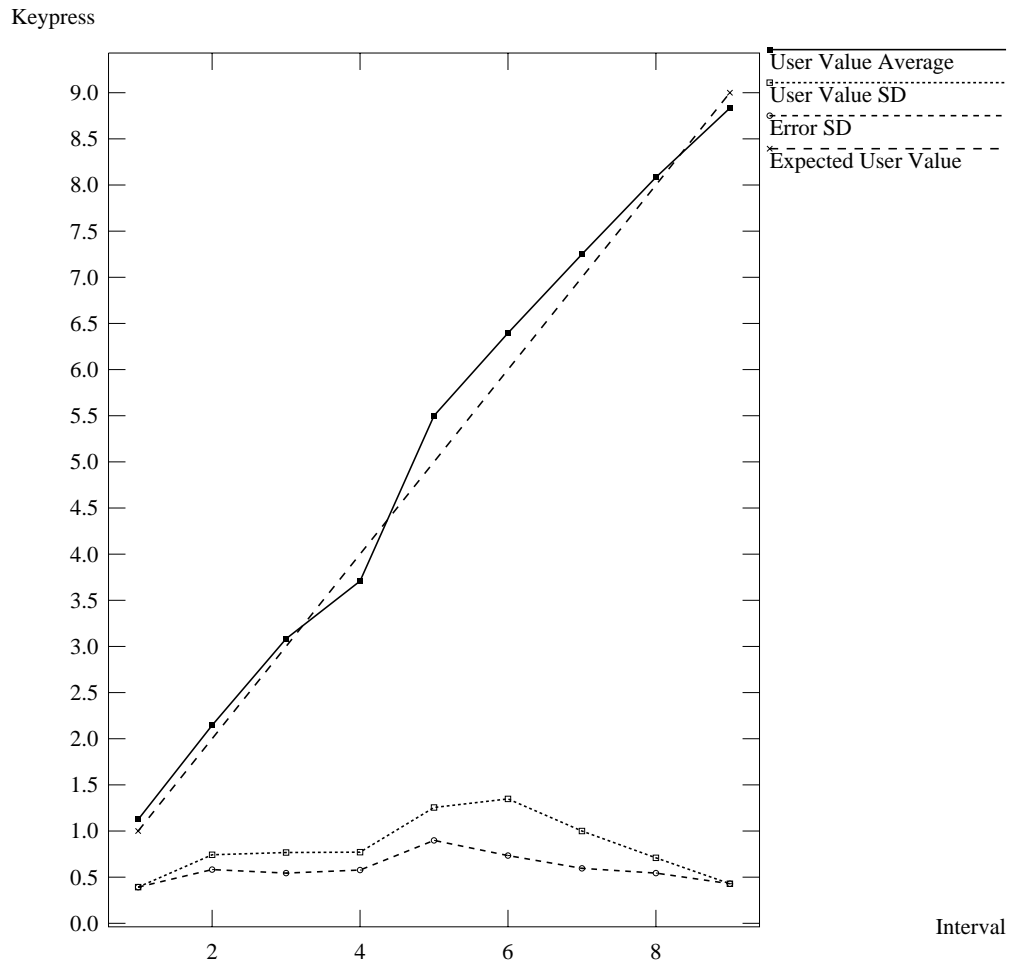


Figure 10.6: Graph of combined results for experiment subsection of block B₂, primary value landfall, primary feature orientation




Figure 10.7: Graph of combined results for control 1 subsection of block B₃, primary value stream function, primary feature hue




Figure 10.8: Graph of combined results for control 2 subsection of block B₃, primary value stream function, primary feature hue




Figure 10.9: Graph of combined results for experiment subsection of block B₃, primary value stream function, primary feature hue




Figure 10.10: Graph of combined results for control 1 subsection of block B₄, primary value stream function, primary feature orientation




Figure 10.11: Graph of combined results for control 2 subsection of block B₄, primary value stream function, primary feature orientation




Figure 10.12: Graph of combined results for experiment subsection of block B₄, primary value stream function, primary feature orientation

Chapter 11

Appendix B

Discrimination Experiment Results

The following tables list the discrimination experiment results. Subjects were asked to detect the presence or absence of a target element in a field of distractor elements. Each subject completed 11 “blocks”, where a block corresponded to a unique target. For each block, the distractors were rectangles rotated 0° and coloured 5R7/8. In one block the target was a rectangle rotated 60° and coloured 5R7/8. The targets for the remaining blocks were rectangles rotated 0° and coloured to be anywhere from 1 to 10 Munsell “hue steps” away from the distractors. The hue steps 1 through 10 correspond to the Munsell hues 10RP 7/8, 5RP 7/8, 10P 7/8, 5P 7/8, 10PB 7/8, 5PB 7/8, 10B 7/8, 5B 7/8, 10GB 7/8 and 5GB 7/8.

Within each block subjects completed 30 “trials”. 15 of the 30 trials were randomly selected to contain the target element. A trial consisted of 36 rectangles (including the target, if present) presented in a random pattern on the screen.

Tables 11.1–11.4 list summary information for all 11 experiment blocks. Each block is subdivided into target present trials, target absent trials, and the combination of both present and absent trials. The tables show number of correct responses n , average response time \bar{x} ,

response time variance σ^2 , and response time standard deviation σ . Subjects were unable to discriminate between the target and distractors for the first two “hue-step” blocks (Munsell hues 10RP 7/8 and 5RP 7/8). Table 11.5 shows the combined results for all subjects.

Target	Present				Absent				Combined			
	n	\bar{x}	σ^2	σ	n	\bar{x}	σ^2	σ	n	\bar{x}	σ^2	σ
60°	13	403	2762	53	15	491	14367	120	28	450	10660	103
1 HS			—				—				—	
2 HS			—				—				—	
3 HS	8	772	40858	202	12	836	41049	203	20	810	39856	200
4 HS	14	501	8105	90	13	529	5507	74	27	515	6798	82
5 HS	14	469	5810	76	15	566	5114	72	29	520	7697	88
6 HS	15	434	2571	51	15	506	8115	90	30	470	6519	81
7 HS	14	401	2490	50	15	495	7219	85	29	449	7065	84
8 HS	15	394	1150	34	15	447	5177	72	30	420	3775	61
9 HS	15	421	8552	92	15	426	3278	57	30	423	5716	76
10 HS	14	394	2426	49	15	548	24496	157	29	474	19516	140

Table 11.1: Summary of discrimination experiment results for subject 1

Target	Present				Absent				Combined			
	n	\bar{x}	σ^2	σ	n	\bar{x}	σ^2	σ	n	\bar{x}	σ^2	σ
60°	15	618	16628	129	15	1191	44119	210	30	905	114061	338
1 HS			—				—				—	
2 HS			—				—				—	
3 HS	11	1193	193292	440	10	1285	217785	467	21	1237	196903	444
4 HS	14	822	120381	347	15	841	16605	129	29	832	64282	254
5 HS	14	735	44779	212	14	834	77901	279	28	784	61605	248
6 HS	15	577	71369	267	15	805	26238	162	30	691	60581	246
7 HS	15	481	2738	52	14	857	162389	403	29	663	113263	337
8 HS	15	477	16007	127	15	620	32697	181	30	549	28801	170
9 HS	15	386	3098	56	14	504	4212	65	29	443	7100	84
10 HS	15	436	10274	101	15	491	5677	75	30	463	8469	92

Table 11.2: Summary of discrimination experiment results for subject 2

Target	Present				Absent				Combined			
	n	\bar{x}	σ^2	σ	n	\bar{x}	σ^2	σ	n	\bar{x}	σ^2	σ
60°	15	555	29160	171	15	557	9987	99	30	556	18851	137
1 HS			—				—				—	
2 HS			—				—				—	
3 HS	10	966	89299	299	15	984	17150	131	25	977	43570	209
4 HS	12	715	38513	196	14	931	35342	188	26	831	47387	218
5 HS	15	496	17275	131	13	584	11293	106	28	537	15999	126
6 HS	15	527	4809	69	15	768	7273	85	30	648	20771	144
7 HS	15	453	15091	123	15	492	2462	50	30	473	8869	94
8 HS	15	413	3246	57	14	518	11575	108	29	464	9826	99
9 HS	15	393	2746	52	15	425	2326	48	30	409	2711	52
10 HS	15	361	1119	33	15	412	2481	50	30	386	2409	49

Table 11.3: Summary of discrimination experiment results for subject 3

Target	Present				Absent				Combined			
	n	\bar{x}	σ^2	σ	n	\bar{x}	σ^2	σ	n	\bar{x}	σ^2	σ
60°	15	368	406	20	15	346	468	22	30	357	553	24
1 HS			—				—				—	
2 HS			—				—				—	
3 HS	11	791	77716	279	12	720	102299	320	23	754	87777	296
4 HS	14	431	8765	94	13	451	4107	64	27	441	6379	80
5 HS	13	369	2039	45	15	343	1765	42	28	355	1987	45
6 HS	15	405	3733	61	15	459	7191	85	30	432	6052	78
7 HS	15	341	420	20	15	368	1941	44	30	354	1328	36
8 HS	15	337	573	24	15	334	2184	47	30	336	1335	37
9 HS	15	337	2244	47	15	335	1521	39	30	336	1819	43
10 HS	14	330	968	31	15	343	1644	41	29	336	1317	36

Table 11.4: Summary of discrimination experiment results for subject 4

Target	Present				Absent				Combined			
	n	\bar{x}	σ^2	σ	n	\bar{x}	σ^2	σ	n	\bar{x}	σ^2	σ
60°	58	489	23003	152	60	646	122813	350	118	569	79353	282
1 HS			—				—				—	
2 HS			—				—				—	
3 HS	40	941	127659	357	49	945	118882	345	89	943	121423	384
4 HS	54	614	67568	260	55	698	56186	237	109	656	63039	251
5 HS	56	519	34780	186	57	577	53383	231	113	549	44621	211
6 HS	60	486	24486	156	60	635	35457	188	120	560	35314	188
7 HS	59	419	7931	89	59	548	72188	269	118	484	43882	209
8 HS	60	405	7507	87	59	479	23500	153	119	442	16667	129
9 HS	60	384	4879	70	59	421	6224	79	119	403	5837	76
10 HS	58	381	5177	72	60	448	14318	120	118	415	10894	104

Table 11.5: Summary of discrimination experiment results for all subjects

Chapter 12

Appendix C

Exposure Duration Experiment Results

The following tables list the exposure duration experiment results. Subjects were shown trials that contained 174 rectangles coloured red and blue. Subjects were asked to estimate the percentage of rectangles coloured blue, to the nearest multiple of 10%. Each subject completed 180 trials. In 90 trials, the rectangles were randomly chosen to be oriented either 0° or 60° . In 45 trials, the rectangles were all oriented 0° . In the remaining 45 trials, the rectangles were all oriented 60° . Each trial was shown for a fixed exposure duration. There were an equal number of trials from five possible exposure durations: 15 milliseconds, 45 milliseconds, 105 milliseconds, 195 milliseconds, and 450 milliseconds.

Tables 12.1–12.5 list summary information for all subjects, grouped by exposure duration. Each duration is divided into control 1 trials (0° orientation), control 2 trials (60° orientation), and experiment trials (random orientation). The tables show the number of trials n , the average estimation error \bar{e} , and the standard deviation of estimation error $\sigma(e)$. Estimation error is the absolute value of the correct response minus the subject's response. Table 12.6 shows the combined results for all subjects.

Exposure	Control 1			Control 2			Experiment		
	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$
15 ms	9	2.88	3.74	9	1.88	2.57	18	1.05	1.66
45 ms	9	0.77	1.27	9	1.00	1.36	18	0.88	1.13
105 ms	9	0.88	1.11	9	0.77	1.06	18	0.55	0.84
195 ms	9	0.44	0.86	9	0.66	0.86	18	0.44	0.76
450 ms	9	0.55	0.93	9	0.88	1.22	18	0.44	0.76

Table 12.1: Summary of exposure duration experiment results for subject 1

Exposure	Control 1			Control 2			Experiment		
	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$
15 ms	9	2.33	3.69	9	1.77	2.50	18	1.77	2.11
45 ms	9	1.44	1.90	9	2.44	3.00	18	1.72	2.07
105 ms	9	1.66	2.80	9	1.00	1.45	18	0.72	0.93
195 ms	9	1.66	3.02	9	0.66	1.11	18	1.16	1.51
450 ms	9	1.00	1.54	9	0.88	1.32	18	0.72	1.05

Table 12.2: Summary of exposure duration experiment results for subject 2

Exposure	Control 1			Control 2			Experiment		
	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$
15 ms	9	1.22	1.96	9	2.33	2.80	18	1.50	2.23
45 ms	9	1.33	1.93	9	0.77	1.27	18	0.94	1.47
105 ms	9	0.77	1.27	9	0.77	1.45	18	0.94	1.30
195 ms	9	0.44	0.70	9	0.55	0.93	18	0.77	1.13
450 ms	9	0.44	0.70	9	0.77	1.06	18	0.88	1.23

Table 12.3: Summary of exposure duration experiment results for subject 3

Exposure	Control 1			Control 2			Experiment		
	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$
15 ms	9	1.33	2.12	9	2.55	3.37	18	1.22	1.71
45 ms	9	0.88	1.22	9	2.00	2.59	18	1.16	1.55
105 ms	9	0.55	0.79	9	0.55	0.93	18	0.77	1.60
195 ms	9	0.66	0.86	9	1.11	1.50	18	0.77	1.18
450 ms	9	0.44	0.86	9	0.77	1.27	18	0.44	0.84

Table 12.4: Summary of exposure duration experiment results for subject 4

Exposure	Control 1			Control 2			Experiment		
	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$
15 ms	9	1.88	2.31	9	2.22	3.00	18	2.05	2.73
45 ms	9	0.88	1.11	9	1.00	1.27	18	0.94	1.35
105 ms	9	0.66	1.11	9	0.88	1.11	18	0.77	1.13
195 ms	9	0.33	0.61	9	0.55	1.06	18	0.66	0.90
450 ms	9	0.55	1.06	9	0.66	1.00	18	0.38	0.64

Table 12.5: Summary of exposure duration experiment results for subject 5

Exposure	Control 1			Control 2			Experiment		
	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$	n	\bar{e}	$\sigma(e)$
15 ms	45	1.93	2.74	45	2.15	2.73	90	1.52	2.07
45 ms	45	1.06	1.46	45	1.44	1.94	90	1.13	1.51
105 ms	45	0.91	1.51	45	0.80	1.16	90	0.75	1.17
195 ms	45	0.71	1.44	45	0.71	1.06	90	0.76	1.10
450 ms	45	0.60	1.01	45	0.80	1.12	90	0.57	0.91

Table 12.6: Summary of exposure duration experiment results for all subjects