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. It is derived from principles arising in an area of cognitive psychology called preattentive processing. Preattentive processing refers to an initial organization of the human visual system based on 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, boundary, 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 simulations being run in the Department of Oceanography. The results showed that rapid and accurate estimation is indeed possible using either hue or orientation. Furthermore, random variation of one of these features resulted in no interference when subjects estimated the numerosity of the other. To determine the robustness 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 paper describes three experiments that investigate the ability of humans to perform high-speed visual estimation. This work is part of an ongoing study of techniques which allow rapid and accurate visualization of large multidimensional datasets.

Scientific visualization in computer graphics is a relatively new field of research. The term "visualization" was used by a 1987 panel sponsored by the National Science Foundation (NSF) discussing how to apply computer science to data analysis problems (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

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.

(Tufte, 1990) and Collins (Collins, 1993) provide 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 visualization 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 trying to explicitly take advantage of the built-in processing of the human visual system. Preattentive vision refers to those visual operations that can be performed prior to focusing attention on any particular region of an image. We hypothesized that results from research on preattentive processing could be used to assist in the design of visualization tools. Here, we use "visualization tools" to refer to software programs which 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 relative number 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 which can rapidly and accurately perform simple tasks such as target detection or boundary detection?
- Can these results be extended to include rapid and accurate numerical estimation?
- How does changing the parameters of display duration and perceived feature difference affect a user's ability to perform numerical estimation?
- Can results from our controlled 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

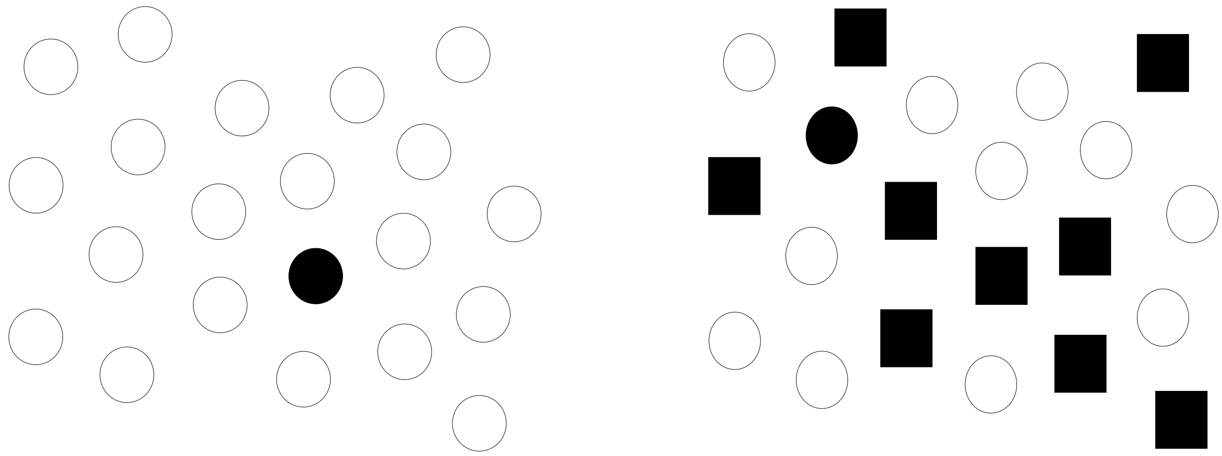

(a) (b)

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

target and boundary detection are important. Results from our 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 found a sharp increase in estimation error when displays were presented for less than 100 milliseconds, when orientation differences were less than 15°, and when hue differences in Munsell space were less than one and a half steps. We used our results to build a tool for visualizing data from salmon migration simulations conducted in the Department of Oceanography. 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

Vision researchers have been working to explain how the human visual system analyses images. One very interesting result 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 which can be performed on large multi-element displays in less than 200 to 250 milliseconds are considered preattentive. This is because eye movements take at least 200 milliseconds to initiate. Random locations of the elements in the displays ensure that attention cannot be prefocused on any particular location. Subjects report that these tasks can be completed with very little effort.

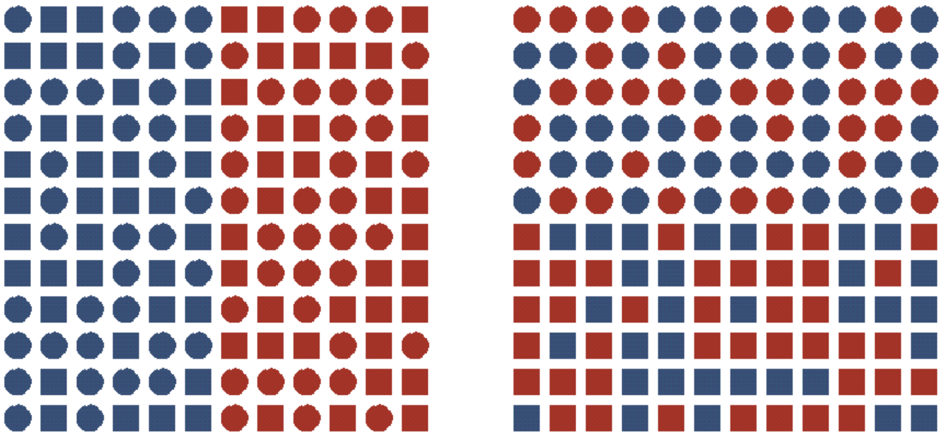

(a) (b)

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

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 poor choice of feature

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)
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)

Fig. 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

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 be used to predict the outcome.

The visual interference discussed above is known as feature domination. It was originally described in research by Callaghan (Callaghan, 1984; Callaghan, 1989). She found that varying hue within a display region interfered with boundary identification 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 domination between intensity and hue (Callaghan, 1984). Variation of intensity interfered with hue segregation. However, variation of hue did not interfere with intensity segregation.

Feature domination is one example of a result which 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 which have been used to perform preattentive 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 which need to perform high-speed target, boundary, or region detection.

A number of scientists have proposed competing theories to explain how preattentive processing occurs, in particular 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 preattentively processed.

3. SALMON MIGRATION SIMULATIONS

The experimental displays we tested were motivated by the need to examine data generated from salmon migration simulations being run in the Department of Oceanography at the University of British Columbia (Thomson et al., 1992; Thomson et al., 1994). Salmon are a well-known type of fish that live, among other areas, on the western Canadian coast. After a period of feeding and growth in the open ocean, 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 reach the coast. The methods used to navigate from the open ocean habitat to the coast are still being researched. Oceanography's simulations are studying the causal effects of ocean currents on sockeye salmon migration patterns. Data representing local ocean current patterns (stream function) and latitudes where each salmon arrived at the B.C. coast (latitude of landfall) were examined during the simulations.

Simulation results are printed on paper using numbers to represent stream function and landfall. Oceanographers perform a number of different analyses on their data. An investigation of their techniques found that many of them corresponded to what appeared to be boundary detection and region size estimation tasks. The oceanographers are interested in using computer-based visualization techniques to browse rapidly through these results. We modeled our experimental task around one of the analysis problems oceanography wanted to study, numerical estimation. This allowed us to increase the likelihood that any application we designed based on our results would be relevant to real-world problems and data.

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 relative number 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?

Estimation is often needed during the analysis of visual displays. If the task is preattentive, visualization tools can be designed which allow users to perform high-speed visual estimation. We decided to examine two preattentive features, hue and orientation. These features are commonly used in existing visualization software. Both hue and orientation have been shown to be preattentive by Triesman, Julész, and others (Julész and Bergen, 1983; Triesman, 1985). Moreover, Callaghan's research has shown that hue exhibits a strong interference effect over form (or orientation) during certain preattentive tasks. 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 must 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 which investigated numerical estimation using either hue or orientation. Two unique orientations were used, 0° rotation and 60° rotation. Two different hues, H_1 and H_2, were chosen from the Munsell colour space. This allowed us to display two-dimensional data elements as coloured, oriented rectangles (Figure 5).

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 which 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 easy to do 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.

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 4). 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. Greys 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.

For our experiments, we chose a pair of hues H_1 and H_2 which satisfied the following two properties:

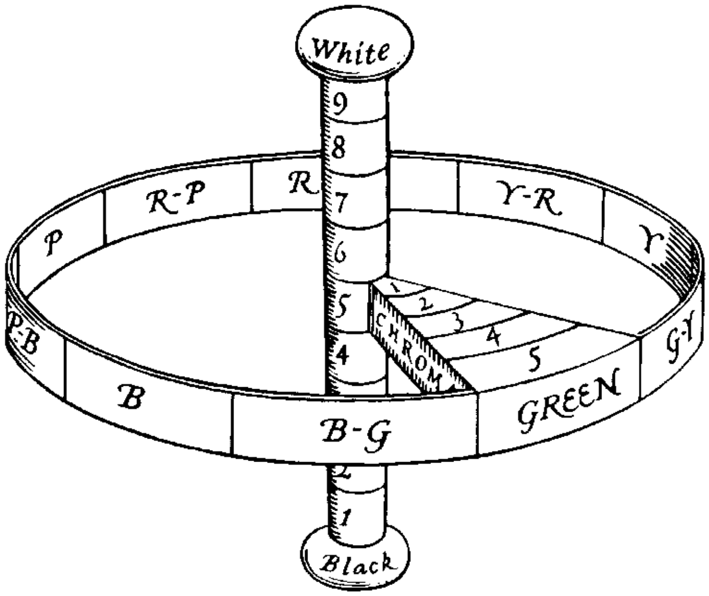

Fig. 4. Munsell colour space, showing its three dimensions hue, value, and chroma (from *Munsell: A Grammar of Color*, New York, New York: Van Nostrand Reinhold Company, 1969)

- *Property 1*: The perceived intensity of H_1 and H_2 was equal (*i.e.* the hues were isoluminant).
- *Property 2*: The perceptual discriminability of H_1 and H_2 was equal to the perceptual discriminability of a rectangle rotated 0° and one rotated 60° (where perceptual discriminability is explained below).

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 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.

Our design allowed us to display data elements with two dimensions D_1 and D_2 (encoded using hue and orientation). Both D_1 and D_2 were two-valued (encoded using H_1 and H_2 or 0° and 60° rotation). This is a very simple example of the general multidimensional visualization problem. In our experiments, D_1 and D_2 corresponded to the landfall and stream function values being studied by the oceanographers.

The experiment was divided into four conditions run in separate blocks of trials: B_1, B_2, B_3, and B_4. The task-relevant data dimension (*i.e.* the dimension the subject was asked to estimate) varied within each block, as did the task-relevant feature (*i.e.* the feature used to encode the task-relevant dimension). This gave us the following design:

- Condition B_1: The task-relevant feature was hue, used to represent landfall; the task-irrelevant feature was orientation, used to represent stream function (Figure 5a).
- Condition B_2: The task-relevant feature was orientation, used to represent landfall; the task-irrelevant feature was hue, used to represent stream function (Figure 5b).
- Condition B_3: The task-relevant feature was hue, used to represent stream function; the task-irrelevant feature was orientation, used to represent landfall (Figure 5c).
- Condition B_4: The task-relevant feature was orientation, used to represent stream function; the task-irrelevant feature was hue, used to represent landfall (Figure 5d).

Varying the task-relevant data dimension between landfall and stream function allowed us to study the effect of different spatial distributions during estimation. Landfall values tended to separate into a small

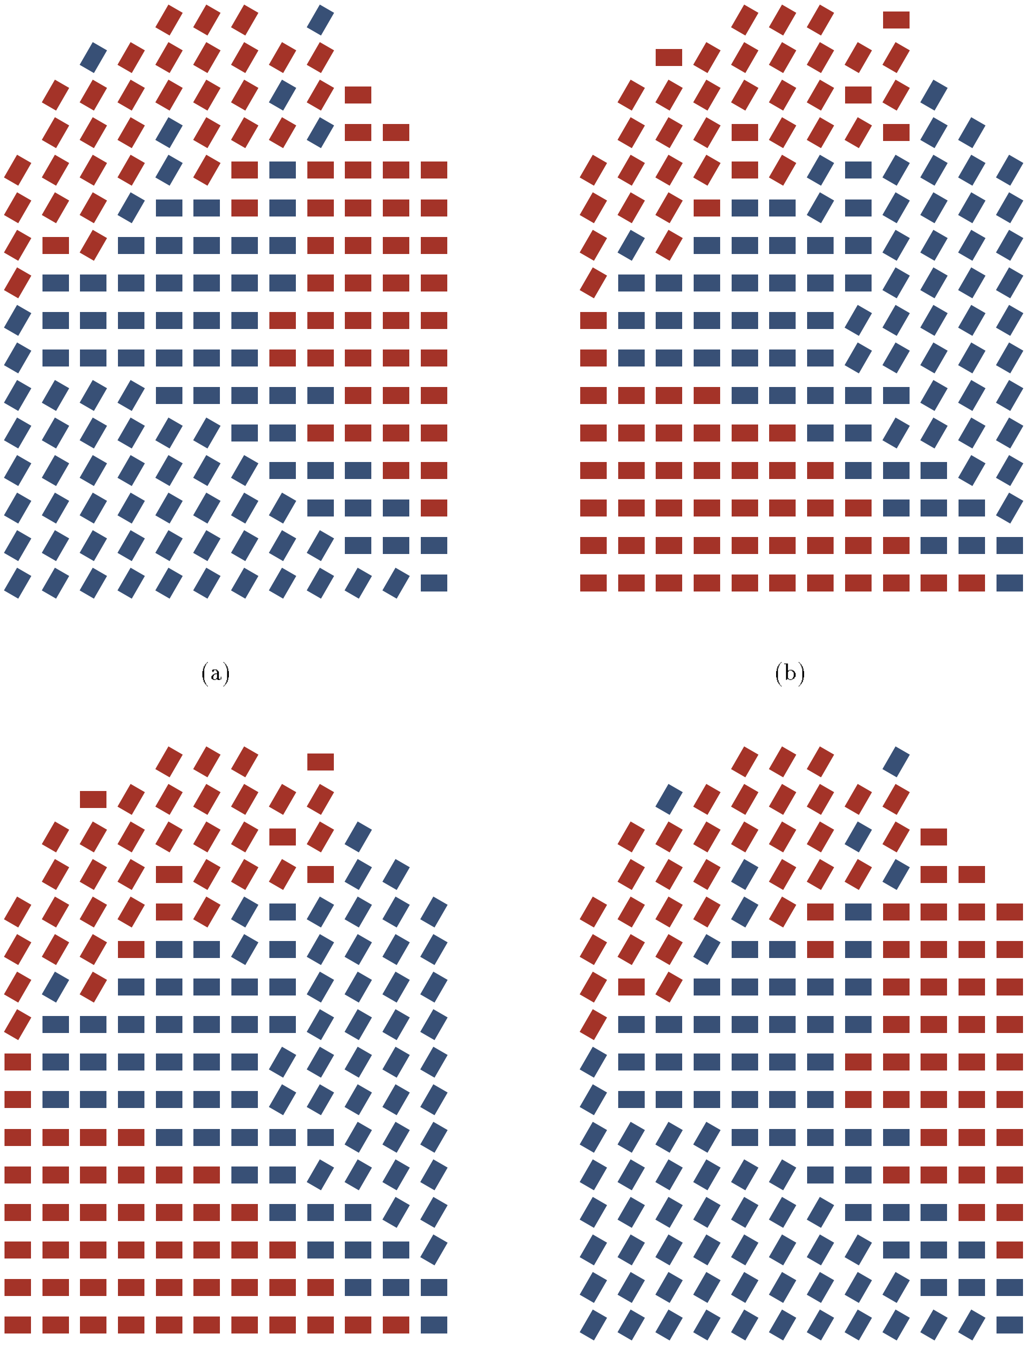

(a) (b) (c) (d)

Fig. 5. Examples of a single data frame from each of the four experimental conditions (in each frame 58% of the rectangles are targets): (a) condition B_1 (landfall represented by hue), user estimates the percentage of elements coloured blue; (b) condition B_2 (landfall represented by orientation), user estimates the percentage of elements rotated 60°; (c) condition B_3 (stream function represented by hue), user estimates the percentage of elements coloured blue; (d) condition B_4 (stream function represented by orientation), user estimates the percentage of elements rotated 60°

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. For example, in conditions B_1 and B_2 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 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 landfall and stream function was removed. This ensured subjects could not infer information about the task-relevant dimension by examining the task-irrelevant dimension.

For each trial in the experiment, subjects were shown a display similar to Figure 5 for 450 milliseconds. The screen was then cleared, and subjects were asked to estimate the number of elements in the display with a specified feature, to the nearest 10%. For example, in conditions B_1 and B_3 subjects were asked to estimate the number of rectangles coloured blue, to the nearest 10%. In conditions B_2 and B_4 they were asked to estimate the number of rectangles oriented 60°.

Within a condition, trials were divided equally between control trials, where the task-irrelevant feature was fixed to a constant value (Figure 6), and experimental trials, where the task-irrelevant feature varied from element to element. Better performance in control versus experimental 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.

At the beginning of the testing session, subjects were shown a sample display frame. The experiment procedure and task were explained. 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 no trial would contain less than 5% or more than 95% of the target rectangles.

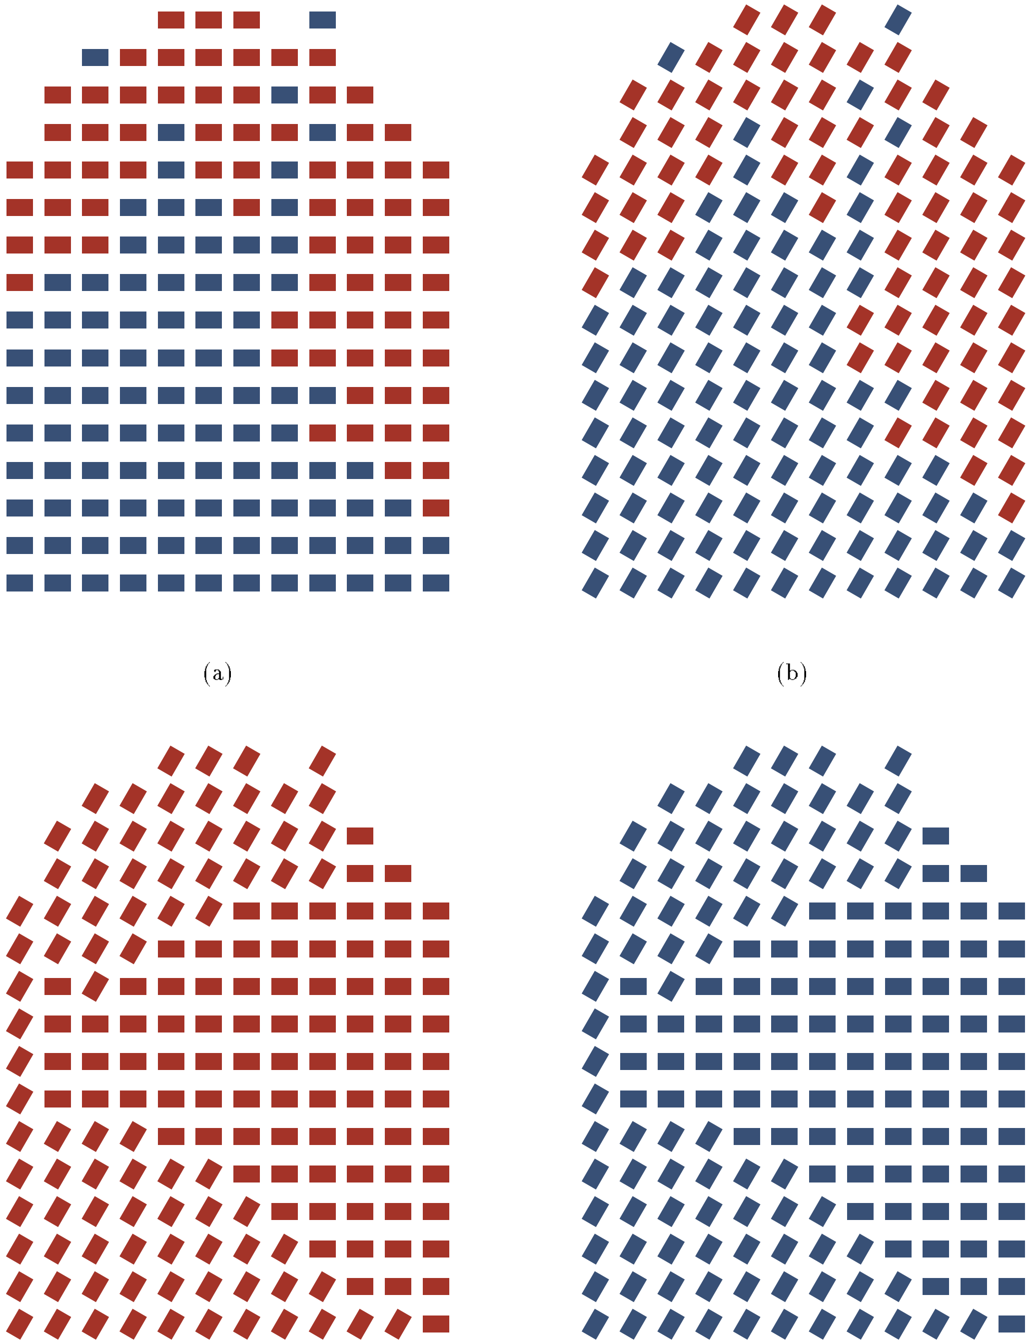

(a) (b)

(c) (d)

Fig. 6. Examples of control displays representing landfall: (a) condition B_1, estimation using hue, stream function represented by constant $0°$ orientation; (b) condition B_1, estimation using hue, stream function represented by constant $60°$ orientation; (c) condition B_3, estimation using orientation, stream function represented by constant red hue; (d) condition B_3, estimation using orientation, stream function represented by constant blue hue

Subjects began the experiment with 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 practice trials. This phase consisted of 18 trials presented in a random order, two for each of the nine possible intervals. This phase was designed to run more like a real experiment block. Trials in which the subjects estimated incorrectly were not redisplayed; subjects were simply told whether their estimate was right or wrong.

Finally, subjects completed two experiment conditions, either B_1 and B_3 (conditions testing hue estimation) or B_2 and B_4 (conditions testing orientation estimation). Each condition consisted of 72 control trials and 72 experimental trials presented in a random order. Subjects were provided with an opportunity to rest after every 48 trials. Data from all four phases were saved for later analysis.

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 by Enns and Rensink to run preattentive psychology experiments (Enns and Rensink, 1991).

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 was less than one interval ($\pm$10%) over all four conditions. Standard deviation of error was also less than one interval. 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 landfall, suggesting that there is evidence of a preference for the spatial arrangement of elements in the estimation task.

4. Accuracy did not differ between control and experimental trials during either hue or orientation estimation. Thus, there was no evidence of feature domination in this task.

Figure 7 graphs mean estimation error across nine intervals for the experimental subsections of all four conditions. ANOVAs showed that mean error was affected by the interval being estimated in each condition (all p-values < 0.05; F-values for control and experimental trials in condition B_1 were $F(8, 207) = 3.49$, $F(8, 207) = 11.0$ and $F(8, 423) = 12.05$; in B_2 they were $F(8, 207) = 6.13$, $F(8, 207) = 5.99$ and $F(8, 423) = 16.52$; in B_3 they were $F(8, 207) = 3.27$, $F(8, 207) = 4.84$ and $F(8, 423) = 9.45$; and in B_4 they were $F(8, 207) = 5.34$, $F(8, 207) = 4.05$ and $F(8, 423) = 13.41$). 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.

We investigated whether a subject's estimation ability differed depending on the feature being estimated. Trials were combined into two groups: trials where subjects estimated using orientation, and trials where subjects estimated using hue. Individual t-tests comparing control and experimental subsections yielded no evidence of any differences in mean error (all p-values > 0.05; t-values for control and experiment subsections were $t(862) = 0.36$, $t(862) = 1.43$ and $t(1726) = 0.45$ respectively). There appears to be no feature preference during estimation.

Spatial distribution of data may affect a subject's estimation ability. For example, it may be easy to perform estimation if the data elements cluster into two distinct groups. It may be more difficult to estimate if the data elements are distributed randomly throughout the display. Subjects estimated using two different data dimensions (landfall and stream function) during the experiment. Both dimensions tended towards their own distinctive spatial patterns. A difference in mean error across the task-relevant dimension would indicate that estimation ability might depend, at least in part, on the spatial distribution of the data being displayed. Trials were combined into two groups: trials where landfall was the task-relevant dimension (conditions B_1 and B_2), and trials where stream function was the task-relevant dimension (conditions B_3 and B_4).

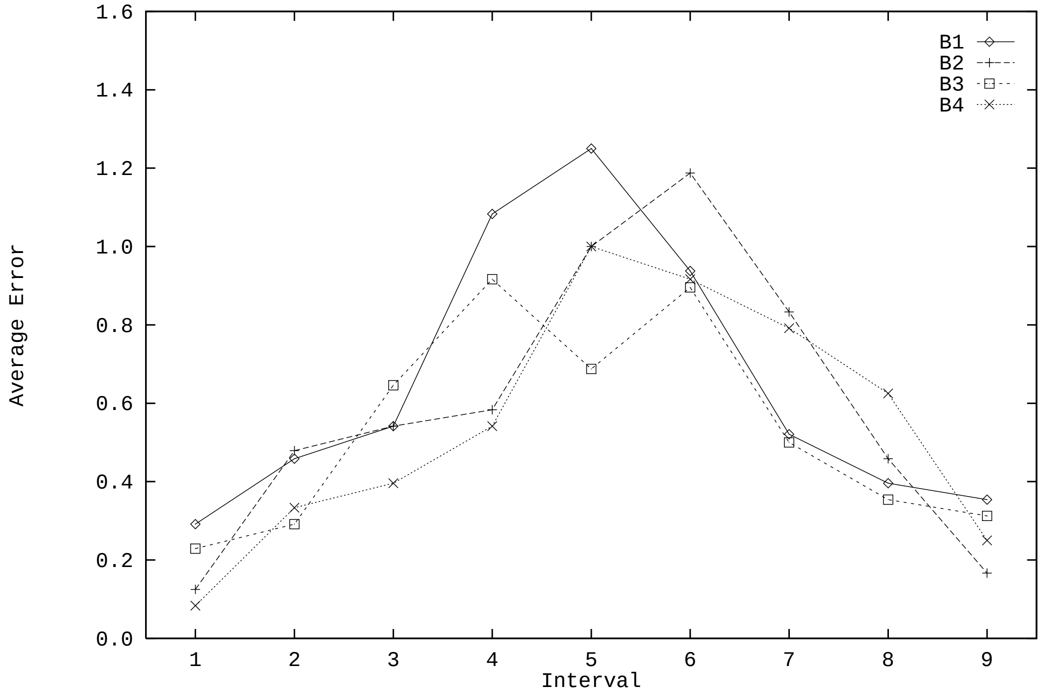

Fig. 7. Graph of average error (absolute difference between a subject's interval estimate and the correct interval) as a function of estimation interval for experimental subsections of conditions B_1, B_2, B_3 and B_4

Our t-tests revealed that accuracy was generally higher for the trials where landfall was the task-relevant dimension ($t(862) = 2.06$, $p < 0.05$, $t(862) = 1.73$, $p < 0.10$, and $t(1726) = 1.84$, $p < 0.10$ for control and experimental subsections 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 examined the question of feature domination during estimation. We found no significant interference effects from the task-irrelevant feature in any of the four conditions. We computed t-values to compare mean error across control and experimental subsections for conditions that used hue as the task-relevant feature. The results showed no evidence of difference in mean error (all p-values > 0.05; $t(862) = 0.03$ and $t(862) = 0.21$ for conditions B_1 and B_3 respectively). Tests for conditions that used orientation as the task-relevant feature yielded similar results (all p-values > 0.05; $t(862) = 0.23$ and $t(862) = 1.15$ for conditions B_2 and B_4 respectively).

5. EXPERIMENT 2: DISPLAY DURATION

Our conclusions in Experiment 1 apply to data displayed for 450 milliseconds. This leaves two important

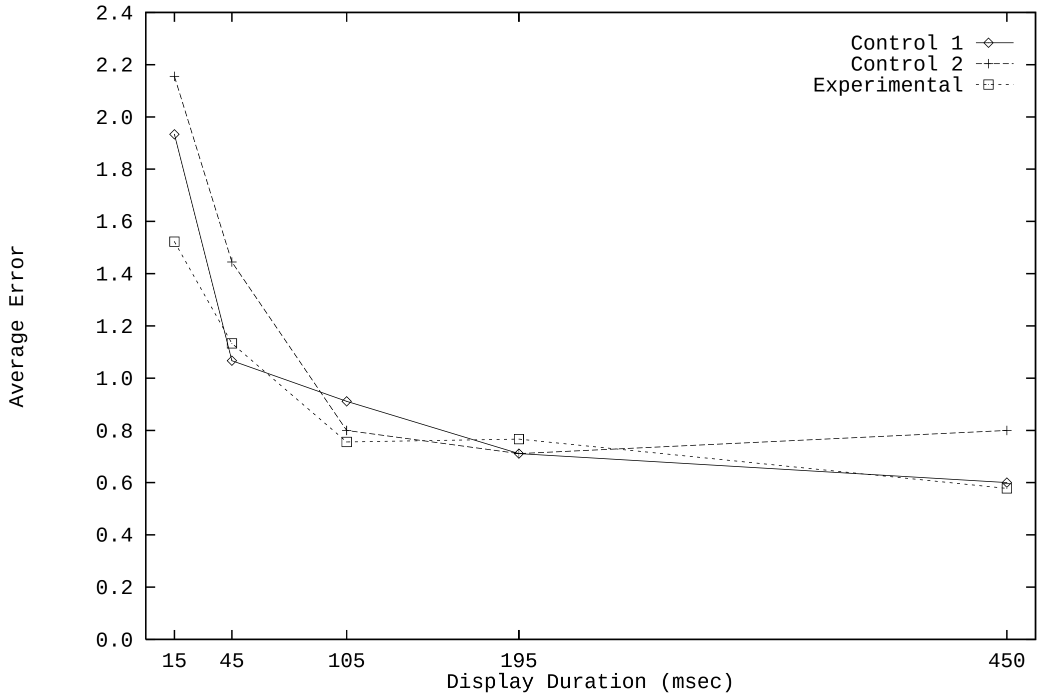

Fig. 8. Graph of average error (absolute difference between a subject's interval estimate and the correct interval) as a function of display duration for combined results from the hue display duration experiment

questions unanswered. First, at what display duration are subjects no longer able to perform accurate estimation? Second, do any feature domination effects begin to appear at lower 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.
2. A focus circle with diameter roughly twice the width of the rectangular elements was displayed for 105 milliseconds.
3. The trial was displayed 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.

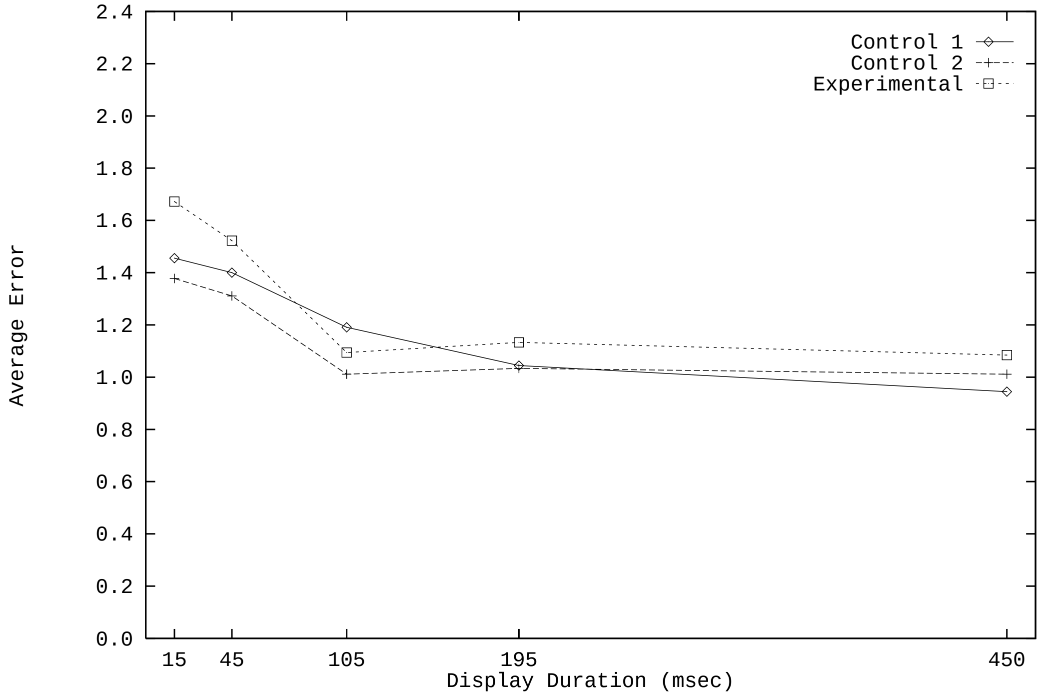

Fig. 9. Graph of average error (absolute difference between a subject's interval estimate and the correct interval) as a function of display duration for combined results from the orientation display duration experiment.

Five subjects estimated the numerosity of elements defined by a blue hue (condition B_1), and 10 subjects estimated the numerosity of elements defined by a 60° rotation (condition B_2). As in Experiment 1, an equal number of trials was used for each interval (10 control and 10 experimental). Trials were split evenly among the five possible display durations, and were presented to the subjects in a random order.

Since Experiment 1 had shown that estimation was relatively accurate at every interval, we simplified the dependent measure by averaging error over all nine intervals. The results are shown in Figure 8 for hue estimation and in Figure 9 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. This indicates that the minimum display duration for robust estimation lies somewhere between 45 and 105 milliseconds.

ANOVAs confirmed that accuracy varied reliably with display duration for estimation using either hue (all p-values < 0.05; $F(4, 221) = 7.53$, $F(4, 221) = 13.59$ and $F(4, 446) = 13.13$ for control and experimental subsections) or orientation ($F(4, 444) = 2.24, p < 0.10$, $F(4, 444) = 1.57, p < 0.20$ and $F(4, 894) =$

7.54, $p < 0.05$ for control and experimental subsections).

Fisher's Protected Least Significant Difference (PLSD) tests were computed to identify display duration pairs with significant differences in average error. For estimation using hue, the duration pairs $(15, 45)$, $(15, 105)$, $(15, 195)$, $(15, 450)$, $(45, 105)$, $(45, 195)$ and $(45, 450)$ were significant in both the control and experimental subsections. As we expected, the significant F-values appear to be due to higher average error from the 15 and 45 millisecond display duration trials.

For estimation using orientation, Fisher's PLSD tests showed that duration pairs $(15, 105)$, $(15, 450)$ and $(45, 450)$ were significant in the control subsections. Duration pairs $(15, 105)$, $(15, 195)$, $(15, 450)$, $(45, 105)$, $(45, 195)$ and $(45, 450)$ were significant in the experimental subsection. Again, the significant F-values appear to be caused by an increase in mean estimation error in the 15 and 45 millisecond trials.

There was no evidence of feature domination during either hue or orientation estimation. As in Experiment 1, results suggested that random variation in orientation did not interfere with numerosity estimates based on hue. The t-values comparing mean estimation error across control and experimental trials had p-values greater than 0.05 at every display duration except 15 milliseconds ($t(178) = 2.18$, $t(178) = 0.76$, $t(178) = 0.69$, $t(178) = 0.40$ and $t(178) = 1.09$ for the display durations 15, 45, 105, 195, and 450 milliseconds). Similar results were found when we checked to see if hue interfered with estimation based on orientation (all p-values > 0.05; $t(358) = 1.64$, $t(358) = 1.04$, $t(357) = 0.05$, $t(357) = 0.69$ and $t(358) = 0.83$ for the five display durations).

6. EXPERIMENT 3: FEATURE DIFFERENCE

Experiment 3 investigated the effects of varying perceived feature differences and display durations. Three conditions were tested using three different hue-orientation pairs during estimation:

1. Condition FD_1: Rectangles were drawn using two hues 5R 7/8 and 5RP 7/8, and two orientations 0° and 5°
2. Condition FD_2: Rectangles were drawn using two hues 5R 7/8 and 10P 7/8, and two orientations 0° and 15°
3. Condition FD_3: Rectangles were drawn using two hues 5R 7/8 and 5PB 7/8, and two orientations 0° and 60°

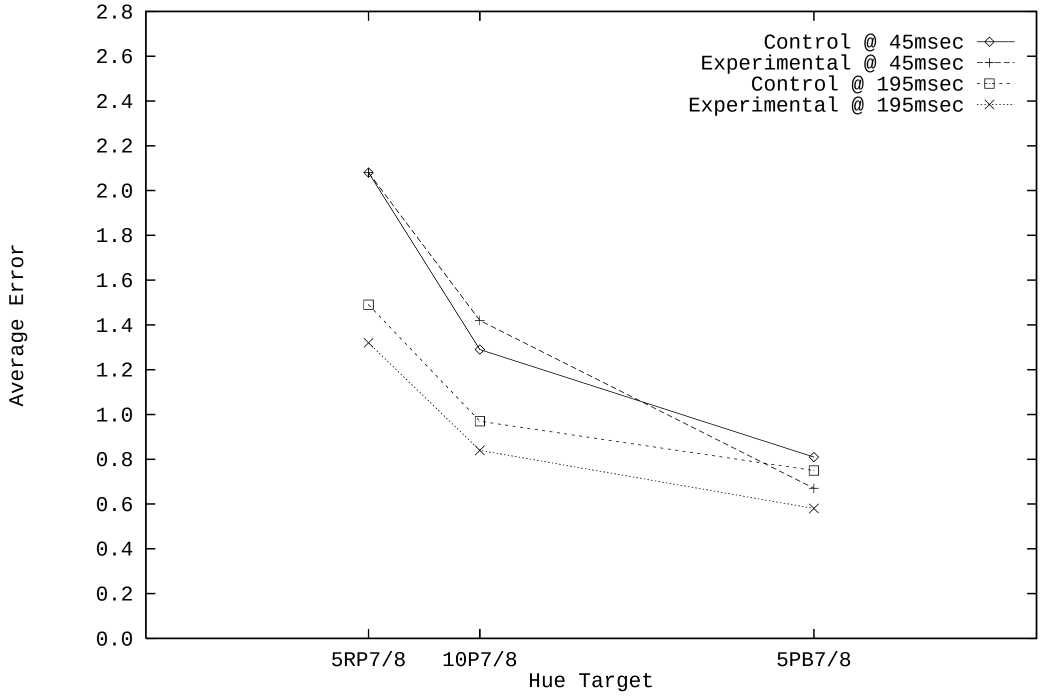

Fig. 10. Graph of average error (absolute difference between a subject's interval estimate and the correct interval) across target hue type for combined results from hue feature difference experiment

The perceptual discriminability between the hues and orientations is smallest in condition FD_1 and largest in condition FD_3, which was essentially a replication of the hues and orientations from the previous experiments. Within each condition the discriminability of the two hues and two orientations had been calibrated to be roughly equal, following the procedures described in Experiment 1. Trials within each 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 numerosity based on hue. The target hue was one of 5RP 7/8, 10P 7/8, or 5PB 7/8, depending on which condition a given trial belonged to. Another six subjects estimated numerosity based on orientation (where target orientation was one of 5°, 15°, or 60°). Trials from the three conditions were intermixed and presented to the subjects in a random order.

Figure 10 graphs mean estimation error for hue trials across the three conditions and both display durations. Figure 11 shows a similar graph for estimation based on orientation. Mean error was below 1.0 for hue estimation using targets 10P 7/8 and 5PB 7/8 at 195 milliseconds, and 5PB 7/8 at 45 milliseconds. For

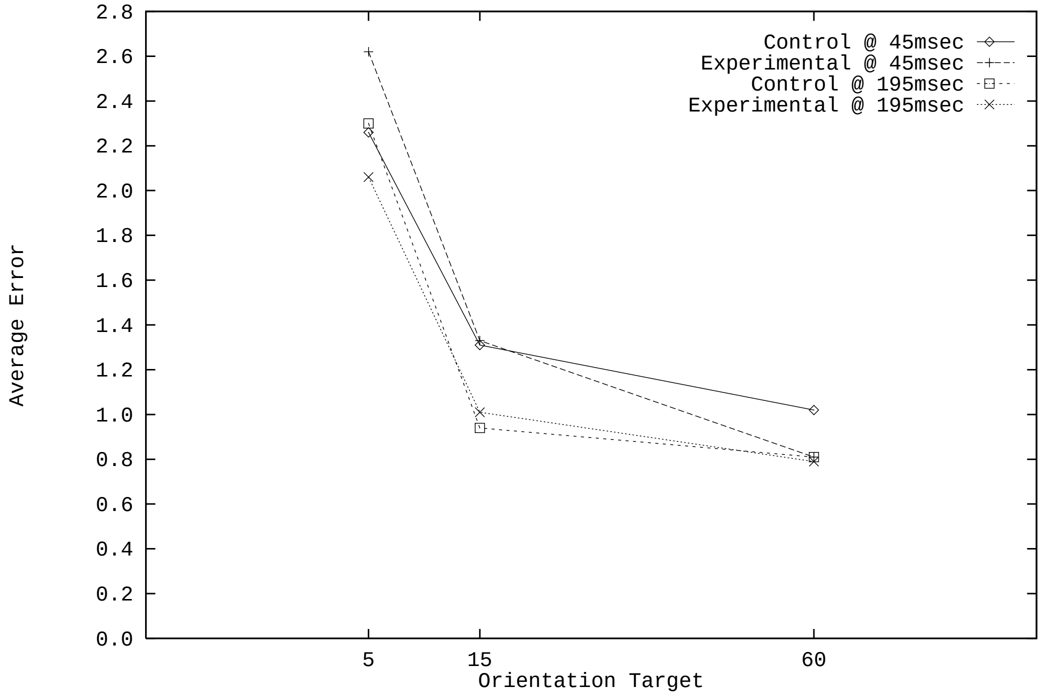

Fig. 11. Graph of average error (absolute difference between a subject's interval estimate and the correct interval) across target hue type for combined results from orientation feature difference experiment

orientation estimation, mean error was below 1.0 for targets oriented at 15° and 60° at 195 milliseconds, and 60° at 45 milliseconds. Outside of these cases, estimation error increased rapidly.

ANOVAs confirmed that mean error during hue estimation was dependent on the target being estimated (all p-values > 0.05; $F(2,321) = 26.92$, $F(2,321) = 30.89$ and $F(2,639) = 61.26$ for 45 millisecond control and experimental subsections; $F(2,319) = 16.42$, $F(2,320) = 12.95$ and $F(2,641) = 31.36$ for 195 millisecond control and experimental subsections). A similar set of ANOVA results was obtained for estimation based on orientation (all p-values > 0.05; $F(2,321) = 20.43$, $F(2,321) = 20.22$ and $F(2,645) = 79.91$ for 45 millisecond control and experimental subsections; $F(2,321) = 45.78$, $F(2,321) = 32.24$ and $F(2,645) = 50.45$ for 195 millisecond control and experimental subsections). Average estimation error increased as discriminability between the target and distractors decreased.

We concluded our analysis by testing for feature domination. There was no evidence that orientation interfered with hue estimation at any display duration. The t-values for all six display duration-target hue pairs had p-values greater than 0.05. Tests for hue interference during orientation estimation were also negative (for all six display duration-target orientation pairs, 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.

These experiments have not addressed two key questions related to multidimensional visualization. First, how can we encode dimensions which 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 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 used a hue which could be separated from the distractors by a straight line in colour space, detection was rapid and accurate. However, when the target used a hue which 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 plan to study both of these effects in the context of our visualization environment.

Tools which 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 Oceanography's salmon migration simulations. 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.

ACKNOWLEDGMENTS

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*, 951–966.

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

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.

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

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

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 and 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.

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 and Applications*, *14*, 2, 61–85.

Thomson, K. A., Ingraham, W. J., Healey, J. 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, J. C., LeBlond, P. H., Groot, C., and Healey, C. G. 1994. The influence of ocean currents on return timing of Fraser River sockeye salmon. *Canadian Journal of Fisheries and Aquatic Sciences*, in press.

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 and 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., Steward, M. I., and O'Connel, K. M. 1992. The role of categorization in visual search for orientation. *Journal of Experimental Psychology: Human Perception & Performance*, *18*, 34–49.

(Triesman and Gelade, 1980; Triesman and Souther, 1986; Nagy et al., 1990; Nakayama and Silverman, 1986; Wolfe and Franzel, 1988)