High-Speed Visual Estimation Using Preattentive Processing

CHRISTOPHER G. HEALEY, KELLOGG S. BOOTH, and JAMES T. ENNS
The University of British Columbia

A new method is presented for performing rapid and accurate numerical estimation. The method is derived from an area of human cognitive psychology called preattentive processing. Preattentive processing refers to an initial organization of the visual field based on cognitive operations believed to be rapid, automatic, and spatially parallel. Examples of visual features that can be detected in this way include hue, intensity, orientation, size, and motion. We believe that studies from preattentive vision should be used to assist in the design of visualization tools, especially those for which high speed target detection, boundary identification, and region detection are important. In our present study, we investigated two known preattentive features (hue and orientation) in the context of a new task (numerical estimation) in order to see whether preattentive estimation was possible. Our experiments tested displays that were designed to visualize data from salmon migration simulations. The results showed that rapid and accurate estimation was indeed possible using either hue or orientation. Furthermore, random variation in one of these features resulted in no interference when subjects estimated the percentage of the other. To test the generality of our results, we varied two important display parameters, display duration and feature difference, and found boundary conditions for each. Implications of our results for application to real-word data and tasks are discussed.

Categories and Subject Descriptors: H.5.2 [**Information Interfaces and Presentation**]: User Interfaces–*ergonomics, screen design (graphics, colour)*; I.3.6 [**Computer Graphics**]: Methodology and Techniques–*ergonomics, interaction techniques*

General Terms: Experimentation, Performance

Additional Key Words and Phrases: boundary detection, cognitive psychology, colour, estimation, icon, human vision, multidimensional data, Munsell, orientation, preattentive, scientific visualization, target detection

1. INTRODUCTION

This article describes three experiments investigating the ability of humans to perform high-speed visual estimation of proportions or percentages (hereafter referred

Authors' addresses: C.G. Healey and K.S. Booth, Imager Computer Graphics Laboratory, Department of Computer Science, 2366 Main Mall, University of British Columbia, Vancouver, British Columbia, V6T 1Z4, Canada; healey@cs.ubc.ca, ksbooth@cs.ubc.ca. James T. Enns, Department of Psychology, 2136 Main Mall, University of British Columbia, Vancouver, British Columbia, V6T 1Z4, Canada; jenns@cortex.psych.ubc.ca.

to as numerical estimation). This work is part of an ongoing study of techniques that allow rapid and accurate visualization of large multidimensional datasets.

Scientific visualization is a relatively new subfield within the larger area of computer graphics. The term "visualization" was first used in this sense by a National Science Foundation (NSF) panel discussing the application of computer science to problems of data analysis [McCormick et al., 1987]. The panel defined the "domain of visualization" to include the development of general purpose tools and the study of research problems that arise in the process.

A variety of methods have been used to convert raw data into a more usable visual format. Both Tufte [Tufte, 1990] and Collins [Collins, 1993] give an interesting review of pre-computer visualization techniques. Many current systems are now being extended to provide user interaction and real-time visualization of results [Bell and O'Keefe, 1987; Hurrion, 1980]. Specialized software tools such as the Application Visualization System, apE, VIS-5D, and the Wavefront Data Visualizer [Hibbard and Santek, 1990; Upson, 1989; Vande Wettering, 1990] have been developed for computer graphics workstations. Diverse solutions for displaying high-dimensional datasets in a low-dimensional environment such as the computer screen have been proposed [Enns, 1990a; Enns, 1990b; Grinstein et al., 1989; Pickett and Grinstein, 1988; Ware and Beatty, 1988].

A recent update on the NSF visualization report described current research being performed at a number of academic institutions [Rosenblum, 1994]. Although many visual presentation techniques have been studied (*e.g.* volume visualization, fluid flow, and perceptual visualization), much less work has focused on formulating guidelines for their design. Our work is intended to address this more general issue.

We began by deliberately trying to take advantage of some of the built-in capabilities of the human visual system. Preattentive vision refers to cognitive operations that can be performed prior to focusing attention on any particular region of an image. We hypothesised that results from research on this ability could be used to assist in the design of visualization tools. Here, we use "visualization tools" to refer to software programs that display data values on a computer screen using simple visual features such as colour, shape, and size. Such tools, if properly designed, should allow users to perform certain types of visual analysis very rapidly and accurately. Tasks of interest include the detection of an element with a unique characteristic, the grouping of similar data elements, and the estimation of the percentage of elements with a given value or range of values. We asked the following specific questions:

- Can results from the existing literature on preattentive processing be used to build visualization tools that permit rapid and accurate performance of simple tasks such as target detection or boundary detection?
- Can these results be extended to include rapid and accurate numerical estimation?
- How do changes in the display duration and the perceived feature difference influence a user's ability to perform numerical estimation?

- Can results from our laboratory experiments be applied to real-world data and tasks?

Previous work in preattentive processing can be applied directly to visualization tools in which high-speed target and boundary detection are important [Healey et al., 1995]. Results from the present study show that preattentive processing can be extended to include high-speed visual estimation. They also point to important boundary conditions on this ability. We used our results to build a tool for visualizing data from salmon migration simulations being conducted in the Department of Oceanography. Anecdotal reports from users of this visualization tool are consistent with our experimental findings.

Before describing our experiments and results, we provide a brief introduction to research on preattentive processing.

2. PREATTENTIVE PROCESSING

One very interesting result of vision research over the past 20 years has been the discovery of a limited set of visual properties that are processed preattentively (*i.e.* without the need for focused attention). Typically, tasks that can be performed on large multi-element displays in 200 milliseconds or less are considered preattentive. This is because eye movements take at least 200 milliseconds to initiate. Any perception that is possible within this time frame involves only the information available in a single glimpse. Random placement of the elements in the displays ensures that attention cannot be prefocused on any particular location. Observers report that these tasks can be completed with very little effort.

A simple example of a preattentive task is the detection of a filled circle in a group of empty circles (Figure 1a). The target object has the visual feature "filled" but the empty distractor objects do not (all non-target objects are considered distractors). A viewer can tell at a glance whether the target is present or absent.

A conjunction target item is one that is made up of two or more features, only one of which is contained in each of the distractors [Triesman, 1985]. Figure 1b shows an example of conjunction search. The target is made up of two features, filled and circular. One of these features is present in each of the distractor objects (filled squares and empty circles). Numerous studies show that the target cannot be preattentively detected, forcing subjects to search serially through the display to find it.

One explicit goal of visualization is to present data to human observers in a way that is informative and meaningful, on the one hand, yet intuitive and effortless on the other. Multidimensional data visualization is concerned with the question "How can we display high-dimensional data elements in a low-dimensional environment, such as on a computer screen or the printed page?" This goal is often pursued by attaching "features" such as hue, intensity, spatial location, and size to each data element. Features are chosen to reveal properties of data elements as well as relationships among them. An ad hoc assignment of features to individual data dimensions may not result in a useful visualization tool. Indeed, too often the tool itself interferes with the viewer's ability to extract the desired information due to

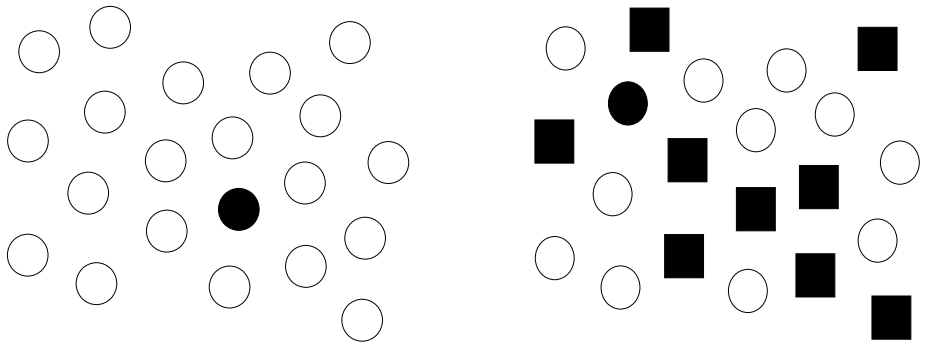

(a) (b)

Figure 1: Examples of two target detection tasks: (a) target can be detected preattentively because it possess the feature "filled"; (b) target cannot be detected preattentively because it has no visual feature that is unique from its distractors.

a poor choice of feature assignment. Direct tests of preattentive vision are often required to determine exactly how different features will interact within a visualization environment.

A good example of this problem is shown in Figure 2. In the first display, hue is used to divide the elements into two groups (*i.e.* a red group and a blue group). Form varies randomly from element to element. Tests show it is easy for subjects to identify the hue boundary as either vertical or horizontal. In the second display, the data-feature mapping has been reversed. Form is used to group the elements, and hue varies randomly across the array. It is much harder for subjects to identify the form boundary in these displays. Moreover, it would be difficult to guess beforehand which data-feature mapping would have provided the best performance. Previous studies in preattentive processing could have been used to predict the outcome.

The visual interference shown in Figure 2 was originally described by Callaghan [Callaghan, 1989]. She found that varying hue within a display region interfered with boundary detection based on form (Figure 2b). It took subjects significantly longer to identify a boundary as horizontal or vertical, relative to a control array where hue was held constant. However, random variation of form did not interfere with hue segregation (Figure 2a). A hue boundary could be identified preattentively, regardless of whether form varied or not. Callaghan found a similar asymmetry in feature interference between intensity and hue [Callaghan, 1984]. Variation of intensity interfered with hue segregation, but variation of hue did not interfere with intensity segregation.

Feature interference is one example of a result that should be carefully considered when designing visualization tools, particularly those which are used to display data regions or boundaries. Studies are currently being conducted to see whether this effect extends to other visualization tasks, such as target detection or numerical estimation.

Figure 3 lists some of the visual features that have been used to perform preat-

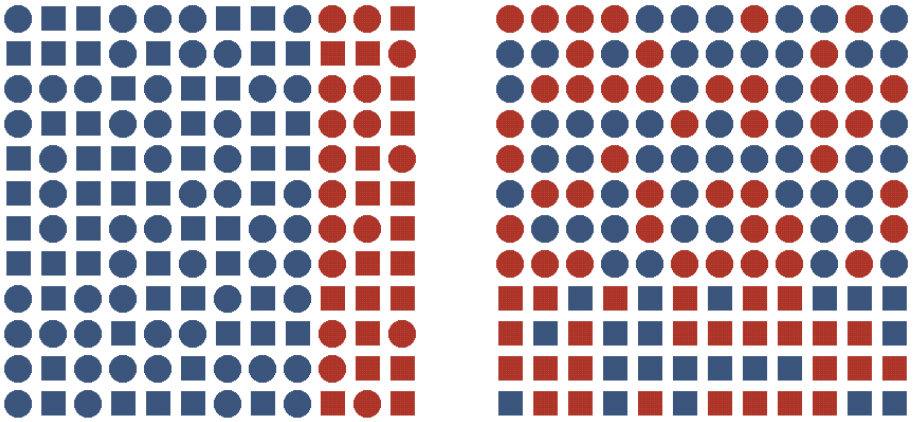

(a) (b)

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

tentive tasks. Literature on preattentive processing describes in detail the various properties of the features, and the tasks for which they can be used. These results can be applied directly to visualization tools that need to perform high-speed target, boundary, or region detection.

A number of scientists have proposed specific theories to explain how preattentive processing occurs. Among the more influential of these have been Triesman's feature integration theory [Triesman, 1985], Julész' texton theory [Julész and Bergen, 1983], Quinlan and Humphreys' similarity theory [Quinlan and Humphreys, 1987], and Wolfe's guided search theory [Wolfe, 1994]. Although the differences in these theories have potentially very important consequences for the actual cognitive operations involved, they need not concern us here, since our interest is in the use of visual features that all of these theories agree can be processed preattentively.

3. SALMON MIGRATION SIMULATIONS

The experimental displays we tested were motivated by the need to examine data generated from salmon migration simulations being conducted in the Department of Oceanography at the University of British Columbia [Thomson et al., 1992; Thomson et al., 1994]. Salmon are a well-known fish that are found, among other areas, on the western coast of Canada. The life history of a salmon consists of four stages: birth, freshwater growth stage, ocean growth stage, and migration and spawning. 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. Upon reaching the ocean, the salmon moves to its "open ocean habitat", where it spends most of its ocean life feeding and growing. For example, 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 Bearing Sea. After a

Feature	Author
line (blob) orientation	Julész & Bergen [1983]; Wolfe [1992]
length	Triesman & Gormican [1988]
width	Julész [1985]
size	Triesman & Gelade [1980]
curvature	Triesman & Gormican [1988]
number	Julész [1985]; Trick & Pylyshyn [1994]
terminators	Julész & Bergen [1983]
intersection	Julész & Bergen [1983]
closure	Enns [1986]; Triesman & Souther [1985]
colour [hue]	Triesman & Gormican [1988]; Nagy & Sanchez [1990]; D'Zmura [1991]
intensity	Beck et al. [1983]; Triesman & Gormican [1988]
flicker	Julész [1971]
direction of motion	Nakayama & Silverman [1986]; Driver & McLeod [1992]
binocular lustre	Wolfe & Franzel [1988]
stereoscopic depth	Nakayama & Silverman [1986]
3-D depth cues	Enns [1990]
lighting direction	Enns [1990]

Figure 3: A list of two-dimensional features that "pop out" during visual search, and a list of authors who describe preattentive tasks performed using the given feature.

period of one to six years, salmon begin their return migration. This consists of an open ocean stage where they swim back to the British Columbia coast and a coastal stage where they swim back to a freshwater stream to spawn. This is almost always the same stream in which they were born. Scientists now know salmon find their stream of birth using smell when they reach the coast.

The ocean phase of salmon migration is not as well understood. It is recognized that it is rapid, well directed, and well timed. Previous work has examined the climate and ocean conditions during migration to see how they affect the position where Fraser River salmon arrive at the British Columbia coast (hereafter point of landfall). The entrance to the Fraser River is located on the southwest coast of B.C., near Vancouver (Figure 4). Generally, if the Gulf of Alaska is warm, salmon will make their point of 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 colder, salmon are distributed further south, making landfall on the west coast of Vancouver Island and approaching the Fraser River primarily via a southern route through the Juan De Fuca Strait. Research is being conducted to determine the factors that drive this interannual variability.

Recent work in plotting ocean currents has provided scientists with a possible explanation for salmon migration patterns. It has been speculated that the variability of ocean currents has impact on where the salmon make their point of landfall. A multi-institutional investigation has been initiated to examine the influences of currents on open ocean return migrations of salmon using the Ocean Surface Cir-

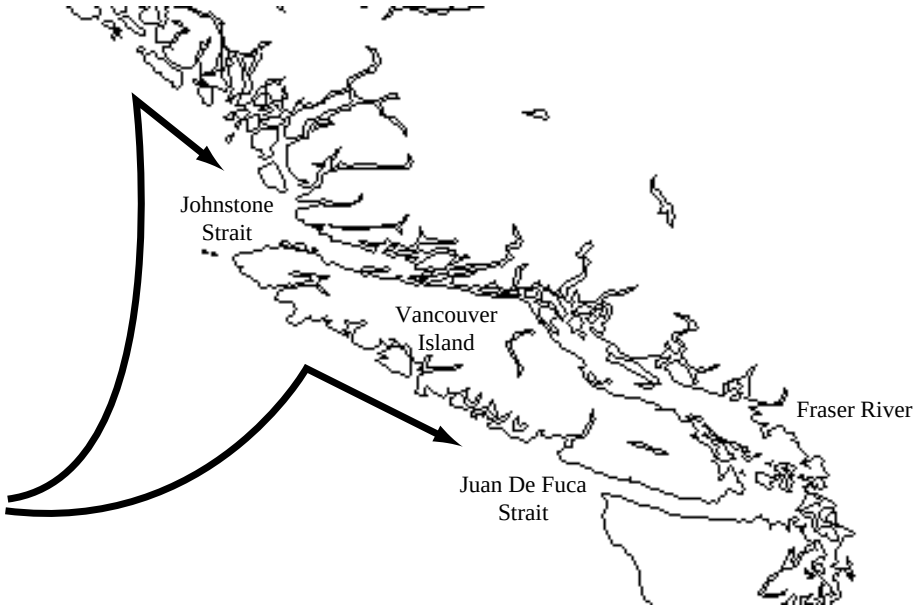

Figure 4: The British Columbia coast. Arrows represent the two possible migration paths for returning salmon.

culation Simulation (OSCURS) model [Ingraham and Miyahara, 1988; Ingraham and Miyahara, 1989]. OSCURS can recreate accurately ocean currents in the North Pacific for any day during the years 1945 to 1990.

Possible return migration paths of the fish are plotted by placing 174 simulated salmon at fixed locations in the open ocean (Figure 5a). The simulated salmon use a "compass-oriented" rule to find their way back to the British Columbia coast. 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, regardless of external forces (*i.e.* ocean currents) that may shift their migration path. OSCURS is used to apply daily ocean currents to the fish as they move towards the coast. Oceanographers record the point of landfall for each salmon (Figure 5b). These are compared with the actual distribution of fish for the given year (provided by the Department of Fisheries and Oceans), to test the accuracy of the simulation.

3.1 Task Selection

We consulted with researchers in Oceanography to identify a suitable task for our psychological experiments. We wanted to choose a task that was common during visualization, but that still addressed a question of potential interest to the oceanographers. Part of this involved using results from the salmon migration simulations during our experiments. Oceanographers hypothesize that a strong northerly current flow will drive most salmon to a point of landfall north of Vancouver Island. It follows that the majority of the fish will pass through the Johnstone Strait to arrive at the Fraser River. This suggests that salmon migration routes can be predicted in

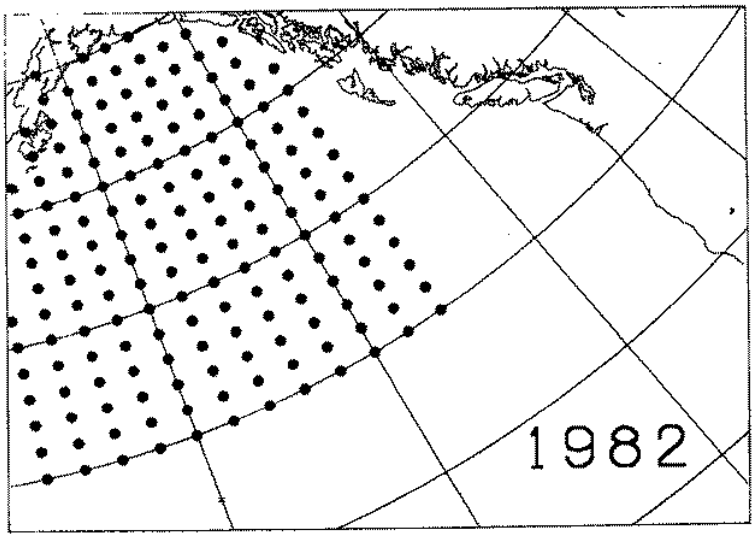

(a)

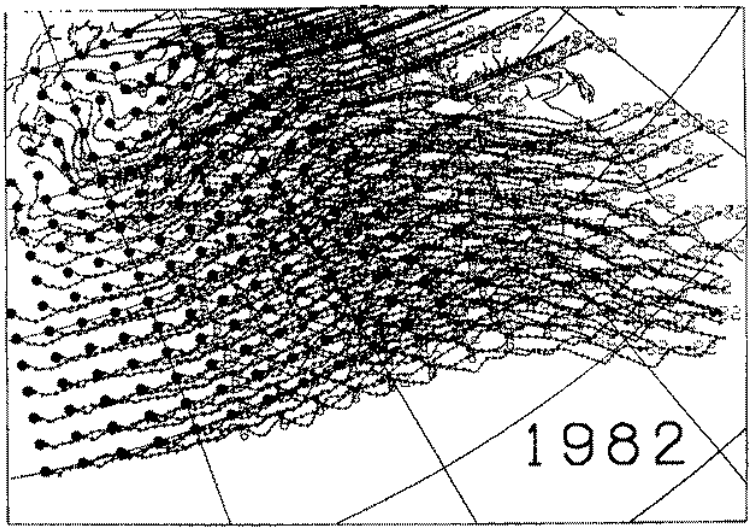

(b)

Figure 5: 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.

part by studying prevailing ocean current patterns. We decided to ask subjects to estimate the percentage of simulated salmon whose point of landfall was north of some fixed latitude, since this was one of the visualization tasks being performed by the oceanographers. Subjects were not informed that the data represented salmon migration results. They were simply asked to estimate the number of elements with a given visual feature.

Relative point of landfall (either north or south of the fixed latitude) was encoded on a two-dimensional map of the open ocean at the spatial position where the salmon started its migration run. A preattentive feature was used to represent relative landfall. For example, during one experiment salmon that landed north of the fixed latitude were coloured blue, while salmon that landed south were coloured red. Subjects were then asked to estimate the percentage of blue elements.

A second question we wanted to investigate was how preattentive features interfered with one another. Callaghan's experiments showed an interference effect during a boundary detection task. We wanted to see if a similar effect existed when subjects were performing estimation. We decided to use stream function as our "interference" attribute, in part because the oceanographers are interested in visualizing these values during their analysis. Given the stream function $\phi(x, y)$, the x and y components of the ocean current vector can be calculated as $V_x = \frac{\partial \phi}{\partial x}$ and $V_y = \frac{\partial \phi}{\partial y}$. The stream function is a scalar function whose gradient is the current (*i.e.* the stream function is the potential function for the vector-valued current field). In our experiments, stream function values were divided into two groups (low and high) and encoded at each spatial location where a salmon started its migration run. A second preattentive feature was used to represent stream function.

4. EXPERIMENT 1: PREATTENTIVE NUMERICAL ESTIMATION

Target detection, boundary detection, and grouping have been studied extensively in the preattentive processing literature [Duncan and Humphreys, 1989; Julész, 1981; Julész and Bergen, 1983; Müller et al., 1990; Triesman, 1985; Triesman and Gormican, 1988]. These results can be extended directly to the visualization problem at hand. Our study sought to go beyond what is known at present by exploring another common task, numerical estimation. We addressed specific instances of two general questions about preattentive features and their use in visualization tools:

- *Question 1*: Is it possible for subjects to rapidly and accurately estimate the percentage of elements in a display within the constraints of preattentive vision? Under what conditions is this possible for the well-studied features of hue and orientation?
- *Question 2*: Does encoding an independent data dimension with a task-irrelevant feature interfere with a subject's estimation ability? If so, which features interfere with one another and which do not?

Although we modelled our experiment task around a question of interest to the oceanographers, we believe estimation is very common during the analysis of vi-

sual displays. If the task is preattentive, visualization tools can be designed which allow users to perform high-speed visual estimation. The visual features we chose, hue and orientation, are commonly used in existing visualization software. Both have been identified as preattentive by Triesman, Julész, and others [Julész and Bergen, 1983; Triesman, 1985]. Moreover, Callaghan's research has shown that hue variation interferes with boundary detection based on form (or orientation). Understanding how hue and orientation interact in a preattentive visualization environment is important. If a visualization tool is being used to display multiple independent data values, interference among features should ideally be eliminated. If a visualization tool is being used to investigate a specific relationship, the "strongest" feature should be used to encode that relationship. Secondary features used to encode additional data values must not interfere with the task-relevant feature.

4.1 Method

We designed experiments that investigated numerical estimation using either hue or orientation. Two unique orientations were used, 0° rotation and 60° rotation. Two different hues which corresponded roughly to red and blue were chosen from the Munsell colour space. This allowed us to display two-dimensional data elements as coloured, oriented rectangles (Figure 6).

We restricted our experiments to two dimensions and two unique values per dimension for a number of reasons. First, focusing on two-dimensional data elements allowed us to test for hue interference effects that had been found by Callaghan in other preattentive tasks. Second, most of the preattentive literature related to our questions has itself been limited to two unique values per feature. This is due in part to the fact that each feature space (in our case, hue and orientation) needs to be divided into values which can be easily distinguished from one another. This is straightforward when only two values are required. True multidimensional visualization tools need to display more than simple two-dimensional data elements. The issues involved in extending our results to more complex datasets are considered in the "Discussion" section of this paper.

For our experiments, we chose a pair of hues that satisfied the following two properties:

- *Property 1*: The perceived intensity of the two hues was equal (*i.e.* the hues were isoluminant).
- *Property 2*: The perceptual discriminability between the two hues was equal to the perceptual discriminability of a rectangle rotated 0° and one rotated 60°.

The method described in Appendix A1 was used to choose a red hue and a blue hue that satisfied these requirements. Our design allowed us to display data elements with two dimensions encoded using hue and orientation. Both dimensions were two-valued (encoded using red and blue or 0° and 60° rotation). This is a very simple example of the general multidimensional visualization problem. In our experiments, we used the following data variable to visual variable mappings:

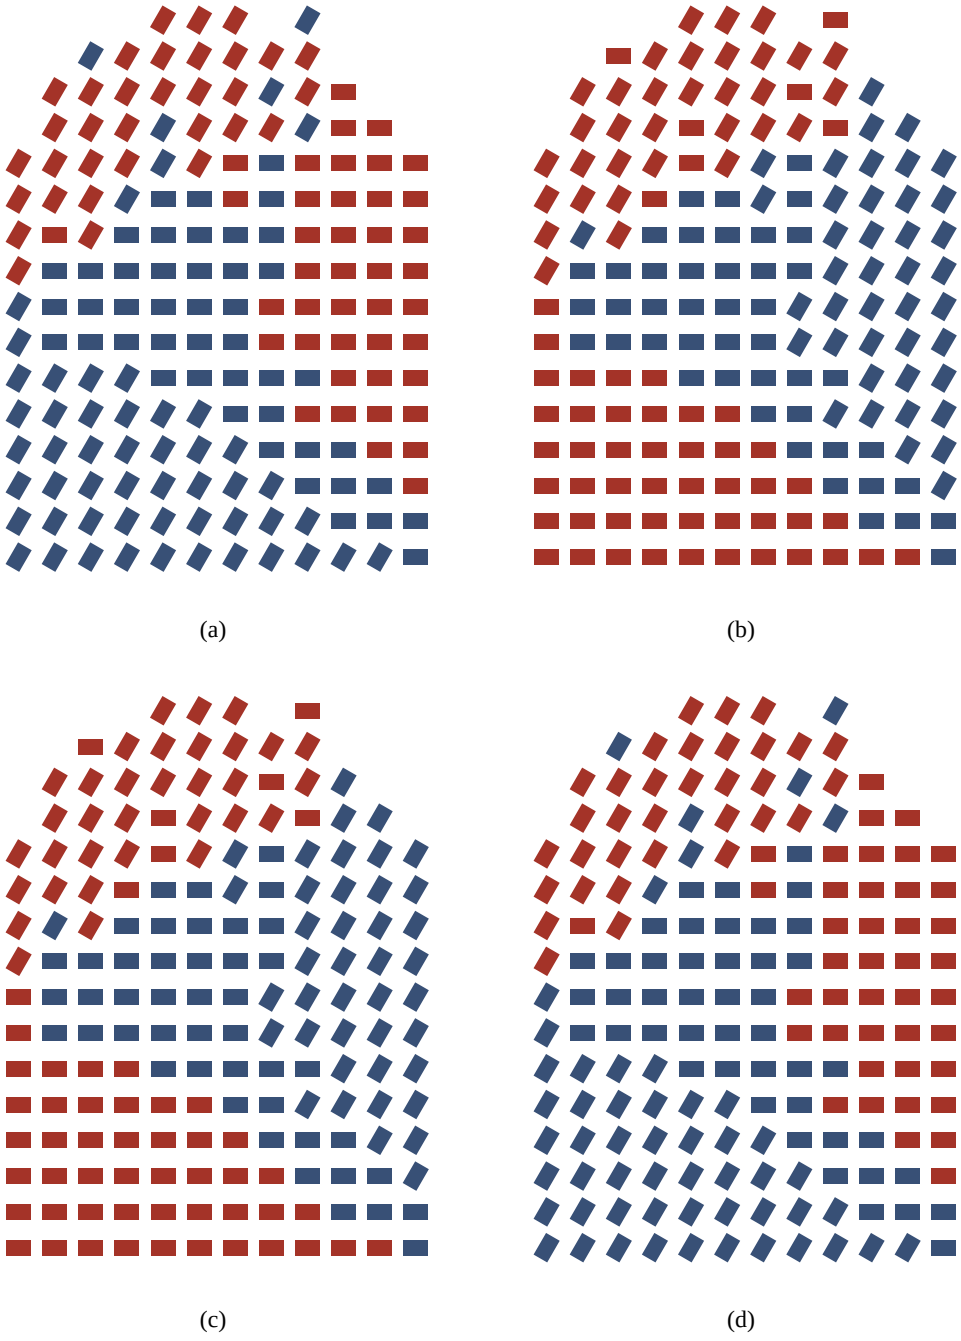

Figure 6: Examples of a single data frame from each of the four mapping conditions (in each frame 58% of the rectangles are targets): (a) Landfall-Hue, user estimated the percentage of elements coloured blue; (b) Landfall-Orientation, user estimated the percentage of elements rotated 60°; (c) Stream Function-Hue, user estimated the percentage of elements coloured blue; (d) Stream Function-Orientation, user estimated the percentage of elements rotated 60°.

- the longitude on which a salmon started its migration run controlled the X-position of the icon representing the salmon
- the latitude on which a salmon started its migration run controlled the Y-position of the icon representing the salmon
- the point of landfall controlled either the hue or the orientation of the icon representing the salmon (the mapping depended on the mapping condition being tested)
- the stream function value controlled either the orientation or the hue of the icon representing the salmon (the mapping depended on the mapping condition being tested)

The experiment was divided into four different mapping conditions run in separate blocks of trials. The task-relevant data dimension (*i.e.* the dimension on which the percentage estimate was based) varied from mapping condition to mapping condition, as did the task-relevant feature (*i.e.* the feature used to encode the task-relevant dimension). This gave us the following design:

- Mapping condition Landfall-Hue: The task-relevant data dimension was point of landfall, represented by hue; the task-irrelevant data dimension was stream function, represented by orientation (Figure 6a).
- Mapping condition Landfall-Orientation: The task-relevant data dimension was point of landfall, represented by orientation; the task-irrelevant data dimension was stream function, represented by hue (Figure 6b).
- Mapping condition Stream Function-Hue: The task-relevant data dimension was stream function, represented by hue; the task-irrelevant data dimension was point of landfall, represented by orientation (Figure 6c).
- Mapping condition Stream Function-Orientation: The task-relevant data dimension was stream function, represented by orientation; the task-irrelevant data dimension was point of landfall, represented by hue (Figure 6d).

Varying the task-relevant data dimension between point of landfall and stream function allowed us to study the effect of different spatial distributions during estimation. Point of landfall values tended to separate into a small number of tightly clustered spatial groups. Stream function values were generally divided into two spatial groups, but with a number of outliers scattered throughout the display. Varying the task-relevant feature allowed us to study the difference between estimation using hue and estimation using orientation. Data were taken directly from the salmon migration simulations. This meant the data values for some trials were modified slightly to meet statistical requirements for the experiment (see Appendix A.2).

For each trial in the experiment, subjects were shown a display similar to those in Figure 6 for 450 milliseconds. The screen was then cleared, and subjects were asked to estimate the percentage of elements in the display with a specified feature, to the nearest 10%. For example, in mapping conditions Landfall-Hue and Stream Function-Hue subjects were asked to estimate the number of rectangles coloured

blue, to the nearest 10%. In mapping conditions Landfall-Orientation and Stream Function-Orientation they were asked to estimate the percentage of rectangles oriented 60°.

Within a mapping condition, trials were divided equally between Constant trials, where the task-irrelevant feature was fixed to a constant value (Figure 7), and Variable trials, where the task-irrelevant feature varied from element to element. Better performance in Constant versus Variable trials would suggest that using a task-irrelevant feature to encode an independent data dimension interferes with estimation based on the task-relevant feature. We tested both for orientation interfering with hue estimation and for hue interfering with orientation estimation.

Twelve subjects with normal or corrected acuity and normal colour vision were tested. The experiments were conducted on a Macintosh II microcomputer equipped with a 13-inch RGB monitor and video hardware capable of displaying 256 colours. The software used was designed and written specifically to run vision experiments [Enns and Rensink, 1991]. Subjects were introduced to each experiment using the method described in Appendix A.3. Subjects then completed two mapping conditions, either Landfall-Hue and Stream Function-Hue (mapping conditions testing hue estimation) or Landfall-Orientation and Stream Function-Orientation (mapping conditions testing orientation estimation). Each mapping condition consisted of 72 Constant trials and 72 Variable trials presented in a random order. Subjects were provided with an opportunity to rest after every 48 trials. Data from the both the introduction and the experiment were saved for later analysis.

4.2 Results

The main dependent variable examined was estimation error, defined as the absolute difference between the subject's interval estimate and the correct interval containing the percentage of target elements present in the display. Statistical analyses using t-tests and analysis of variance (ANOVA) F-tests revealed the following findings:

1. Average estimation error ranged from 0.55 to 0.65 (less than one interval) across the four mapping conditions. Standard deviation of error ranged from 0.26 to 0.36. Subject responses were all clustered close to the correct interval, suggesting that rapid and accurate numerical estimation can be performed under these conditions.
2. Subjects estimated equally well using either hue or orientation, suggesting that there is no subject preference for either feature in the estimation task.
3. Subjects' estimates were more accurate when the task-relevant data dimension was stream function, suggesting that there is evidence of a preference for the spatial arrangement of elements in the estimation task.
4. Accuracy did not differ between Constant and Variable trials during either hue or orientation estimation. Thus, there was no evidence of feature interference in this task.

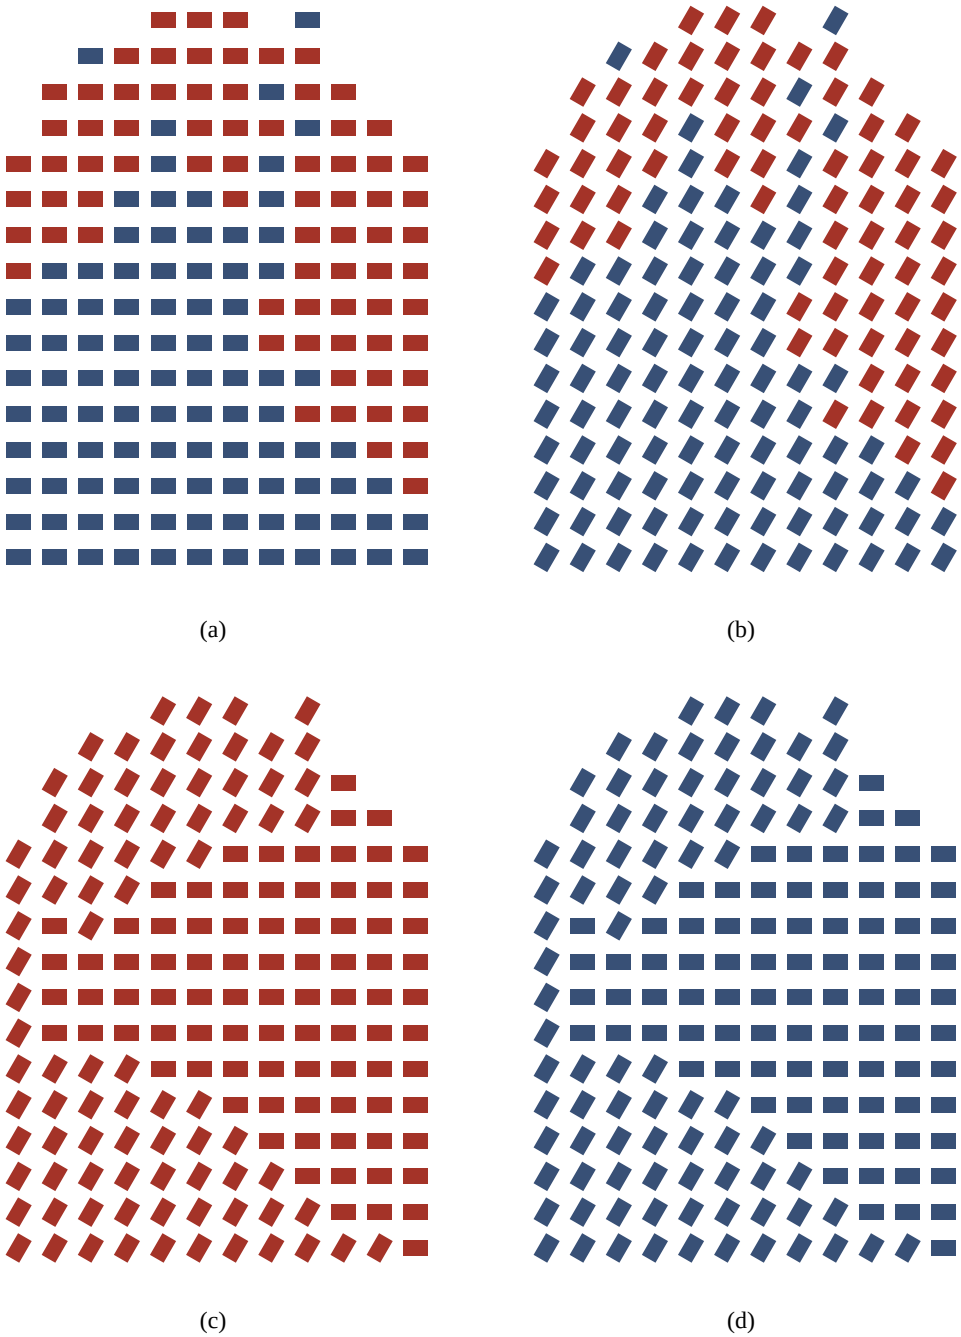

Figure 7: Examples of Constant trials from two of the four mapping conditions (the same orientations and hues are used for Constant trials from mapping conditions Stream Function-Hue and Stream Function-Orientation): (a) Landfall-Hue, stream function represented by constant 0° orientation; (b) Landfall-Hue, stream function represented by constant 60° orientation; (c) Landfall-Orientation, stream function represented by constant red hue; (d) Landfall-Orientation, stream function represented by constant blue hue.

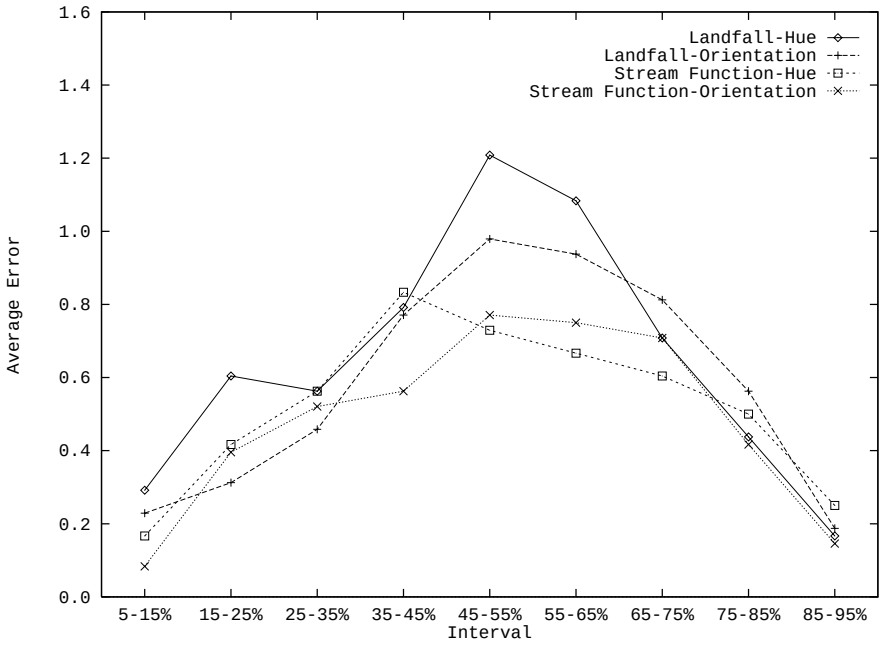

Figure 8: Graph of average error (absolute difference between a subject's interval estimate and the correct interval) as a function of estimation interval for the Constant trials of mapping conditions Landfall-Hue, Landfall-Orientation, Stream Function-Hue, and Stream Function-Orientation. ANOVAs comparing mean error across estimation interval were significant in every mapping condition (all p-values < 0.05); F-values for the Constant trials in Landfall-Hue were $F(8, 423) = 12.48$, $MSe = 0.452$; in Landfall-Orientation they were $F(8, 423) = 11.95$, $MSe = 0.372$; in Stream Function-Hue they were $F(8, 423) = 6.84$, $MSe = 0.334$; and in Stream Function-Orientation they were $F(8, 423) = 8.32$, $MSe = 0.359$.

Figures 8 and 9 graph mean estimation error across nine intervals for the Constant and Variable trials of all four mapping conditions. ANOVAs showed that mean error was affected by the interval being estimated in each condition. The shape of the average error graph was symmetric, with a maximum at the middle intervals and a minimum at either extreme. This phenomenon is often referred to as "end effect" and is simply thought to reflect the fact that, when subjects have less than perfect information, there is greater freedom of choice in guessing at the center of a scale than at either end.

The subjects' ability to perform estimation did not depend on the visual feature begin estimated. Trials were combined into two groups: trials where subjects estimated using orientation (mapping conditions Landfall-Orientation and Stream Function-Orientation), and trials where subjects estimated using hue (mapping conditions Landfall-Hue and Stream Function-Hue). Individual t-tests comparing Constant and Variable trials yielded no evidence of any differences in mean error (all p-values > 0.05; t-values for Constant and Variable trials were $t(1726) = 1.69$ and $t(1726) = 0.59$ respectively). There appears to be no feature preference during estimation.

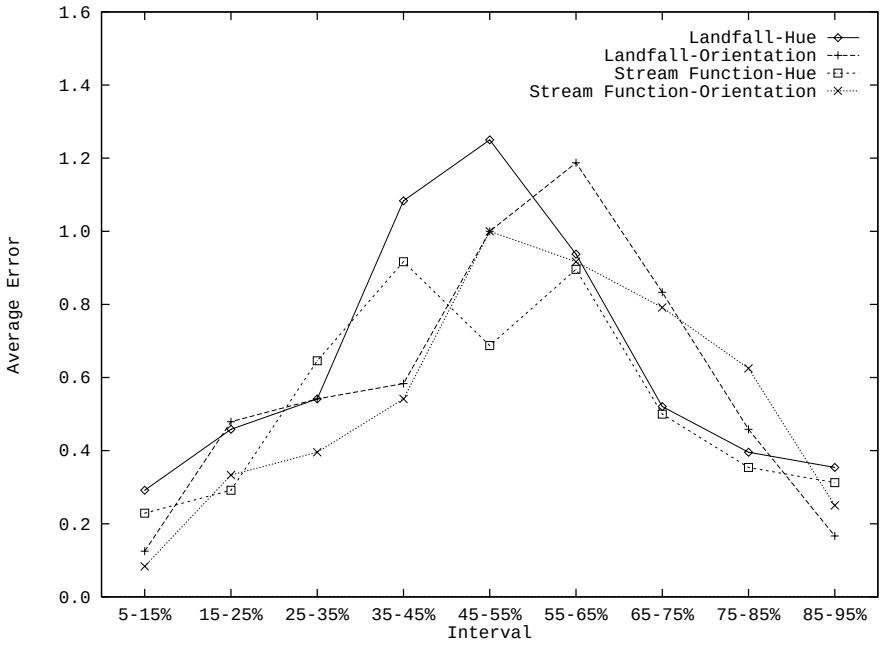

Figure 9: Graph of average error (absolute difference between a subject's interval estimate and the correct interval) as a function of estimation interval for the Variable trials of mapping conditions Landfall-Hue, Landfall-Orientation, Stream Function-Hue, and Stream Function-Orientation. ANOVAs testing mean error across estimation interval were significant in every mapping condition (all p-values < 0.05); F-value for the Variable trials in Landfall-Hue was $F(8, 423) = 12.05$, $MSe = 0.486$; in Landfall-Orientation it was $F(8, 423) = 16.52$, $MSe = 0.367$; in Stream Function-Hue it was $F(8, 423) = 9.45$, $MSe = 0.348$; and in Stream Function-Orientation it was $F(8, 423) = 13.41$, $MSe = 0.349$.

The accuracy of estimation did depend on the spatial distribution of the data points. Estimation was more accurate for trials where stream function (which had a less compact spatial distribution) was the task-relevant dimension. Trials were combined into two groups: trials where point of landfall was the task-relevant dimension (mapping conditions Landfall-Hue and Landfall-Orientation), and trials where stream function was the task-relevant dimension (mapping conditions Stream Function-Hue and Stream Function-Orientation). Our t-tests revealed that accuracy was significantly higher for stream function ($t(1726) = 3.50$, $p < 0.05$ and $t(1726) = 1.84$, $p < 0.10$ for Constant and Variable trials respectively). Additional experiments that explicitly control the change in spatial distribution are needed before we can state specifically its effect on the estimation task.

Finally, we found no significant interference effects in any of the four mapping conditions. Neither feature interfered with the other. To examine this question, we computed t-values to compare mean error across Constant and Variable trials for mapping conditions that used hue as the task-relevant feature. The results showed no evidence of difference in mean error (all p-values > 0.05; $t(862) = 0.04$ and $t(862) = 0.28$ for mapping conditions Landfall-Hue and Stream Function-

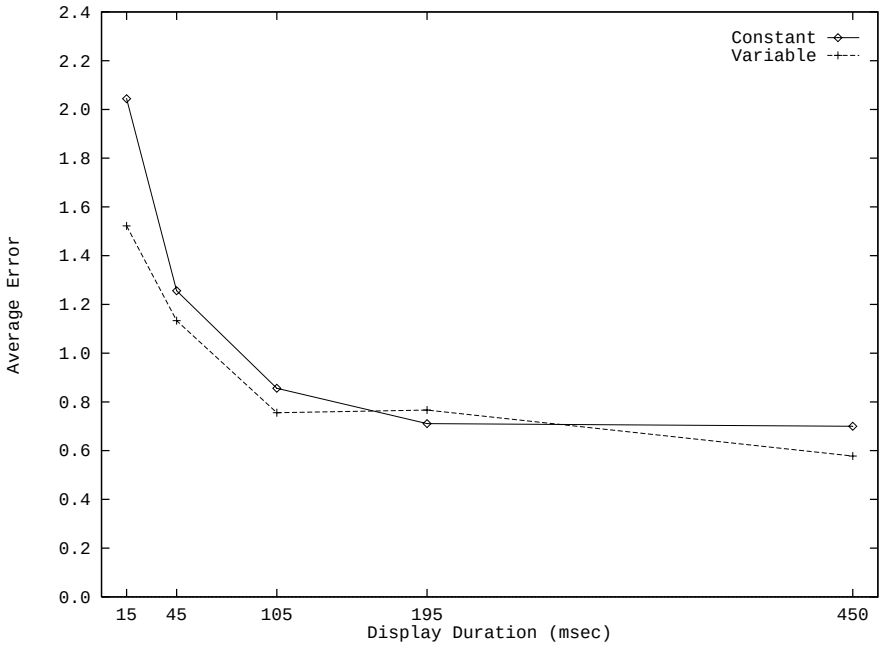

Figure 10: Graph of average error (absolute difference between a subject's interval estimate and the correct interval) as a function of display duration for combined results from the hue display duration experiment. ANOVAs comparing mean error across display duration were significant (all p-values < 0.05) with $F(4, 445) = 19.77$, $MSe = 1.464$ and $F(4, 446) = 13.13$, $MSe = 0.979$ for Constant and Variable trials. Fisher's PLSD tests identified the duration pairs (15,45), (15,105), (15,195), (15,450), (45,105), (45,195) and (45,450) as significant in both the Constant and Variable subsections. t-values comparing mean error between Constant and Variable trials were significant only at the 15 millisecond display duration (all other p-values > 0.05) with $t(178) = 2.18$, $t(178) = 0.76$, $t(178) = 0.69$, $t(178) = 0.40$, and $t(178) = 1.09$ for the display durations 15, 45, 105, 195, and 450 milliseconds.

Hue respectively). Tests for mapping conditions that used orientation as the task-relevant feature yielded similar results (all p-values > 0.05; $t(862) = 0.30$ and $t(862) = 1.47$ for mapping conditions Landfall-Orientation and Stream Function-Orientation respectively).

5. EXPERIMENT 2: DISPLAY DURATION

Our conclusions in Experiment 1 apply to data displayed for 450 milliseconds. This leaves two important questions unanswered. First, at what display duration are subjects no longer able to perform accurate estimation? Second, do any feature interference effects begin to appear at shorter display durations?

In Experiment 2 display duration was randomly varied among five possible values: 15, 45, 105, 195, and 450 milliseconds. Fifteen subjects with normal or corrected acuity and normal colour vision were tested in a manner similar to Experiment 1. Trials were presented to subjects in the following way:

1. A blank screen was displayed for 195 milliseconds.

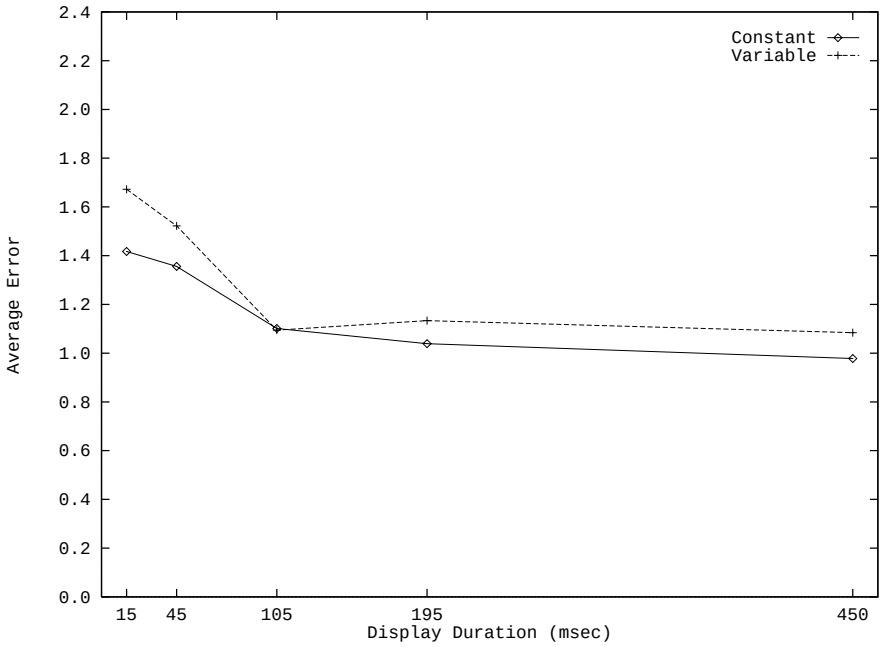

Figure 11: Graph of average error (absolute difference between a subject's interval estimate and the correct interval) as a function of display duration for combined results from the orientation display duration experiment. ANOVAs comparing mean error across display duration were significant (all p-values < 0.05) with $F(4, 894) = 3.65$, $MSe = 1.895$, and $F(4, 894) = 7.54$, $MSe = 1.810$ for Constant and Variable trials. Fisher's PLSD tests identified the duration pairs (15,105), (15,195), (15,450) (45,195) and (45,450) as significant in the Constant subsection, and (15,105), (15,195), (15,450), (45,105), (45,195) and (45,450) as significant in the Variable subsection. t-values comparing mean error between Constant and Variable trials were not significant (all p-values > 0.05) with $t(358) = 1.64$, $t(358) = 1.04$, $t(357) = 0.05$, $t(357) = 0.69$, and $t(358) = 0.83$ for the display durations 15, 45, 105, 195, and 450 milliseconds.

2. A focus circle with diameter roughly twice the width of the rectangular elements was displayed for 105 milliseconds.
3. The trial was shown for its display duration (one of 15, 45, 105, 195, or 450 milliseconds).
4. A "mask" of randomly oriented grey rectangles was displayed for 105 milliseconds.
5. The screen was blanked, and subjects were allowed to enter their estimations.

Five subjects estimated the percentage of elements defined by a blue hue (mapping condition Landfall-Hue), and 10 subjects estimated the percentage of elements defined by a 60° rotation (mapping condition Landfall-Orientation). As in Experiment 1, an equal number of trials was used for each interval (10 Constant and 10 Variable). Trials were split evenly among the five possible display durations, and were presented to the subjects in a random order.

We found that the minimum display duration for robust estimation using either hue or orientation lay somewhere between 45 and 105 milliseconds. Since Exper-

iment 1 had shown that estimation was relatively accurate at all percentage levels, we simplified the dependent measure by averaging error over all nine intervals. The results are shown in Figure 10 for hue estimation and in Figure 11 for orientation estimation. Inspection of these figures shows that estimation accuracy was reasonably stable at all durations of 105 milliseconds and higher. Below that duration, error values increased rapidly. ANOVAs confirmed that accuracy varied reliably with display duration for estimation using either hue (Figure 10) or orientation (Figure 11). Fisher's Protected Least Significant Difference (PLSD) tests were computed to identify display duration pairs with significant differences in average error. As we expected, the statistical significance reflects the higher average error from the 15 and 45 millisecond display duration trials for both hue and orientation estimation.

There was no evidence of feature interference for either estimation based on hue or estimation based on orientation. As in Experiment 1, results suggested that random variation in orientation did not interfere with numerical estimation based on hue. The t-values comparing mean estimation error across Constant and Variable trials had p-values greater than 0.05 at every display duration except 15 milliseconds (Figure 10). Similar results were found when we checked to see if hue interfered with estimation based on orientation (Figure 11).

6. EXPERIMENT 3: FEATURE DIFFERENCE

Experiment 3 was designed to address two additional questions related to numerical estimation. First, how much of a feature difference is necessary to allow accurate estimation? Second, how is this difference affected by display duration? Three mapping conditions were tested using three different hue-orientation pairs during estimation:

1. Mapping condition Small: Rectangles were drawn using two hues 5R 7/8 and 5RP 7/8, and two orientations 0° and 5°
2. Mapping condition Medium: Rectangles were drawn using two hues 5R 7/8 and 10P 7/8, and two orientations 0° and 15°
3. Mapping condition Large: Rectangles were drawn using two hues 5R 7/8 and 5PB 7/8, and two orientations 0° and 60°

The perceptual discriminability between the hues and orientations is smallest in mapping condition Small and largest in mapping condition Large. This latter condition was essentially a replication of the hue and orientation values tested in the previous experiments. Within each mapping condition the discriminability of the two hues and two orientations were calibrated to be roughly equal, following the procedures described in Appendix A.1. Trials within each mapping condition were randomly displayed at two display durations: 45 milliseconds and 195 milliseconds. Otherwise, the details of this experiment were identical to the previous two experiments.

Six subjects estimated percentages based on hue. The target hue was one of 5RP 7/8, 10P 7/8, or 5PB 7/8, depending on which mapping condition a given trial

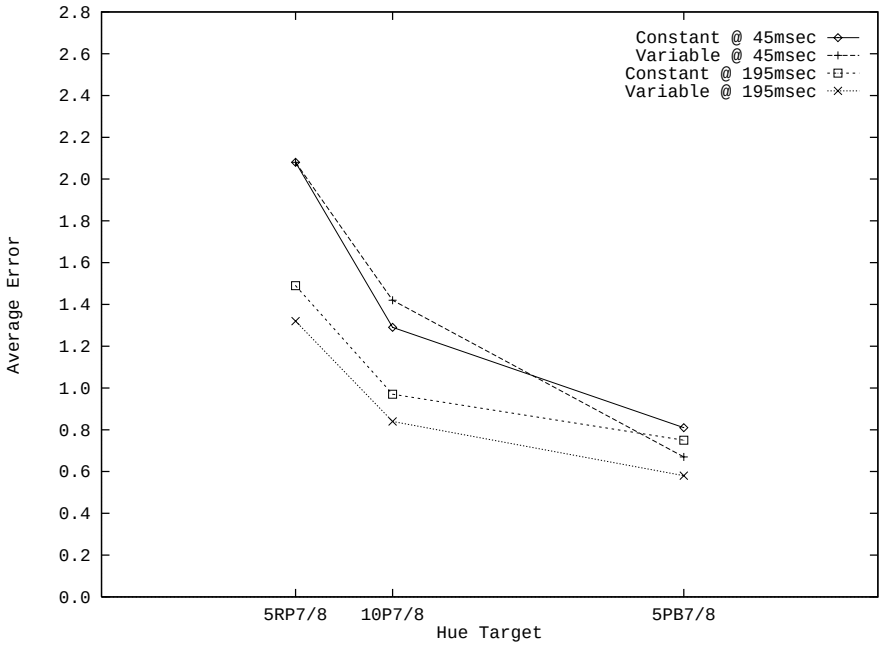

Figure 12: Graph of average error (absolute difference between a subject's interval estimate and the correct interval) across target hue type for combined results from hue feature difference experiment. ANOVAs comparing mean error across hue difference were significant at both exposure durations (all p-values < 0.05); F-values for Constant and Variable trials were $F(2, 645) = 54.16$, $MSe = 1.579$ and $F(2, 639) = 61.26$, $MSe = 1.719$ for 45 millisecond trials; they were $F(2, 642) = 26.63$, $MSe = 1.166$ and $F(2, 641) = 31.36$, $MSe = 0.964$ for 195 millisecond trials.

belonged to. Another six subjects estimated percentages based on orientation (where target orientation was one of 5°, 15°, or 60°). Trials from the three mapping conditions were intermixed and presented to the subjects in a random order.

Subjects were able to perform accurate hue estimation at 195 milliseconds using targets 10P 7/8 and 5PB 7/8, and at 45 milliseconds using 5PB 7/8. Similar results were found for orientation estimation (accurate for targets oriented at 15° and 60° at 195 milliseconds and 60° at 45 milliseconds). Figure 12 graphs mean estimation error for hue trials across the three mapping conditions and both display durations. Figure 13 shows a similar graph for estimation based on orientation. Outside of the above cases, estimation error increased rapidly. ANOVAs confirmed that mean error during hue estimation was dependent on feature difference, with all p-values less than 0.05 (Figure 12). A similar set of ANOVA results was obtained for estimation based on orientation (Figure 13). Average estimation error increased as discriminability between the target and distractors decreased.

Finally, there was no evidence of feature interference during either hue or orientation estimation. The t-values comparing mean estimation error across Constant and Variable trials for all six display duration-target hue pairs were not significant (all p-values were greater than 0.05). Tests for hue interference during orientation estimation were also negative (for all six display duration-target orientation pairs,

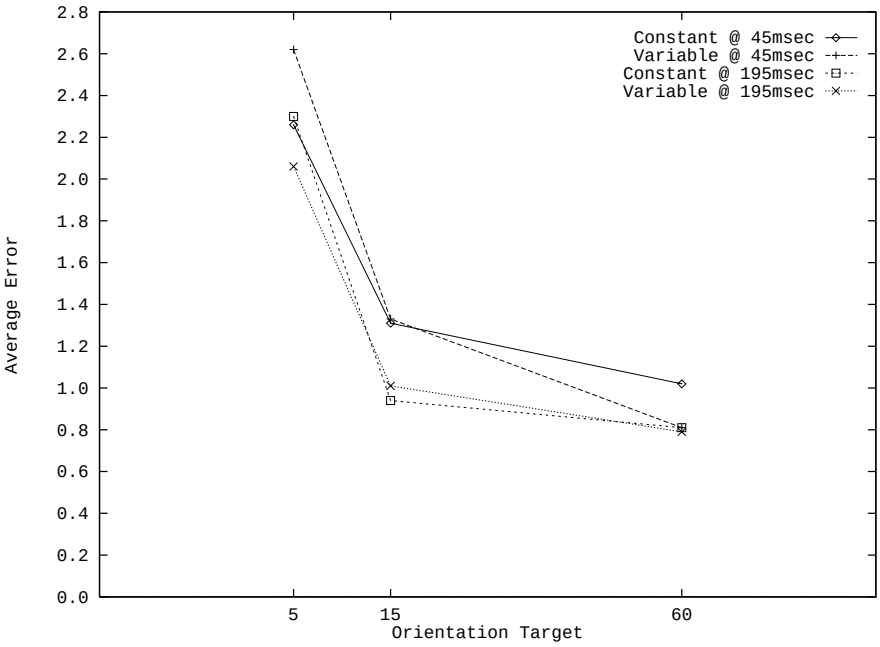

Figure 13: Graph of average error (absolute difference between a subject's interval estimate and the correct interval) across target orientation type for combined results from orientation feature difference experiment. ANOVAs comparing mean error across hue difference were significant at both exposure durations (all p-values < 0.05); F-values for Constant and Variable trials were $F(2, 645) = 40.74$, $MSe = 2.232$ and $F(2, 645) = 79.91$, $MSe = 2.198$ for 45 millisecond trials; they were $F(2, 321) = 77.07$, $MSe = 1.929$ and $F(2, 645) = 50.45$, $MSe = 1.955$ for 195 millisecond trials.

p-values > 0.05).

7. DISCUSSION

The results from Experiment 1 confirm that rapid and accurate numerical estimation can be performed on large multi-element displays using either hue or orientation. Experiment 2 showed further that the estimation task could be performed preattentively, since similar high levels of accuracy were observed down to 105 milliseconds. Both of these results are implicitly assumed in many existing visualization tools, but have not been experimentally verified in the literature. Finally, Experiment 3 showed that within a preattentive display limit of 200 milliseconds, target elements could be reasonably similar to distractor elements (15° rotation or one and a half hue steps in Munsell space) while still allowing for accurate estimation. This provides a new, quantitative guideline for choosing display properties for visualization tools. The absence of any significant interference effects in all of the experiments suggests that it is indeed possible to develop effective multidimensional visualization tools for numerical estimation based on these features.

Before addressing several implications of our findings it is important to consider why users of any visualization tool would want to perform visual numerical esti-

mation at all. The computer could calculate the precise number of values falling within some range. This would, of course, be the easiest solution to the problem if the range were known in advance. The difficulty lies in the exploratory nature of much of the work involving analysis of large datasets. Often the range is not known in advance. We see high-speed preattentive visualization as a very effective tool in helping to formulate and constrain a set of questions that could then be investigated using the power of the computer. To illustrate, consider the problem faced by oceanographers studying salmon migration. They begin their analysis by searching for displays (each display corresponds to migration data for a single year) where more than a fixed percentage of the simulated salmon landed north of Vancouver Island. Such displays are tagged as potential years of interest. This part of the analysis could be automated, although the oceanographers often vary their cutoff percentage while viewing the displays. Years of interest are compared for the percentage and spatial location of low and high stream function values. This often involves estimation on various user-chosen subregions of the display. These subregions could not be pre-chosen for comparison without first identifying and examining the years of interest. Next, the oceanographers go back to the original data and search for displays with a stream function pattern similar to those found in the years of interest. Displays with a similar makeup of stream function values but a different percentage of point of landfall values must be explained in the context of their migration hypothesis. Finally, point of landfall values are compared statistically to historical salmon distributions provided by the Department of Fisheries and Oceans. This provides a computational measure of how well the simulation (and the hypothesis on which it is based) "fits" actual salmon migration patterns. It would not be possible to perform this type of exploratory data analysis by simply pre-programming the computer to search for displays with a user-chosen range of point of landfall and stream function values.

One important concern is how our techniques generalize to other problem environments that need to use high-speed estimation. Our results show that subjects did not use counting to perform numerical estimation, since the most liberal reports of the time needed to count more than 100 elements are 250 milliseconds per element [Trick and Pylyshyn, 1994]. This suggests we should be able to increase the number of elements being displayed while still maintaining accurate estimation. Additional experiments are needed to determine a maximum possible display size under these conditions. This does not necessarily limit our techniques to datasets of this size, however. An efficient method for displaying very large datasets (*e.g.* millions of elements) may not place them all on a single display frame. If the dataset can be divided into subgroups, each subgroup can be displayed one after another in a movie-like fashion. Results from our research have shown that target and boundary detection can be performed on dynamic sequences of 400 element display frames [Healey et al., 1995]. Moreover, interference effects that existed in static displays were also present in this dynamic environment. These findings suggest that other preattentive tasks like estimation might extend to a real-time visualization setting. This would dramatically increase the potential size of the datasets being analysed. For example, a one million element dataset split into 400

element frames and viewed at five frames per second (*i.e.* 200 milliseconds per frame) could be displayed in approximately ten minutes. Even given time for additional inspection of individual frames, the analysis task could be finished in under an hour. Oceanographers are hoping to use this technique to monitor the effects of currents, temperature, and salinity on daily movements of salmon as they migrate back to the B.C. coast. In this problem domain, each data subgroup represents salmon positions, as well as current, temperature, and salinity information for a specific day during the migration run.

Our experiments have not addressed two key questions related to multidimensional visualization. First, how can we encode dimensions that are more than two-valued (*i.e.* truly multi-valued or continuous dimensions)? Second, how can we extend our methods beyond two-dimensional data elements? We are currently studying different techniques to try to address each of these questions.

An obvious method for representing a multi-valued dimension is to divide the corresponding feature space into more than two unique values. Research in preattentive processing has already studied some aspects of this problem. For example, it is easy to detect a tilted rectangle in a field of flat or upright rectangles (or vise-versa). However, in a situation where both the target and distractors are tilted (*e.g.* target is tilted 15° and distractors are tilted 75°), detection is much more difficult. Wolfe suggests orientation might be divisible into only four categories: steep, flat, left, and right [Wolfe et al., 1992]. Another example is the division of hue feature space. Experiments have studied the effect of dividing hue into multiple values. When a target contained a hue that could be separated from the distractors by a straight line in colour space, detection was rapid and accurate. However, when the target contained a hue that was collinear in this space with its distractors, detection was significantly more difficult [D'Zmura, 1991]. This suggests that hue feature space can be divided into multiple values, but that the target (or targets) must be linearly separable from their distractors. We are currently studying both of these effects in the context of our visualization environment.

Tools that support the visualization of multiple data dimensions must deal with a potential interaction between some or all of the features being used to represent the dimensions. Rather than trying to avoid this, we can sometimes control the interaction and use it to our advantage. For example, Pickett and Grinstein have used texture to represent high-dimensional data; each dimension controls one aspect of a texture element displayed to the user. Another promising avenue of investigation involves emergent features. An emergent feature is created by grouping several simpler features together. The emergent feature cannot be predicted by examining the simpler features in isolation. A careful choice of simple features will allow a target element or a region of similar data elements to be detected preattentively [Pomerantz and Pristach, 1989], thereby signalling a correlation of variables in the data.

The data values used in our experiments were derived from salmon migration simulations designed by oceanographers. We are currently applying our techniques to a number of new problem environments, including the visualization of results from spatial database queries and output from medical scanners (*e.g.* CT, MRI, and

PET). We hope this will provide further support for the exploitation of preattentive processing in the design of multidimensional visualization tools.

APPENDIX

This section describes in detail the following three procedures:

1. The method used to choose the two hues that were used during our experiments. Both hues were chosen from the Munsell colour space, and were constrained to meet two experimental properties related to luminance and perceptual discriminability.
2. The method used to modify Oceanography's simulation results to meet statistical requirements for all three experiments.
3. The method used to introduce subjects to each of the three experiments.

A.1 Method for Choosing Hues

For our experiments, we chose a pair of hues from the Munsell colour space that satisfied the following two properties:

- *Property 1*: The perceived intensity of the two hues was equal (*i.e.* the hues were isoluminant).
- *Property 2*: The perceptual discriminability between the two hues was equal to the perceptual discriminability of a rectangle rotated 0° and one rotated 60°.

The Munsell colour space was originally proposed by Albert H. Munsell in 1898 [Birren, 1969]. 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, value, and chroma (Figure 14). Hue refers to ten uniquely identifiable colours such as red, blue, or blue-green. Individual hues are further subdivided into ten subsections. The number before the hue specifies its subsection (*e.g.* 5R, 2B, or 9BG). Value refers to a colour's lightness or darkness. It ranges from one (black) to nine (white). Chroma defines a colour's strength or weakness. Grays are colours with a chroma of zero. A chroma's range depends on the hue and value being used. A Munsell colour is specified by "hue value/chroma". For example, 5R 6/6 would be a relatively strong red, while 5BG 9/2 would be a weak cyan.

One feature of the Munsell colour space is that colours with the same value are isoluminant. Property 1 was satisfied by picking hues that had the same value in Munsell space. We chose value 7, because that slice through Munsell space provided a large number of displayable colours for a variety of different hues.

Property 2 was satisfied by running a set of calibration tests involving a simple target detection task. Subjects were asked to detect the presence or absence

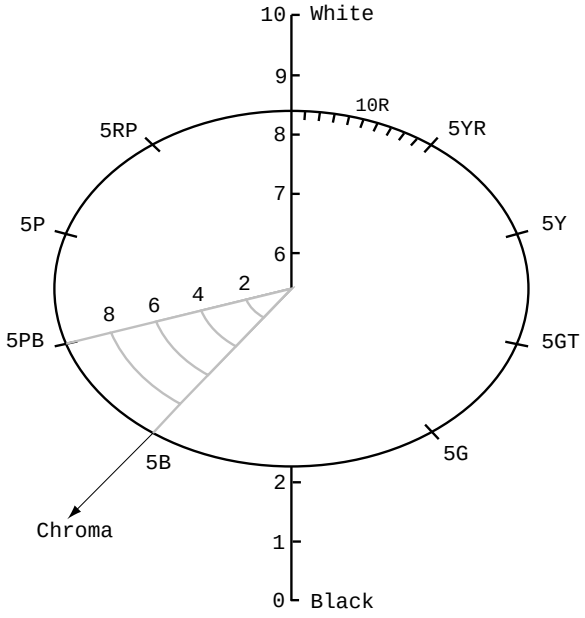

Figure 14: Munsell colour space, showing its three dimensions hue, value, and chroma (from *Munsell: A Grammar of Color*, New York, New York: Van Nostrand Reinhold Company, 1969).

of a rectangle rotated 60° in a field of distractor rectangles rotated 0°. Both the target and distractor rectangles were coloured 5R 7/8. Accuracy in this task was very high (approximately 98%), so the average correct response time ($\overline{RT} = 569$ milliseconds) was used as a measure of discriminability.

The test was then modified to make hue the feature that varied between target and distractor. In the second test, the target was a rectangle coloured 10RP 7/8. The distractors were rectangles coloured 5R 7/8. This made the target half a "hue step" from the distractors in Munsell space. Both the target and distractor rectangles were rotated 0°. The average reaction time for detection was computed from the trials in which subjects responded correctly. The hues used for the target and distractors during the second test were very similar. Because of this, accuracy was lower (approximately 74%) and average correct response times were larger ($\overline{RT} = 943$ milliseconds) than in the test using orientation.

A stepped series of tests were then run with increasingly large hue differences. In each successive test, the target was moved another half "hue step" away from the distractors (*i.e.* 5RP 7/8, 10P 7/8, and so on). This process continued until accuracy and average correct response time were equal to or below that measured for the orientation test. A target coloured Munsell 5PB 7/8 satisfied this criteria (accuracy was 100%, $\overline{RT} = 560$ milliseconds). The result was two isoluminant hues (Munsell 5R 7/8 and Munsell 5PB 7/8) having a perceptual discriminability equal to that of a 60° rotation counter-clockwise from horizontal.

A.2 Method for Modifying Simulation Data

The data values used during our experiments were taken directly from the oceanographers' salmon migration simulations. The data values for some trials were modified to meet certain statistical requirements for the experiment. For example, in the mapping conditions Landfall-Hue and Landfall-Orientation of Experiment 1 point of landfall values were modified to ensure that the following conditions held:

1. An equal number of trials had a given percentage of data elements with a point of landfall value of "north" (*i.e.* 4 trials where 5-15% of the data elements had a value of "north", 4 trials where 15-25% of the data elements had a value of "north", and so on up to 85-95%). This allowed us to compare estimation ability across a uniform range of percentages.
2. Any dependence that might have existed between point of landfall and stream function was removed. This ensured subjects could not infer information about the task-relevant dimension by examining the task-irrelevant dimension.

The stream function data values were modified in a similar way for the mapping conditions Stream Function-Hue and Stream Function-Orientation of Experiment 1. The same modified trials were used during Experiments 2 and 3.

A.3 Method for Introducing the Experiment

The procedure and task of each experiment were introduced to the subjects before they began a testing session. For example, before starting Experiment 1 subjects were shown a sample display frame. The procedure and task of the experiment were described. Subjects were then shown how to enter their estimations. This was done by typing a digit on the keyboard between 1 and 9, which corresponded to the interval they estimated contained the target feature: interval 1 (5-15%), interval 2 (15-25%), and so on up to interval 9 (85-95%). Subjects were told that no trial would contain less than 5% or more than 95% of the target rectangles.

The first phase of the introduction consisted of a set of nine practice 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%. These practice trials were designed to calibrate the subjects' responses and to give them an idea of the experiment procedure and the speed of the trials. 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 18 practice trials presented in a random order, two for each of the nine possible intervals. This phase was designed to run more like the real experiment. Trials in which the subjects estimated incorrectly were not redisplayed; subjects were simply told whether their estimate was right or wrong.

After completing the introduction, subjects begin the actual experiment. A similar explanation and introduction session was given to subjects before Experiments

2 and 3.

ACKNOWLEDGEMENTS

The authors wish to thank K.A. Thomson, W.J. Ingraham, and P.H. LeBlond of the Department of Oceanography at the University of British Columbia for their advice and comments on the application of our work to their salmon migration simulations. The authors also wish to thank W.B. Cowan of the Computer Graphics Laboratory at the University of Waterloo for his advice on colour during the early stages of this work.

REFERENCES

BELL, P. C. AND O'KEEFE, R. M. (1987). Visual interactive simulation—history, recent developments, and major issues. *Simulation 49*, 3, 109–116.

BIRREN, F. (1969). *Munsell: A Grammar of Color*. Van Nostrand Reinhold Company, New York, New York.

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

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

COLLINS, B. M. (1993). Data visualization—has it all been seen before? In *Animation and Scientific Visualization*, Earnshaw, R. and Watson, D., Eds., 3–28. Academic Press, New York, New York.

DUNCAN, J. AND HUMPHREYS, G. W. (1989). Visual search and stimulus similarity. *Psychological Review 96*, 3, 433–458.

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

ENNS, J. T. (1990a). The promise of finding effective geometric codes. In *Proceedings Visualization '90*, 389–390, San Francisco, California.

ENNS, J. T. (1990b). Three-dimensional features that pop out in visual search. In *Visual Search*, Brogan, D., Ed., 37–45. Taylor & Francis, New York, New York.

ENNS, J. T. AND RENSINK, R. A. (1991). VSearch Colour: Full-colour visual search experiments on the Macintosh II. *Behavior Research Methods, Instruments, & Computers 23*, 2, 265–272.

GRINSTEIN, G., PICKETT, R., AND WILLIAMS, M. (1989). EXVIS: An exploratory data visualization environment. In *Proceedings Graphics Interface '89*, 254–261, London, Canada.

HEALEY, C. G., BOOTH, K. S., AND ENNS, J. T. (1995). Real-time multivariate data visualization using preattentive processing. *ACM Transactions on Modeling and Computer Simulation 5*, 3, 190–221.

HIBBARD, B. AND SANTEK, D. (1990). The VIS-5D system for easy interactive visualization. In *Proceedings Visualization '90*, 28–35, San Francisco, California.

HURRION, R. D. (1980). An interactive visual simulation system for industrial management. *European Journal of Operations Research 5*, 86–93.

INGRAHAM, W. J. AND MIYAHARA, R. K. (1988). Ocean surface current simulations in the North Pacific Ocean and Bearing Sea (OSCURS numerical model). Technical Report NMFS F/NWC-130, National Oceanic and Atmospheric Association.

INGRAHAM, W. J. AND MIYAHARA, R. K. (1989). OSCURS numerical model to ocean surface current measurements in the Gulf of Alaska. Technical Report NMFS F/NWC-168, National Oceanic and Atmospheric Association.

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

JULÉSZ, B. AND BERGEN, J. R. (1983). Textons, the fundamental elements in preattentive vision and perception of textures. *The Bell System Technical Journal 62*, 6, 1619–1645.

MCCORMICK, B. H., DEFANTI, T. A., AND BROWN, M. D. (1987). Visualization in scientific computing—a synopsis. *IEEE Computer Graphics & Applications 7*, 7, 61–70.

MÜLLER, H. J., HUMPHREYS, G. W., QUINLAN, P. T., AND RIDDOCH, M. J. (1990). Combined-feature coding in the form domain. In *Visual Search*, Brogan, D., Ed., 47–55. Taylor & Francis, New York, New York.

NAGY, A. L., SANCHEZ, R. R., AND HUGHES, T. C. (1990). Visual search for color differences with foveal and peripheral vision. *Journal of the Optical Society of America A 7*, 1995–2001.

NAKAYAMA, K. AND SILVERMAN, G. H. (1986). Serial and parallel processing of visual feature conjunctions. *Nature 320*, 264–265.

PICKETT, R. AND GRINSTEIN, G. (1988). Iconographic displays for visualizing multidimensional data. In *Proceedings of the 1988 IEEE Conference on Systems, Man, and Cybernetics*, 514–519, Beijing and Shenyang, China.

POMERANTZ, J. AND PRISTACH, E. A. (1989). Emergent features, attention, and perceptual glue in visual form perception. *Journal of Experimental Psychology: Human Perception & Performance 15*, 4, 635–649.

QUINLAN, P. T. AND HUMPHREYS, G. W. (1987). Visual search for targets defined by combinations of color, shape, and size: An examination of task constraints on feature and conjunction searches. *Perception & Psychophysics 41*, 5, 455–472.

ROSENBLUM, L. J. (1994). Research issues in scientific visualization. *IEEE Computer Graphics & Applications 14*, 2, 61–85.

THOMSON, K. A., INGRAHAM, W. J., HEALEY, M. C., LEBLOND, P. H., GROOT, C., AND HEALEY, C. G. (1992). The influence of ocean currents on the latitude of landfall and migration speed of sockeye salmon returning to the Fraser River. *Fisheries Oceanography 1*, 2, 163–179.

THOMSON, K. A., INGRAHAM, W. J., HEALEY, M. C., LEBLOND, P. H., GROOT, C., AND HEALEY, C. G. (1994). Computer simulations of the influence of ocean currents on Fraser River sockeye salmon (*oncorhynchus nerka*) return times. *Canadian Journal of Fisheries and Aquatic Sciences 51*, 2, 441–449.

TRICK, L. AND PYLYSHYN, Z. (1994). Why are small and large numbers enumerated differently? A limited capacity preattentive stage in vision. *Psychology Review 101*, 80–102.

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

TRIESMAN, A. AND GELADE, G. (1980). A feature-integration theory of attention. *Cognitive Psychology 12*, 97–136.

TRIESMAN, A. AND GORMICAN, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review 95*, 1, 15–48.

TRIESMAN, A. AND SOUTHER, J. (1986). Illusory words: The roles of attention and top-down constraints in conjoining letters to form words. *Journal of Experimental Psychology: Human Perception & Performance 14*, 107–141.

TUFTE, E. R. (1990). *Envisioning Information*. Graphics Press, Cheshire, Connecticut.

UPSON, C. (1989). The application visualization system: A computational environment for scientific visualization. *IEEE Computer Graphics & Applications 9*, 4, 30–42.

VANDE WETTERING, M. (1990). apE 2.0. *Pixel 1*, 4, 30–35.

WARE, C. AND BEATTY, J. C. (1988). Using colour dimensions to display data dimensions. *Human Factors 30*, 2, 127–142.

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

WOLFE, J. M. AND FRANZEL, S. L. (1988). Binocularity and visual search. *Perception & Psychophysics 44*, 81–93.

WOLFE, J. M., FRIEDMAN-HILL, S. R., STEWART, M. I., AND O'CONNELL, K. M. (1992). The role of categorization in visual search for orientation. *Journal of Experimental Psychology: Human Perception & Performance 18*, 1, 34–49.