UC Irvine
UC Irvine Electronic Theses and Dissertations

Title

Deriving Attention Filters from Statistical Summary Representations to Investigate Mechanisms of Feature-Based Attention for Grayscale

Permalink

https://escholarship.org/uc/item/93v3g21z

Author

Yang, Howard

Publication Date

2017

Copyright Information

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Deriving Attention Filters from Statistical Summary Representations to Investigate Mechanisms

of Feature-Based Attention for Grayscale

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Psychology

by

Howard Yang

Dissertation Committee:
Professor Charles Chubb, Chair
Professor Michael D'Zmura
Associate Professor Alyssa Brewer

2017

DEDICATION

To

Megin

You have always been behind me. I am so grateful to have you in my life. I would not be where I am today if not for you.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

CURRICULUM VITAE

Howard Yang

Department of Cognitive Sciences
University of California, Irvine
Irvine, CA 92697
Email: howary1@uci.edu

EDUCATION

2010   B.A. in Psychology, University of California, Los Angeles

2014   M.A. in Psychology, University of California, Irvine

2017   Ph.D. in Psychology, University of California, Irvine

TEACHING EXPERIENCE

2010-2017   Teaching Assistant at the University of California, Irvine for fifteen quarters

TEACHING ASSISTANTSHIPS

Spring 2017   Organizational and Leadership Behavior – Jacqueline Lewis

Winter 2017   Abnormal Psychology – Jacqueline Lewis

Spring 2016   Sports Psychology – Jacqueline Lewis

Fall 2015   Psychology Fundamentals – Donald Hoffman

Spring 2015   Learning and Behavior – Charles Chubb

Winter 2015   Psychology Fundamentals – Thomas D'Zmura

Spring 2014   Social Psychology – Samuel Gilmore

Winter 2014   Writing About Memory – Charles Chubb

Fall 2014   Social Science Lab – Joanne Christopherson

Spring 2013   Personality Theories – Virginia Mann

Winter 2013   Psychology Fundamentals – Charles Wright

Fall 2012   Language and the Brain – Greg Hickok

Spring 2012   Vision – George Sperling

Winter 2012   Human Memory – Christine Lofgren

Fall 2011   Introduction to Psychology – Christine Lofgren

ABSTRACT OF THE DISSERTATION

Deriving Attention Filters from Statistical Summary Representations to Investigate Mechanisms
of Feature-Based Attention for Grayscale

By

Howard Yang

Doctor of Philosophy in Psychology

University of California, Irvine, 2017

Professor Charles Chubb, Chair

An attention filter is a neural process acting across space to modulate the effectiveness with which
a target feature (e.g., blackness) is processed. I investigated the range of attention filters achievable
for grayscale in estimating two summary statistics: (1) centroids and (2) mean orientation (MO)
judgments. Chapter 1: in estimating centroids of dots varying in grayscale, participants can
achieve two types of attention filters: 1) One gives equal weight and 2) one gives weight graded
in proportion to contrast intensity. Comparison of results from different display types suggested
three grayscale centroid filters accessible: (1) one assigns equal weight to all grayscales; (2) one
assigns weight near 0 to negative contrasts, weight proportional to contrast intensity to positive
contrasts; (3) and one assigns weight near 0 to positive Weber contrasts, weight proportional to
contrast intensity to negative contrasts.

Chapter 2 investigated whether vision can access the same grayscale-selective attention filters for
estimating centroids as MO of bars. Each task had three attention conditions: Attend-Light (equal

weight to positive contrasts, weight 0 to negative contrasts), Attend-Dark (equal weight to negative contrasts, weight 0 to positive contrasts), and Attend-All (equal weight to all). Centroid filters matched equal targets well; MO filter accuracy was somewhat worse. Efficiencies, or estimated proportion of items used by an ideal detector, were more than 100% higher for centroids than MO. The pattern of Efficiencies across tasks and attention conditions suggested different strategies between tasks: centroid results suggested participants were applying a filter to the input and computing the centroid of the output; MO results suggested participants were selecting a subsample of bars and basing responses on that sample.

Chapter 3: attention filters giving weight inverse to dot contrast intensity can be achieved. This requires preattentive vision to dampen saliency of the extreme contrast and sharply tune to the low contrast. This implicates a preattentive mechanism sharply tuned to dim items, with discrimination dropping more rapidly with increasing contrast intensity. In selective attention conditions, filters gave high weight to target polarity's lowest contrast dots, but weight near 0 to the distractor polarity. This implicates separate mechanisms responsive either to only positive or negative contrasts.

Chapter 1

Does Human Vision have a Mechanism to Linearly Grade Feature-based Attention to Contrast?

Abstract

An attention filter is a neural process that modulates the strength with which different visual features operate to control performance in a task requiring the participant to extract information about a particular target quality (e.g., blackness) that is distributed broadly across space. This study investigated the attention filters participants achieved for variations in achromatic contrast. Stimuli comprised spatially random clouds of dots varying in Weber contrast. Dark-only (Light-only) displays included only negative (positive) Weber contrasts; Full-set displays included both polarities. In a given condition, the participant strove to mouse-click the centroid, weighting dots in accordance with a specified target filter based on one of two weighting rules: Equal-weight (weight dots equally) or Graded-weight (weight dots in proportion to contrast magnitude). With Dark-only (Light-only) displays, participants strove in separate conditions to apply the two weighting rules to all dots. With full-set displays, participants strove in separate conditions to apply the two weighting rules to (1) all dots, (2) all light dots ignoring dark dots, and (3) all dark dots ignoring light dots. Participants achieved attention filters well-matched to target filters in both Equal-weight and Graded-weight conditions with Light-only and Dark-only displays. Comparison of attention filters achieved with single-polarity vs. Full-set displays suggests that human vision affords direct access to three attention filters selective for gray-scale: one assigns graded weight to light dots and weight 0 to dark; another assigns graded weight to dark dots and

weight 0 to light; the third assigns equal weight to all dots. Other attention filters can be achieved by combining these basic ones.

Introduction

It is well documented that human observers are able to deploy attention broadly across space to select information carried by a specific visual feature (e.g., Baldassi and Verghese, 2005; Haenny et al., 1988; Lankheet and Verstraten, 1995; Ling et al., 2009; Liu et al., 2007; Martinez-Trujillo and Treue, 2004; Maunsell et al., 1991; Muller et al., 2006; Saenz et al., 2003; Serences and Boynton, 2007; Shih and Sperling, 1996; Treue and Martinez-Trujillo, 1999). Any given deployment of such feature-based attention (FBA) aims to heighten sensitivity for a particular visual target property T. A number of basic questions arise concerning any such FBA-deployment, including:

1. How effective is the FBA-deployment in sensitizing the participant to T?

2. Is the FBA-deployment sensitive to information other than T? If so,

    A. which non-target features influence the FBA-deployment? and

    B. how exactly do non-target features influence FBA?

These questions led Drew, Chubb, and Sperling, 2010 to conceptualize FBA in terms of *attention filters*. An attention filter is a process, initiated by a participant in the context of a task requiring FBA, that operates broadly across space to modulate the relative effectiveness with which different features in the retinal input influence performance.

Drew et al. (2010) investigated the attention filters human observers achieved for estimating the centroid of random clouds of dots varying in Weber contrast on a neutral gray background. In different attention conditions, participants strove to mouse-click the centroid of

the dot cloud giving equal weight to 1) all dots (*Attend-to-all-equal*), 2) all dots lighter than the background, ignoring all dots darker than the background (*Attend-to-light-equal*), and 3) all dots darker than the background, ignoring all dots lighter than the background (*Attend-to-dark-equal*). From a participant's performance in a given attention condition, Drew et al. (2010) were able to estimate the attention filter *f(c)* achieved by the participant in that condition. For each Weber contrast *c*, *f(c)* gives the relative weight exerted by a dot of Weber contrast *c* in determining the participant's response.

Participants achieved strikingly different attention filters in the Attend-to-all-equal, Attend-to-light-equal and Attend-to-dark-equal attention conditions. In the Attend-to-all-equal condition, participants achieved attention filters that gave nearly equal weight to all dots. In each of the Attend-to-light-equal and Attend-to-dark-equal conditions, participants achieved attention filters that gave roughly three to four times as much weight to target vs. distractor dots.

In a related study, Drew et al. (2008) investigated whether participants were able to achieve attention filters for centroid estimation that weighted dots in proportion to the absolute value of dot Weber contrast. In this study, participants were tested in the Attend-to-all-equal, Attend-to-light-equal and Attend-to-dark-equal described above. In addition, however, they were tested in three other analogous conditions, the Attend-to-all-graded, Attend-to-light-graded, and Attend-to-dark-graded conditions in which (as in the Attend-to-all-equal, Attend-to-light-equal and Attend-to-dark-equal conditions) the target dots were respectively all dots in the display, just the dots lighter than the background, or just the dots darker than the background. However, instead of striving to give equal weight to all target dots (as in the "equal" conditions), in the graded attention conditions the participant strove to weight target dots in proportion to the absolute value of dot Weber contrast. For example, in the Attend-to-all-graded and Attend-to-dark-graded conditions,

the participant strove to give four times as much weight to dots of Weber contrast -1.0 as compared to dots of Weber contrast -0.25.

For X = ("all," "light," "dark"), Drew, Chubb, Ehrlich, Rubin, & Sperling, (2008) found that participants achieved very similar attention filters in the Attend-to-X-graded conditions as they did in the Attend-to-X-equal conditions. This led them to conclude that observers do not have access to attention filters that provide a graded response to Weber contrast.

The participants in the study of Drew et al. (2008) were given very little practice in each attention condition prior to testing. To subject the conclusions of Drew et al. (2008) to a more stringent test, the current study used an experimental protocol that included with each type of attention (equal, graded): (1) substantial initial, general training in the centroid task, (2) training and testing in conditions using stimuli composed exclusively of dots with only positive (only negative) contrast polarities, and (3) additional training and testing using stimuli composed of dots with mixed positive and negative contrast polarities.

Methods

Observers

The six participants (3 males and 3 females) were either volunteer lab members or paid UCI undergraduates. All methods used were approved by the UC Irvine Institutional Review Board.

Materials & display device

A Viewsonic VGA CRT monitor connected to a Macintosh computer was used to produce displays. The screen size measured 37 cm x 28 cm, screen resolution was 1024x768 pixels and the refresh rate was 85 Hz. The viewing distance of 61 cm was maintained with a chin rest.

Stimuli & trial procedure

The mean luminance of the background (Weber contrast 0) was 32.3 cd/m^2. The stimuli were spatially random clouds of dots varying in Weber contrast. The Weber contrasts of dots were drawn from a set of values approximately equal to -1.0, -0.75, -0.5, -0.25, 0.25, 0.5, 0.75, 1.0. Black dots of Weber contrast $\approx$ -1.0 actually had luminance of 0.3 cd/m^2. White dots of Weber contrast $\approx$ 1.0 had luminance of 64.5 cd/m^2. Each square 7x7 pixel dot subtended 0.21° visual angle. The dots' positions were drawn from a circular bivariate Gaussian density with standard deviation of 100 pixels (3.01° of visual angle). The mean of the Gaussian cloud was also subject to a random perturbation with a circular, bivariate normal distribution with mean equal to the center of the display with standard deviation of 30 pixels (.904° visual angle). The stimulus field within which dots appeared was 512x512 pixels. This region subtended 14.94° of visual angle at the viewing distance of 61 cm and it was circumscribed by a thin, square, black frame centered in an otherwise gray display.

The events that occurred on an experimental trial are indicated in Fig. 1.1. Before each trial, the participant fixated the center of this stimulus field and pressed a button to initiate the trial, which was marked by the onset of the cue frame. The stimulus was then displayed for 200 ms, during which time the border disappeared. After the dots disappeared, the box reappeared along with a small, black, cross-shaped mouse-cursor in the center of the box. The participant used the mouse to move the cursor to click on the response location. Visual feedback was then given; the feedback display showed the dots from the current trial, a bull's-eye composed of three concentric black rings indicating the correct response location, and a small dot indicating the location selected by the participant. The participant then pressed space bar to proceed to the next trial; this allowed him/her to study the feedback as long as desired.

Figure 1.1. The sequence of events in a trial. After the participant initiated the trial with a button-press, they were cued with a thin square black frame marking the boundary within which dots appeared for 1,000 msecs. (a) After it disappeared, the stimulus was presented for 200 ms (b) before disappearing and being replaced by the square frame with a cross-hair cursor in the middle. (c) The participant moved the cursor (indicated here by the arrow) which made the frame disappear; they mouse-clicked on the response location (d) which was followed immediately by (e) feedback comprising a bull's-eye at the location of the correct response, cross-hairs at the location of the participant's response, and the stimulus from that trial. The participant examined this display as long as desired before initiating the next trial.

Experimental conditions

A given experimental condition is marked by (1) the *Weber-contrast set* used to generate the dots in a given stimulus cloud, and (2) the *target filter* used to provide feedback in that condition. Three Weber-contrast sets were used: the "light-only" set comprising Weber contrasts 0.25, 0.5, 0.75 and 1.0; the "dark-only" set comprising Weber contrasts -0.25, -0.5, -0.75 and -1.0; and the "full" set comprising all 8 Weber contrasts in the light-only and dark-only sets.

I. With each of the light-only and dark-only sets, the participant was tested in two attention conditions:

A. Equal-weight: Within a given block, stimuli included either 2 dots of each Weber contrast (8 dots in all) or 4 dots of each Weber contrast (16 dots in all), and the target filter gave equal weight to dots of all Weber contrasts.

B. Graded-weight: 8-dot stimuli included either 2 dots of each Weber contrast, whereas 16-dot stimuli included 4 dots. The target filter weighted each dot in proportion to the absolute value of its Weber contrast.

II. All stimuli that used the full set of Weber contrasts included 2 dots of each Weber contrast. Six different attention conditions used full set stimuli:

1. Attend-to-all-equal: the target filter gave equal weight to dots of all Weber contrasts.

2. Attend-to-all-graded: the target filter weighted dots in proportion to the absolute values of their Weber contrasts.

3. Attend-to-light-equal: the target filter gave equal weight to all dots with positive Weber contrasts and zero weight to all dots with negative Weber contrasts.

4. Attend-to-light-graded: the target filter weighted all dots with positive Weber contrasts in proportion to their Weber contrasts and gave weight zero to all dots with negative Weber contrasts.

5. Attend-to-dark-equal: the target filter gave equal weight to all dots with negative Weber contrasts and zero weight to all dots with positive Weber contrasts.

6. Attend-to-dark-graded: the target filter weighted all dots with negative Weber contrasts in proportion to the absolute values of their Weber contrasts and gave weight zero to all dots with positive Weber contrasts.

Protocol for training and testing

General outline

Every participant went through the following sequence of training and testing phases:

I. Basic centroid training

II. Rule-1 Attention Training and Testing (For 3 of 6 participants, rule-1 used the "equal" weighting rule; for the other 3 participants, rule-1 used the "graded" weighting rule.)

    A. 2-dot Training using Rule-1

    B. Testing using Rule-1 (data is retained from all sessions)

       1. Dark-only

       2. Light-only

       3. Full-set

         (a) Attend-to-all (16 dots)

         (b) Attend-to-light (16 dots)

         (c) Attend-to-dark (16 dots)

III. Rule-2 Attention Training: All stages parallel those for Rule-1, but participants used Rule-2 instead which was the weighting rule not already used.

I. Basic centroid training.   Each participant was first trained to estimate the centroids of dot-clouds comprising solely black dots. The purpose of this training phase was two-fold: (1) to maximally improve the participant's centroid computation (i.e., reduce response error insofar as possible) and thereby (2) to minimize idiosyncratic differences between the centroid computations used by different participants. In this phase of training, on each trial, all dots were black (Weber contrast -1.0). In a given block of 22 trials, 10 trials had clouds that contained 8 dots, 10 trials had clouds that contained 16 dots, and 2 trials had stimuli that consisted of a single dot. Each participant ran at least 132 trials.

Phase II: Rule-1, Attention Training and Testing. Prior to testing in a given attention condition, the participant received centroid task training whose purpose was to optimize performance in that specific attention condition. The stimuli were designed to become progressively more challenging during the course of training and testing.

2-dot training using Rule-1. Prior to testing, in each of the Rule-1 and Rule-2 phases of the experiment, the participant received training in which each display included only two dots. For conditions requiring equal weighting, the Weber contrasts of the two dots in each display were never identical. A block of trials included three trials using each of the 28 possible, non-identical pairs of Weber contrasts for a total of 84 trials per block. On all of these trials, the participant strove to mouse-click the midpoint between the two dots in the display.

The 2-dot training for graded weighting occurred in three separate blocks. These three blocks were run in the following order to allow difficulty to increase gradually. The first block included pairs of dark dots only. The second block included pairs of light dots only. The third block included pairs of dots in which one dot was dark and one dot was light. The first two conditions used Weber contrast pairs that were either identical (4 combinations) or not (6 combinations) for a total of 10 Weber contrast pairs. Each of these 10 Weber contrast pairs was repeated on 10 different trials for a total of 100 trials per block. In each of conditions 1 and 2, participant P6, who was highly experienced and highly skilled in many variants of the centroid task, ran only one block. The other participants all ran at least two blocks. The stopping rule was improvement to $\alpha = 0.05$ and subsequent asymptoting.

The third condition used dot pairs that included one light and one dark dot. All 16 Weber possible combinations were used [(-1.0, -0.75, -0.5, -0.25) x (0.25, 0.5, 0.75, 1.0)]. Each Weber

contrast pair was used in 10 different trials for a total of 160 trials per block. Again, P6 ran one block; all others ran at least two.

Testing using Rule-1: Light-only and dark-only conditions. All participants were tested using the light-only and dark-only Weber-contrast sets before they were tested using the full Weber-contrast set. Although we were interested in comparing the results of the light-only-equal vs. light-only-graded (and dark-only-equal vs. dark-only-graded) attention conditions, a secondary motivation for including these conditions was because it seemed likely that prior experience in these less demanding conditions with only one luminance polarity might improve performance in the various attention conditions using the full set (two polarities of) Weber contrasts. For this reason, all participants were tested first in the single-polarity conditions. In each of the light-only-equal, light-only-graded, dark-only-equal and dark-only-graded conditions, the participant was tested in separate blocks of 100 trials. In each of the light-only-equal, light-only-graded, dark-only-equal and dark-only-graded conditions with 8-dot displays and with 16-dot displays (8 conditions in all), all participants except P6 ran at least 3 blocks; P6 (highly experienced) ran only one block in each condition.

Testing using Rule-1: Full-set conditions. Following testing in the single-polarity conditions, the participant was tested on the full Weber-contrast set. The full-set condition blocks were run in the following order: (i) attend-to-all, (ii) attend-to-light, and (iii) attend-to-dark. Blocks comprised 100 trials. All participants except P6 ran at least two blocks in all conditions. P6 ran two blocks in all conditions except the Attend-to-light-graded, Attend-to-light-equal, and Attend-to-dark-graded conditions in which he ran only one block.

Data analysis. In each attention condition, the participant strove to click on the centroid of the dots in the display weighting the dots in accordance with just one particular target filter f_{targ}. (In

the Attend-to-dark-graded condition, for example, $f_{targ}(w) = -w$ for all Weber contrasts $w < 0$, and $f_{targ}(w) = 0$ for all $w > 0$.) Thus, on a trial comprising n dots in which w_i is the Weber contrast of the i^{th} dot in the display and x_i and y_i are the x- and y-coordinates of its location, the target location ($Targ_x$, $Targ_y$) of the centroid is defined by

$$Targ_X = \frac{1}{W}\sum_{i=1}^{n} f_{targ}(w_i)x_i \quad \text{and} \quad Targ_Y = \frac{1}{W}\sum_{i=1}^{n} f_{targ}(w_i)y_i \qquad (2)$$

W is the sum of $f_{targ}(w_i)$ taken over all dots $i = 1,2,...,n$ in the display.

Typically, however, the participant cannot achieve this goal. We use a simple linear model (Equations 3) that well-approximates the x- and y-coordinates of the participant's response (R_X and R_Y):

$$R_X = \frac{D}{W}\sum_{i=1}^{n} f(w_i)x_i + (1-D)x_{default} + Noise_X \qquad (3a) \quad \text{and}$$

$$R_Y = \frac{D}{W}\sum_{i=1}^{n} f(w_i)y_i + (1-D)y_{default} + Noise_Y \qquad (3b)$$

(Sun, Chubb, Wright, and Sperling, 2016). In Eqs. 3, $x_{default}$ and $y_{default}$ are the x- and y-coordinates of the location toward which the participant's response is assumed to revert if he/she extracts only partial information on a given trial, D (Data-Drivenness) is a real number between 0 and 1, f is a real-valued function of Weber contrast, W is the sum of $f(w_i)$ over all items i in the display, and $Noise_X$ and $Noise_Y$ are normally distributed random variables with mean 0 and standard deviation σ.

The function f is called the *attention filter* achieved by the participant in the given condition. For the set of Weber contrasts used to define the dots in a given attention condition, the values $f(w)$, are called the filter weights of Weber contrasts. It should be noted that f is only defined up to an arbitrary multiplicative constant. For plotting purposes, we impose the constraint that the filter weights $f(w)$ must sum to 1.

The parameter D is the *Data Drivenness* of participant's response. If $D = 1$, then the participant's response on a given trial is determined exclusively by the items in the stimulus cloud and random noise. At the other extreme, if $D = 0$, then the participant's response on a given trial is influenced not at all by the stimulus but instead by the combined effects of random noise plus a tendency to click on a fixed location ($x_{default}$, $y_{default}$).

Estimating model parameters. To estimate the model parameters $x_{default}$, $y_{default}$, $f(c)$, D and σ for a participant in an attention condition in which the dots in the displays are all drawn from a set C = $\{w_1, w_2, \ldots, w_N\}$ of Weber contrasts, we proceed as follows. First draw from the data those trials in which the display was "fully-loaded," i.e., the displays in which the full complement of Weber contrasts was used. In the dark-only conditions, C contained the four Weber contrasts -1.0, -0.75, -0.5, -0.25 whereas in the light-only conditions, C contained Weber contrasts 1.0, 0.75, 0.5, 0.25. In the blocks from dark-only and light-only conditions that used 8-dot (16-dot) displays, the fully-loaded displays included 2 dots (4 dots) of each of the four Weber contrasts in C. In the full-set conditions, C was the set of all 8 Weber contrasts ±1.0, ±0.75, ±0.5, ±0.25, and these fully-loaded displays contained 2 dots with each of these eight Weber contrasts.

(Appendix A describes the procedures for estimating the model parameters.)

Results

Figure 1.4 shows model fits for all 6 participants in all of the Dark-only and Light-only (single-polarity) conditions. The upper (lower) 6 panels show the results for the conditions using 8-dot (16-dot) displays. A given panel contains the results from 4 separately blocked conditions. On the left (right) side of a given panel, the two plotted lines show the attention filters achieved by the participant in the Dark-only (Light-only) conditions; the black (white) line shows the attention filter achieved in the Equal-weight (Graded-weight) condition. The thin black line with square markers (white line with triangular markers) shows the target filter for the condition using the Equal (Graded) weighting rule. On each side of each panel the upper numbers give the Efficiency characterizing the attention filter achieved in the Equal-weight condition (in black) and the Graded-weight condition (in white), and the lower two numbers give the corresponding Data-Drivenness values. The mean Data-Drivenness values of the 6 participants are shown in Fig. 1.2. The most important facts to note about the results plotted in Fig. 1.2 are:

1. All six participants achieve clearly distinct attention filters in the Equal-weight vs. Graded-weight conditions with both the Dark-only and Light-only displays.

2. The attention filters achieved with 16-dot displays are very similar in form to the corresponding attention filters achieved with 8-dot displays. This is also reflected by the average attention filters shown in Fig. 1.2c.
   It should be noted that

3. The Equal-weight filters are generally flat for all participants with all types of displays; however, they show a slight dip for Weber contrast 0.25 in the Light-only (8-dot and 16-dot) displays and for Weber contrast -0.25 in the dark-only displays.

4. Efficiencies tend to be higher for the 8-dot than for the 16-dot displays.

5. Participants P3 and P6 are more highly skilled in these tasks than the other four participants. In nearly all cases, P3 and P6 achieve higher Efficiencies than do the other participants. In addition, the attention filters they achieve match the target filters more closely than do the target filters achieved by the other participants.

6. The Graded-weight attention filters achieved by the less skilled participants (P1, P2, P4 and P5) show similar asymmetries between the Dark-only vs. Light-only displays. Each of these participants achieves an attention filter in the Dark-only Graded-weight condition that is steep and concave relative to the target filter. By contrast, for each of these participants, the slope of the attention filter achieved in the Light-only Graded-weight condition undershoots the slope of the target filter significantly.

The attention filters achieved by all six participants in the conditions using the Full-set (both polarities) stimuli are shown in Fig. 1.3. For each participant the Attend-to-all, Attend-to-dark and Attend-to-light Equal-weight attention filters (the three attention filters plotted in black) are all dramatically different. The same is true of the Attend-to-all, Attend-to-dark and Attend-to-light Graded-weight attention filters (the three attention filters plotted in white). The dramatically different attention filters for the 3 types of FBA instructions confirm the previous findings of Drew et al. (2010) that participants can effectively modulate the relative salience of dots based on contrast polarity.

Different participants vary in the degree to which they are able to achieve distinct, high-Efficiency, Graded-weight vs. Equal-weight attention filters in the Attend-to-all, Attend-to-dark and Attend-to-light conditions. The most successful is P6, each of whose attention filters

approximates the target filter for the given condition very accurately and whose Efficiency is, in each condition, higher than the Efficiency achieved in that condition by any other participant. The performance of this participant proves that with proper training at least some participants can achieve clearly distinct and highly efficient attention filters in all six of the Graded-weight vs. Equal-weight, Attend-to-All, Attend-to-dark and Attend-to-light conditions.

Assessing the cost of filtering out distractors.

Let F be the attention filter achieved by a participant with the Full-set display in either the Attend-to-dark or Attend-to-light condition (either Equal-weight or Graded-weight condition), and define the function

$$F_{full-set} = \frac{F}{F(w_1) + F(w_2) + F(w_3) + F(w_4)} \quad (6)$$

where w_1, w_2, w_3, and w_4 are the target Weber contrasts in that condition. $F_{full-set}$ is the restriction of F to its target Weber contrasts, rescaled by Eq. to sum to 1. This makes it easy to compare $F_{full-set}$ to the attention filter $F_{target-only}$ achieved with an 8-dot Dark-only or Light-only display.

Fig. 1.4 show the comparisons between single-polarity and full-set attention filters. Note that if the participant were able to perfectly filter out the distractor dots in full-set stimuli, the remaining target stimuli would be identical to 8-dot single-polarity stimuli. In this case, and if the process of filtering did not consume necessary resources, the function $F_{full-set}$ would be expected to superimpose onto $F_{target-only}$.

Fig. 1.4 shows that the prediction equal filters for single-polarity and attention-filtered dual polarity stimuli holds quite well for all subjects in 3 of the 4 classes of stimuli and attention conditions: both dark and light graded attention conditions, and for attention-to-light equal attention conditions. However, this prediction fails dramatically for the equal attention to dark

15

conditions for four of the six participants P1, P2, P3, and P5 and the other two subjects show a similar but weaker tendency.

When attempting to weight all contrasts of dark dots equally and to ignore light dots in full-set stimuli, 4 of 6 participants weight the dark dots almost proportionally. This is unexpected because there is no difficulty in weighting both dark dots and light dots equally. In the mirror situation, attempting to weight light dots equally while ignoring dark dots, even the 4/6 who were unable to equally weight dark dots do a quite credible job of equally weighting light dots, although there is a slight tendency towards proportional weighting. The mean data of the 11 participants in Drew et al (2010) in the same three dual-polarity attention conditions (attend equally to all, dark, light) are strikingly similar although those authors did not comment on it. These data suggest that participants can form an attention filter to equally weight all dots independent of polarity, but when they are required to form a selective filters for dark or for light, they cannot form a good equal-weight single-polarity filter. (i.e., the equal-weight-all filter is unique and is not the sum of two single-polarity filters).

The right sides of the 6 panels in the upper half of Fig. 1.5 show the corresponding plots for the Attend-to-light, Equal-weight condition. Although none of the participants show anything like the dramatic mismatches seen in the left sides of four of these panels, for each participant, the black plot on the right side of the panel is slightly steeper than the white plot indicating that attempting to filter out dots of negative Weber contrast induces a distortion of the attention filter away from the target filter.

There is little or no evidence of any such systematic distortion of the attention filters achieved with Full-set displays in the Graded-weight conditions shown in the lower half of Fig.

16

1.5. In all cases, $F_{full\text{-}set}$ (the black curve) superimposes neatly onto $F_{full\text{-}set}$ (the white curve).
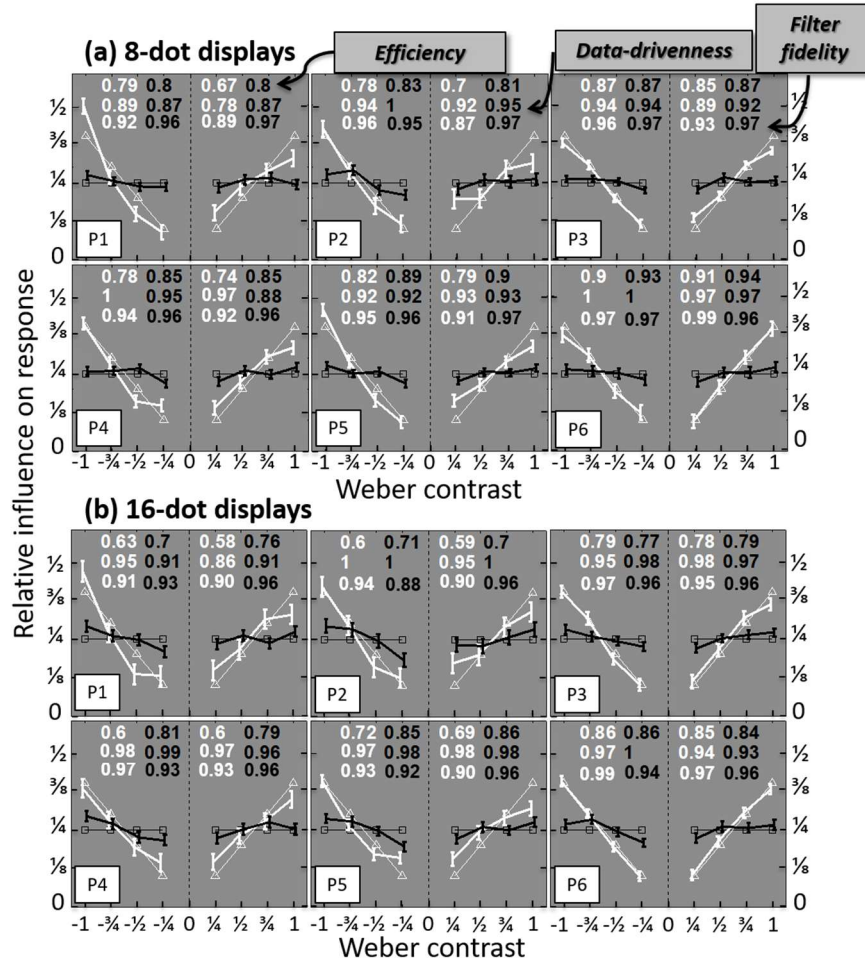
Figure 1.2. Attention filters of individual participants in single-polarity displays: Dark-only conditions (left side of each panel) and Light-only conditions (right side of each panel). Thin lines indicate target filters. Black (white) solid lines.indicate attention filters achieved in Equal-weight (Graded-weight) conditions. Error bars show 95% confidence intervals. At the top of each panel: Eff is Efficiency, DD is Data-Driveness, and FF is filter fidelity. Black numbers: Equal-weight condition; White numbers: Graded-weight condition.
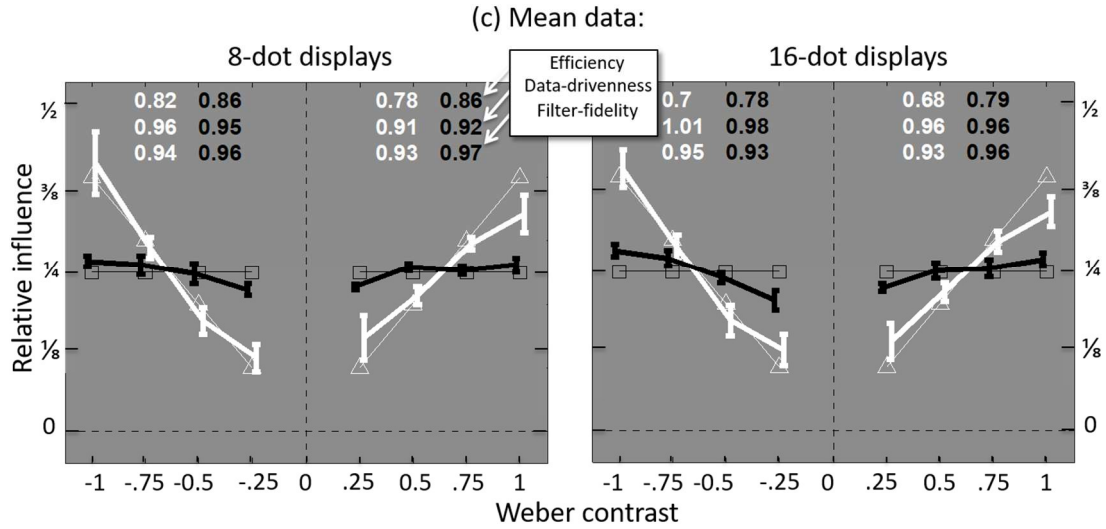
Figure 1.3. The average results from Dark-only conditions (left side of each panel) and Light-only conditions (right side of each panel) for 8-dot (left panel) and 16-dot (right panel) displays. Attention filters achieved in Equal-weight (Graded-weight) conditions are shown by black (white) solid lines. Error bars show 95% confidence intervals of the mean across the 6 participants. The target filters for Equal-weight (Graded-weight) attention conditions are indicated by the thin black lines with square markers (white lines with triangular markers). On each side of each panel the upper numbers give the Efficiency characterizing the attention filter achieved in the Equal-weight condition (in black) and the Graded-weight condition (in white), the two middle (lower) numbers give the corresponding Data-Drivenness (Filter-fidelity) values.

18
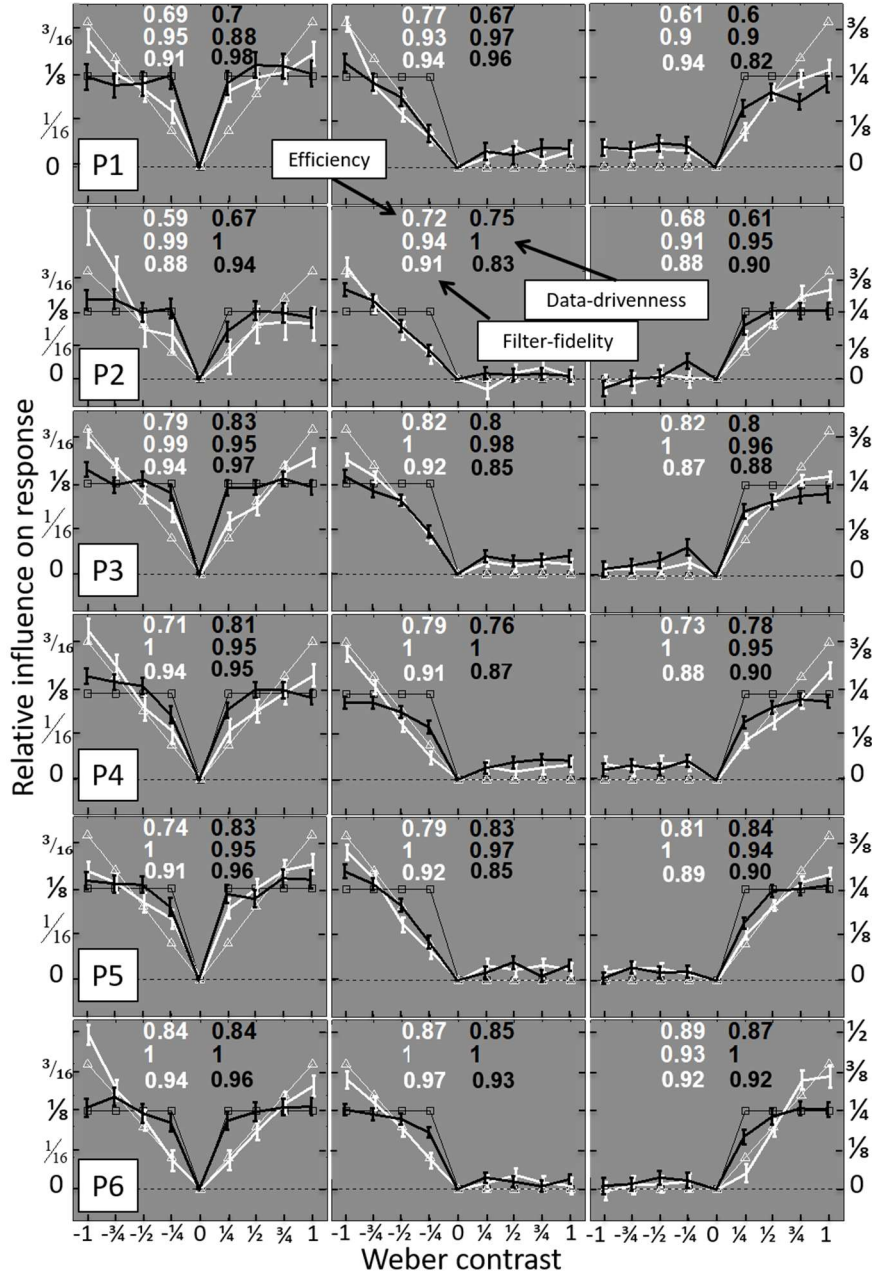
Figure 1.4. Attention filters for the conditions using the Full-set stimuli. Each row gives the results for one participant. The target filters for Equal-weight (Graded-weight) attention conditions are indicated by the black lines with open squares (white lines with open triangles). Attention filters achieved in Equal-weight (Graded-weight) conditions are shown by black (white) solid lines. Error bars show 95% confidence intervals. Left, middle and right panels show results from Attend-to-all, Attend-to-dark and Attend-to-light conditions respectively. The top (bottom) two numbers in each panel give the Efficiencies (Data-Drivenness values) achieved in the Equal-weight (black) and Graded-weight (white) attention conditions. To facilitate comparison, the vertical units in the left panel differ from those used in the center and right panels. The vertical scale on the left (right) gives units that make the target filter weights in the leftmost panel (center and rightmost panels) sum to 1.
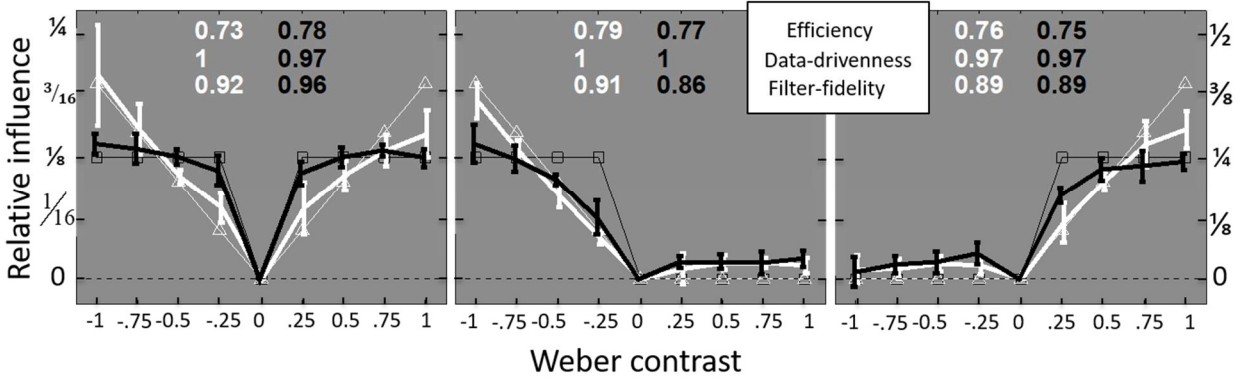
Figure 1.5. Average attention filters achieved for the conditions using the Full-set display. The target filters for Equal-weight (Graded-weight) attention conditions are indicated by the black lines with open squares (white lines with open triangles). Attention filters achieved in Equal-weight (Graded-weight) conditions are shown by black (white) solid lines. Left, middle and right panels show results from Attend-to-all, Attend-to-dark and Attend-to-light conditions respectively. The average Efficiency, Data-Drivenness and Filter-fidelity achieved in each condition are also indicated in font of the same color as the line used to plot the attention filter. To facilitate comparison, the vertical units in the left panel differ from those used in the center and right panels. The vertical scale on the left (right) gives units that make the target filter weights in the leftmost panel (center and rightmost panels) sum to 1. Error bars give 95% confidence intervals for the mean across the 6 participants.
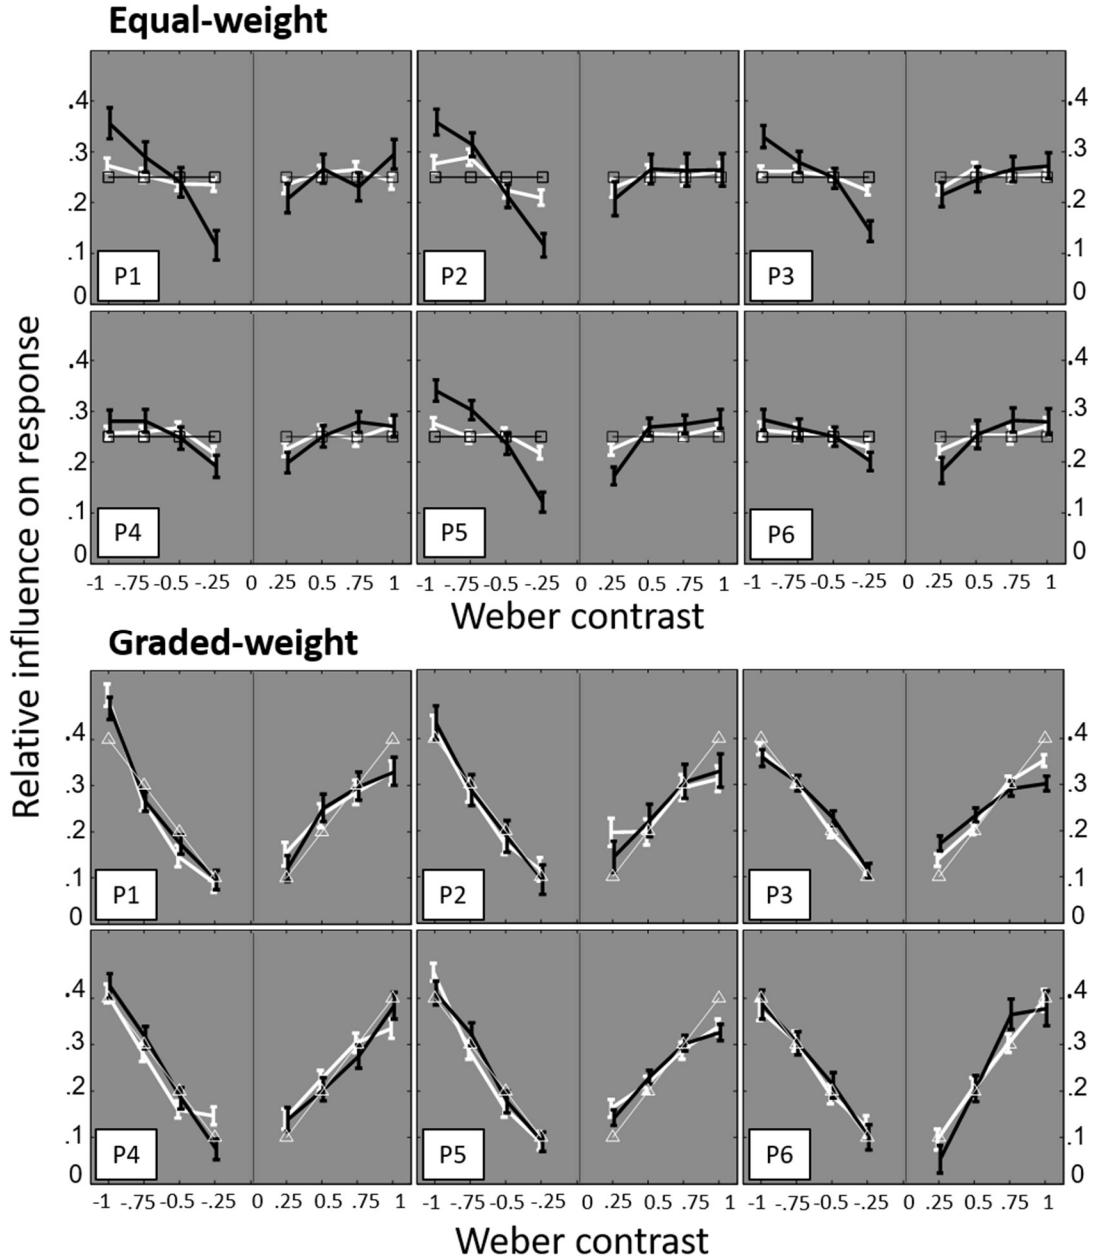
Figure 1.6. *The effect of distractors on attention filters.* Plotted in white on the left (right) side of a given panel is the Equal-weight (upper half of figure) or Graded-weight (lower half of figure) attention filter achieved by the participant with 8-dot, Dark-only (Light-only) displays. Plotted in white on the right side of each display is the corresponding attention filter achieved with Light-only displays. Plotted in black in each half-panel is the function $F_{normalized}$ given by Eq. (6). This function reflects the relative sensitivity to target Weber contrasts of the attention filter achieved by the participant with Full-set displays in the Attend-to-dark (left side of panel) or Attend-to-light (right side of panel), Equal-weight condition. If the participant were able to completely filter out the distractors from the Full-set display, then the black line should superimpose on the white line.

Discussion

The central question addressed by the current study is: what constraints are imposed by human vision on the attention filters participants can achieve for sparse displays of different Weber contrasts?

What the results with Dark-only and Light-only displays show. We begin by observing that all participants achieve clearly distinct attention filters in the Equal-weight vs. Graded-weight conditions with Dark-only and Light-only displays, in each case with high Efficiency and high Data-Driven-ness. This finding immediately implies that human vision affords easy access to at least two distinct attention filters selective for dot Weber contrast and possibly more.

It is easy to see how there might be as many as four attention filters. For each of the attention filters achieved with the Dark-only display, the sensitivity of that filter to positive Weber contrasts is unconstrained by the data from that condition (because no dots with positive Weber contrasts occurred in Dark-only displays); thus, the results are consistent with any possible extension of the attention filter to the positive Weber contrasts. Similarly, the sensitivity of each of the attention filters achieved with the Light-only display is consistent with any possible extension to the negative Weber contrasts. For example, one possibility is that each of the attention filters achieved with Dark-only and Light-only displays assigns weight 0 to dots of the untested polarity.

How could there be only two filters? The single polarity results are consistent with the idea that people have access to (A) one attention filter that gives weight to all Weber contrasts in accordance to the black plot that combines the results on the right and left side of a given panel, and (B) a second attention filter that gives weight in accordance to the corresponding white plot. Under this scenario, the participant uses filter (A) in the Equal-weight conditions with both the

Light-only and Dark-only displays and uses filter (B) in the graded-weight conditions with both types of displays.

Additional insight can be derived from comparing the single polarity data with the Full-set results. We shall argue that our results support the following conclusions:

Human vision affords easy access to

1. An attention filter that gives roughly equal weight to all Weber contrasts (except for those of very low absolute value)

2. An attention filter that gives weight 0 to negative Weber contrasts and graded weight to increasingly positive Weber contrasts;

3. An attention filter that gives weight 0 to positive Weber contrasts and graded weight to increasingly negative Weber contrasts.

but not to

1. An attention filter that gives weight 0 to negative Weber contrasts and equal weight to positive Weber contrasts, nor to

2. An attention filter that gives weight 0 to positive Weber contrasts and equal weight to negative Weber contrasts.

Individual differences are important. Some participants are more skillful than others in deploying a range of different attention filters. Participant P6 stands out in the current study; in each condition, P6 achieves an attention filter very closely matched to the target filter with high Efficiency and high Data-Drivenness. If all of our participants performed like P6, we would have to conclude that human vision affords nearly complete freedom in achieving all of the different

target filters tested in the current study. Other participants, however, show striking limitations in the attention filters they can achieve.

Evidence for a single Equal-weight attention filter activated by both polarities. Most participants seem to be unable (or partially restricted in their ability) to achieve Equal-weight attention filters selective for either positive or negative Weber contrast alone. This limitation in performance is seen clearly when we compare the attention filters achieved in the Equal-weight condition with Dark-only displays to the attention filters achieved in the Attend-to-dark-equal condition (with full-set displays). Let $F_{Dark\text{-}only\text{-}equal}$ be the sensitivity function characterizing the attention filter used by the participant to perform the centroid task with Dark-only displays in the Equal-weight attention condition. Because there are no dots of positive Weber contrast in Dark-only displays, it is impossible to tell from the results with Dark-only displays alone how the attention filter would respond to light dots. If, however, $F_{Dark\text{-}only\text{-}equal}$ assigned the value 0 to all positive Weber contrasts, then this filter would be optimal for the Attend-to-dark-equal condition (with full-set displays), in which case we would expect the attention filter achieved by the participant in the Attend-to-dark-equal task to match $F_{Dark\text{-}only\text{-}equal}$ across the set of negative Weber contrasts.

The results plotted in the left sides of the upper six panels of Fig. 1.6 emphatically reject this prediction for 4 of the 6 participants. The attention filters achieved by the other two participants in the Attend-to-dark-equal condition are also deflected slightly away from the filters that they achieved in the Equal-weight condition with the Dark-only display. We conclude that the attention filters $F_{Dark\text{-}only\text{-}equal}$ assign positive weight to dots with positive Weber contrast.

We can draw similar inferences concerning the attention filters used in the Equal-weight condition with Light-only displays. Although the mismatches are less dramatic, the results plotted

in the right sides of the upper six panels of Fig. 1.6 show that for all participants the attention filters $F_{Attend-to-light-equal}$ achieved in the Attend-to-light-equal condition (with full-set displays) are all deflected in the same way from the attention filter $F_{Light-only-equal}$ achieved in the Equal-weight condition with Light-only displays: In all cases, $F_{Attend-to-light-equal}$ gives lower relative weight to dots with Weber contrast 0.25 than does $F_{Light-only-equal}$. This implies that participants cannot directly import the attention filter that they used in the Equal-weight condition with Light-only displays for use in the Attend-to-light-equal condition. It follows that $F_{Light-only-equal}$ assigns non-zero values to negative Weber contrasts.

These observations suggest that participants are forced to use attention filters in the Attend-to-dark-equal (Attend-to-light-equal) condition that differ from those they used in the Equal-weight condition with Dark-only (Light-only) displays. The most parsimonious account proposes that participants are actually using the same attention filter for the Equal-weight conditions both with Light-only and Dark-only displays and that this attention filter is sensitive to both positive and negative Weber contrasts. If this were true, then the hypothetical attention filter would be ideally suited for the Attend-to-all-equal condition. Consistent with this idea, all 6 participants achieve attention filters that are remarkably well-matched to the target filter in the Attend-to-all-equal condition.

On the other hand, participants P6 and P4 both achieve attention filters in the Attend-to-dark-equal and Attend-to-light-equal conditions that approximate the target filters fairly well. This suggests that some participants may be able to decouple the responses to positive vs. negative Weber contrasts of the neural system that is used to achieve Equal-weight attention filters.

Black-white asymmetry: an interaction between polarity and condition-specific target-filter match accuracy.

One well-documented constraint with systematic influence on sensitivity functions is black-white asymmetry: particularly for intense Weber contrasts, vision is more sensitive to decrements from mean luminance than increments of equal magnitude. (Cf. Lu & Sperling (2012) for an overview.) Much recent work has found many physiological bases, including the faster processing speed of "black processing" off-bipolar cells compared to "white processing" on-bipolar cells maintained by thalamocortical processing (Jin et al., 2011), the predominance of off-cells in macaque V1 layers 2/3 providing output to higher cortical areas (Yeh et al., 2009), and asymmetries between on and off channel pathways in saturating non-linearities likely originating in photoreceptors (Kremkow et al., 2014). The current study found filter weights for black dots of Weber -1.0 were remarkably higher than for white dots of Weber +1.0. This constrained filter shapes such that across different stimuli and attention rules slopes for negative contrasts were steeper than for positive contrasts. Overall, graded filters were more accurate for negative contrasts than positive and equal filters were more accurate for positive Webers than negative.

The preceding discussion suggests that human vision affords easy access to a single attention filter that gives equal weight to dots of both polarities. A different picture emerges if we focus on the attention filters achieved with the Graded-weight rule. Let $F_{Dark\text{-}only\text{-}graded}$ ($F_{Light\text{-}only\text{-}graded}$) be the sensitivity function characterizing the attention filter achieved by a participant in the Graded-weight condition with Dark-only (Light-only) displays. The fact that the white and black plots superimpose in both sides of all six panels in the lower half of Fig. 1.6 suggests that each participant has been able to apply the same attention filter in the Attend-to-dark-graded (Attend-to-light-graded) condition as he/she used in the Graded-weight condition with Dark-only (Light-only) displays. This implies, however, that $F_{Dark\text{-}only\text{-}graded}$ ($F_{Light\text{-}only\text{-}graded}$) assigns weight near 0 to all positive (negative) Weber contrasts. We conclude that human vision provides ready access to

an attention filter that gives weight 0 to all negative (positive) Weber contrasts and weight that

increases in graded fashion to increasingly positive (negative) Weber contrasts.

Chapter 1

Experiment 2

Attention Filters that Give Graded Weight to a Large Set of Weber Contrasts are Easily Achieved but Improvement Saturates Rapidly

Abstract

Experiment 2 occurred before Experiment 1. Experiment 2 used massed centroid training with complex stimuli in only three display conditions that all comprised eight contrasts: (1) 8-dots, dark-only; 2) 8-dots, bright-only, and 3) 16-dots, comprising bright and dark) rather than the careful training that started from trivially easy conditions to ensure proficiency at each stage of training. However, the attention filters observed here had similar accuracy as filters from Experiment 1 had in matching targets: both were reasonably accurate but the filter to target comparison typically showed patent mismatch as well. Subsequent to Experiment 2 (here), four subjects also went on to complete Experiment 1 yet showed no improvement. The two new subjects in Experiment 1 showed attention filters with accuracy similar to the four old subjects that completed thousands of trials in Experiment 2 first and then the careful systematic procedure of Experiment 1. However, average Graded filters for single polarity displays with eight contrasts (not seen in Experiment 1) were remarkably accurate despite the nuanced gradient of different target weights. In addition to recapitulating performance seen in Experiment 1, this chapter measured attention cost by comparing filter accuracy of Uniform filters from conditions using displays matched in contrast number (8) that either comprised both polarities or only one. Filter

accuracy for the double polarity condition had slightly more accurate match to target filter and less variance, suggesting that simultaneous attention to both polarities occurred without attention cost.

Introduction

This centroid study occurred before Experiment 1 above and had the same basic aim of measuring the accuracy with which participants achieved graded attention filters for contrast while computing centroids of dot clouds comprising mixtures of light and dark contrasts. Subjects trained by estimating dot centroids following the uniform attention rule initially before undertaking the graded attention rule. Unlike Experiment 1, here we exclusively trained subjects to give graded weight using Dark-only and Bright-only displays with the reasoning it would be more manageable to train on before being tested on displays with both luminance polarities. However, every display here had eight contrasts and either eight or 16 dots. It was possible subjects had not learned how to give graded weight because displays were too complex. Hence, Experiment 1's very gradual training protocol was conceived to ensure stimuli and attention demands were always near their abilities.

Importantly, the fact Experiment 1 employed a training regimen of levels from trivial to advanced was useful for investigating whether or not improvement in graded attention filters saturates. Hence, we sought to compare the results of Experiment 2 with Experiment 1.

Methods

Eight subjects were all lab volunteers or paid UCI undergraduates. Four were male, four were female. All had normal vision or vision corrected to normal. All protocol were approved by UCI Institutional Review Board.

Except where explicitly stated, Experiment 2's methods adumbrate those found in Experiment 1.

Stimuli
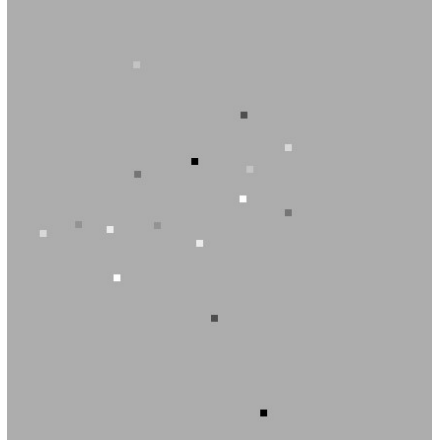
Figure 1.7 Example of a Full-set display trial with 16-dots, two each of eight Weber contrasts.
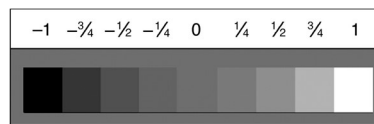
Figure 1.8. Eight Weber contrasts found in the full-axis displays. 0 was the background contrast level.

Three Weber-contrast sets were used, each comprising 8 contrasts—

The first two sets were half-axis and had one dot of each contrast including

1. the "Light-only L" (large) set comprising Weber contrasts: {0.125, 0.25, 0.375, 0.5, 0.625, 0.75, 0.875 and 1.0};

2. the "Dark-only L" (large) set comprising Weber contrasts: {-0.125, -0.25, -0.375, -0.5, -0.625, -0.75, -0.875 and -1.0};

and the set of both luminance polarities, with two dots of each Weber contrast comprising

3. the "Full-axis" set comprising every other Weber contrast in the Light-only L set and dark-only L sets: {-1, -0.75, -0.5, -0.25, 0.25, 0.5, 0.75, 1.0}.

Training and testing protocol

Pre-Training

      I.   Centroid

Initial training for task

    a. 16-extreme black dots centroid (100 trials)

    b. 16-extreme white dots centroid (100 trials)

Attention Training: Instructions always given to follow a weighting-rule

Session 1: Equal weight given to all dot contrasts

    II.

    a. 8 dots, dark (light) only

    b. 8 dots, light (dark) only

Testing

    III.   16 dots, both dark (8) and light (8)

    …continue II, III with sequence following balanced Latin square

Session 2: Dots weighted in proportion to contrast intensity

    …also follow II, III following balanced Latin square

Session 3: Equal weight

…

The chart of training and testing above shows subjects new to the study underwent brief pre-training in part I with simple homogeneous dot clouds of -1 Weber contrast for 100 trials before repeating this with dots of +1 Weber contrast. Centroid error measured in pixel distance was examined after to ensure the subject understood the task and was performing within a normal range of centroid accuracy. All subjects performed with reasonable accuracy.

After initial pre-training, stimuli in every condition consisted of 8 dot contrasts. Subjects ran permuted block sequences for training (cf. II a and II b), and testing (III) to ensure a balanced Latin square so no two subjects followed the same sequence of conditions. Two stimuli sets were used

in separate training blocks with 150 trials of 8 dots, 1 each of 8 Weber contrasts, either II a) all dark or II b) all bright.

Testing blocks had 150 trials of stimuli with 16 dots, 2 each of 8 Weber contrasts, 4 dark and 4 bright. Graded and equal attention conditions were never intermixed in a single session. All subjects went through the described sequence giving uniform attention in Session 1 before repeating the experimental sequence giving proportionate attention in Session 2. Attention conditions were thus alternated each session.

Like Experiment 1, subjects viewed their sensitivity function after each block as feedback to aid improvement.

Important differences with the incrementally more difficult training in the other experiment: The previous chapter detailed how its experimental protocol ensured attainment of increasing proficiency over the course of discrete training steps from presumably too easy to be beneficial to conditions that were highly challenging. These steps consisted of not just passing criterion tests for improvement and improvement saturation, but simpler stimuli designed to make these improvement goals at a given step easier and more achievable. For instance, in the extensive protocol training to weight attention in proportion to contrast occurred initially with dot pairs and progressed to four contrasts. Remarkably, in contrast the current study's simplest attention training displays never had less than eight contrasts and eight dots.

Number of trials run by subject in each condition:

Subject 1 ran 1,200 trials in the Full-set graded weight condition; 1,200 trials in the Full-axis uniform weight condition; 2,250 trials in the dark-only L graded weight condition; 2,150 trials in the Bright-only L set graded weight condition; 2,250 trials in the Dark-only L set uniform weight condition; and 2,250 trials in the Bright-only L set uniform weight condition.

Subject 2 ran 1,800 trials in the Full-axis graded weight condition; 1,500 trials in the Full-axis uniform weight condition; 1,900 trials in the Dark-only L set graded weight condition; 1,900 trials in the Bright-only L set graded weight condition; 1,950 trials in the Dark-only L set uniform weight condition; and 2,250 trials in the Bright-only L set uniform weight condition.

Subject 3 ran 1,200 trials in the Full-axis graded weight condition; 600 trials in the Full-axis uniform weight condition; 200 trials in the Dark-only L set graded weight condition; 200 trials in the Bright-only L set graded weight condition; 200 trials in the Dark-only L set uniform weight condition; and 200 trials in the Bright-only L set uniform weight condition.

Subject 4 ran 750 trials in the Full-axis graded weight condition; 450 trials in the Full-axis uniform weight condition; 1,050 trials in the Dark-only L set graded weight condition; 1,050 trials in the Bright-only L set graded weight condition; 1,650 trials in the Dark-only L set uniform weight condition; and 1,650 trials in the Bright-only L set uniform weight condition.

Subject 5 ran 750 trials in the Full-axis graded weight condition; 1,650 trials in the Full-axis uniform weight condition; 2,550 trials in the Dark-only L set graded weight condition; 300 trials in the Bright-only L set graded weight condition; 2,550 trials in the Dark-only L set uniform weight condition; and 300 trials in the Bright-only L set uniform weight condition.

Subject 6 ran 1,200 trials in the Full-axis graded weight condition; 450 trials in the Full-axis uniform weight condition; 1,400 trials in the Dark-only L set graded weight condition; 1,050 trials in the Bright-only L set graded weight condition; 1,400 trials in the Dark-only L set uniform weight condition; and 1,050 trials in the Bright-only L set uniform weight condition.

Subject 7 ran 1,650 trials in the Full-axis graded weight condition; 750 trials in the Full-axis uniform weight condition; 1,050 trials in the Dark-only L set graded weight condition; 1,050 trials

in the Bright-only L set graded weight condition; 700 trials in the Dark-only L set uniform weight condition; and 700 trials in the Bright-only L set uniform weight condition.

Subject 8 ran 300 trials in the Full-axis graded weight condition; 300 trials in the Full-axis uniform weight condition; 300 trials in the Dark-only L set graded weight condition; 300 trials in the Bright-only L set graded weight condition; 300 trials in the Dark-only L set uniform weight condition; and 300 trials in the Bright-only L set uniform weight condition.

Results

Both Uniform and Graded attention filter Fidelities, or filter accuracies, for Full-axis displays were overall high. Uniform filters in all subjects were better matches to targets filters than graded filters were. (Fig. 1.9) However, attention filters for Dark-only L set and Bright-only L set displays bore a different pattern: half the subjects had
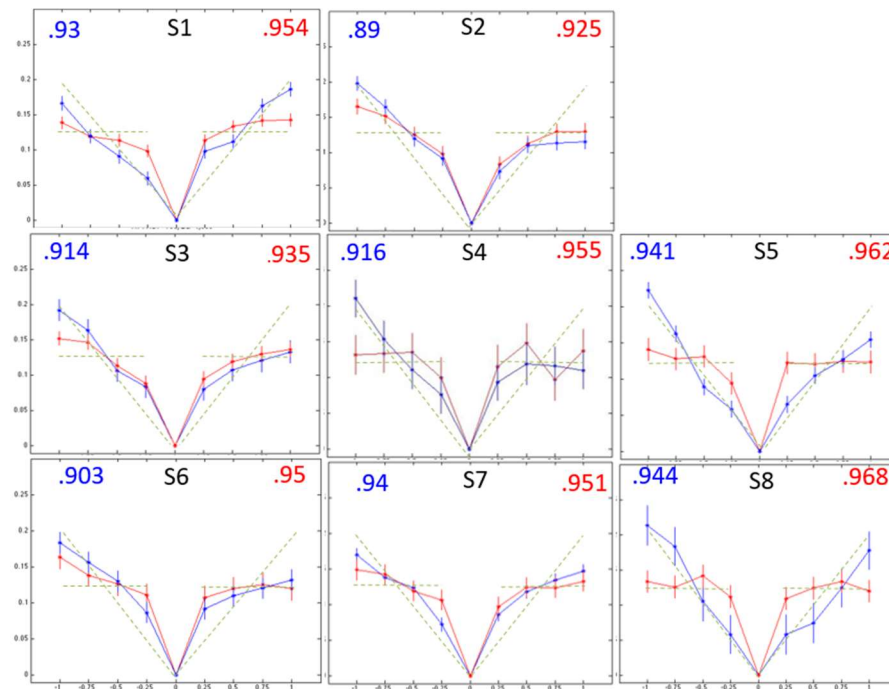
Figure 1.9. Each panel shows each of eight subjects' attention filters for Full-axis displays. Uniform attention filters (red) and graded attention filters (blue) are superimposed. Error bars indicate 95%

confidence intervals. Targets are drawn with green dashed lines. Filter fidelity, or accuracy of filter match to target, is indicated in the color of the filter.
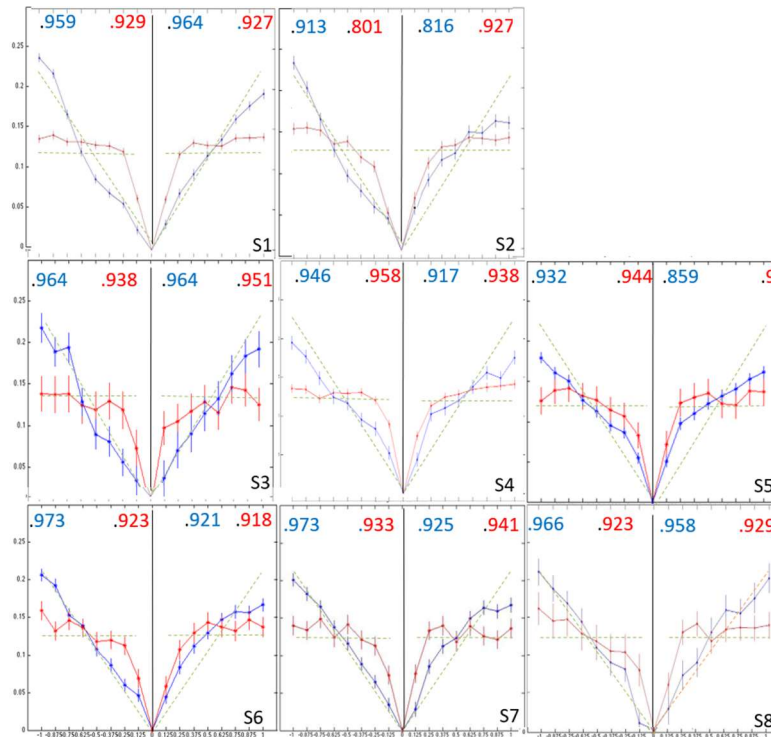
Figure. 1.10. Each panel shows each of eight subjects' attention filters for Dark-only L set and Bright-only L set displays. Uniform attention filters (red) and graded attention filters (blue) are superimposed. Error bars indicate 95% confidence intervals. Targets are drawn with green dashed lines. Filter fidelity, or accuracy of filter match to target, is indicated in the color of the filter.

higher Filter-Fidelities in the Uniform attention conditions than the Graded attention conditions;

the other half showed the opposite pattern. (Fig. 1.10) Importantly, each subject's Graded and

Uniform filters showed more significant differences in Attend-Dark and Attend-Bright conditions

than in Full-axis conditions. (Cf. figs. 1.9, 1.10) As seen with individual subjects, average

Uniform filters were more aligned with targets than were Graded filters. (Fig. 1.11)
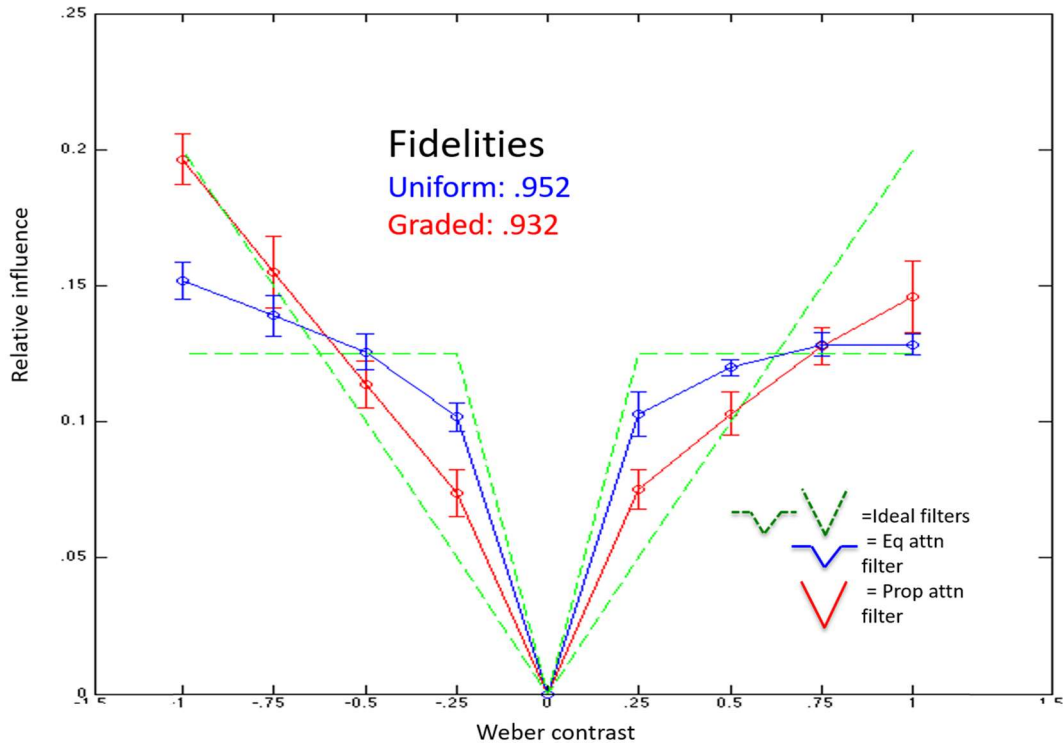
Figure 1.11. Fidelities, or filter accuracies, for average of uniform attention filters (red) and graded attention filters (blue) conditions for Full-axis display conditions of both luminance polarities. (Average of eight subjects.) Error bars represent 95% confidence intervals. Target filters are dashed green lines.
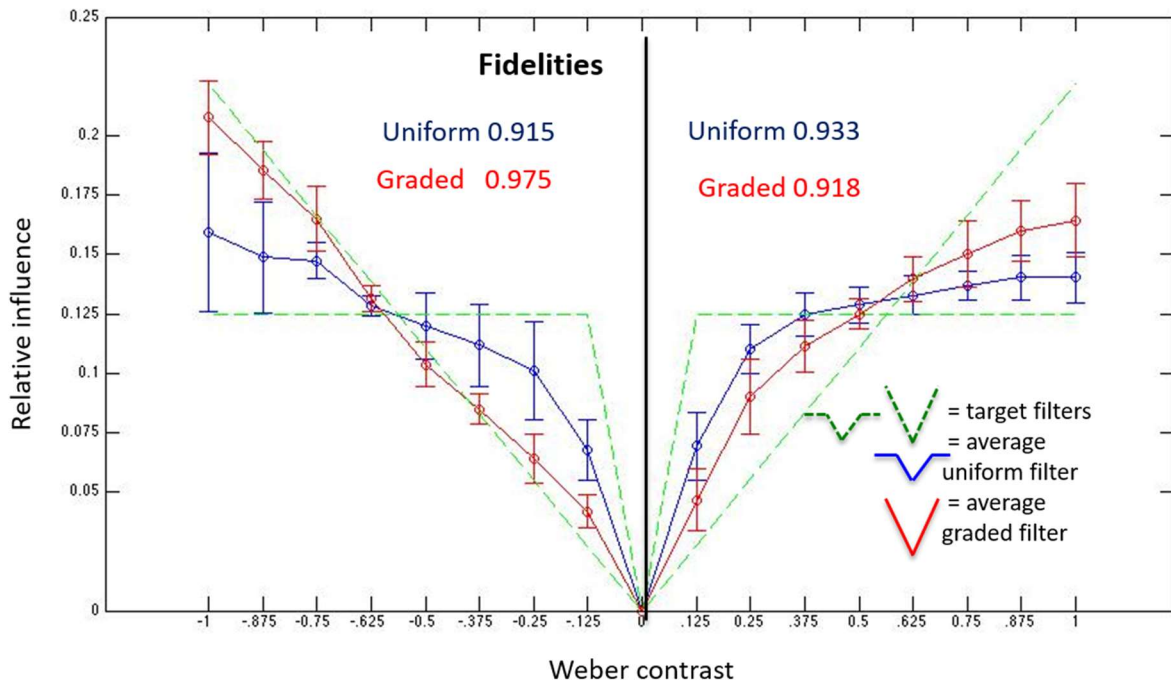
Figure 1.12. Attention filters for all single polarity displays superimposed. Uniform (blue) and graded (red) attention filters for Dark-only L set and Bright-only L set displays (8 dots, 1 dot of each contrast). Target filters for uniform and graded attention conditions are drawn in green dashed line.

Filter-Fidelities for the Dark-only L set and Bright-only L set displays are comparable to those achieved from Experiment 1's Dark-only and Bright-only displays. The filters' match with their targets is remarkably accurate, particularly the Graded attention filter for Dark-only L displays. Note: L set is an abbreviation for "Large set", which had eight Weber contrasts versus the four of Experiment 1's displays. (Fig. 1.12) It is easy to appreciate that it may well be harder to achieve Graded filters for displays that have more contrasts because that means there are more features to discriminate, attend to and accurately assign different weights in accordance with the graded attention rule. Hence, the high accuracy of the observed filters are far from entailed by Experiment 1's results.
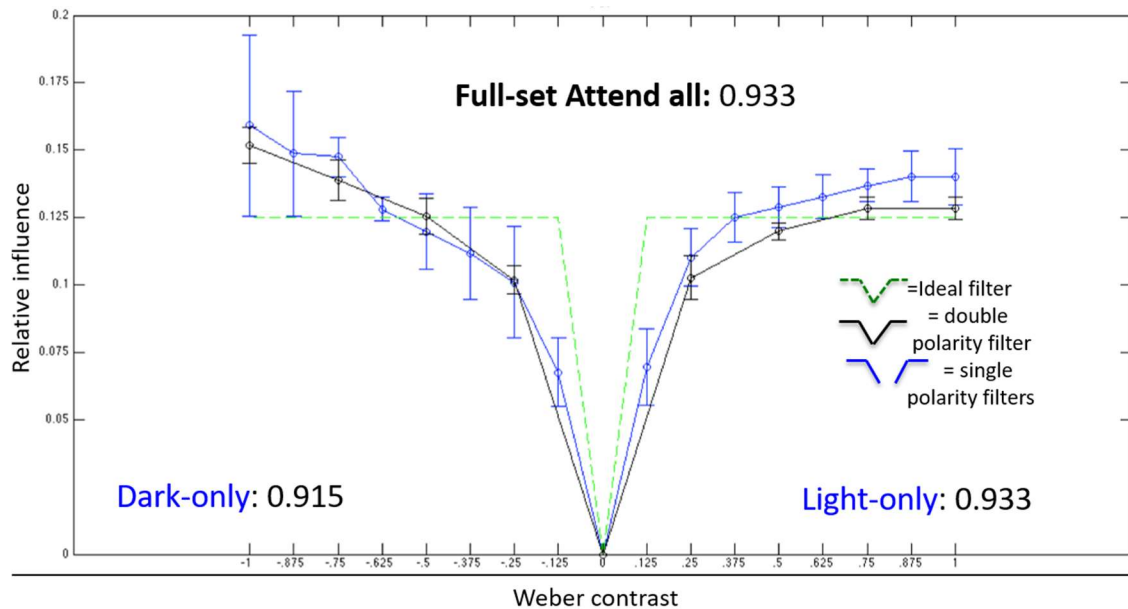
Figure 1.13. Uniform attention filters comparing Full-axis filters with combined Dark-only L set and Light-only L set filters. (Average of eight subjects.) 8-dot displays either Dark-only L set (blue) (comprising Weber set: {-1.0, -0.875, -0.75, -0.625, -0.5, -0.375, -0.25, -0.125}) and Bright-only L set (blue) (comprising Weber set: {0.125, 0.25, 0.375, .05, 0.625, 0.725, 0.875, 1.0}) were superimposed on filter for 16-dot Full-axis displays (black) comprising Weber set {-1.0, -0.75, -0.5, -0.25, 0.25, 0.5, 0.75, 1.0). Error bars show 95% confidence intervals. Target filter is green dashed line.

Fidelity, or filter accuracy, for average Uniform filters of eight contrast displays was only marginally higher for 16-dot Full-axis displays than for the combination of 8-dot displays Dark-only L set and Light-only L set. However, error bars were much smaller for Full-axis display filters showing less individual differences. Hence, Uniform weighting of two polarities compared to one occurred without a cost. (Fig. 1.13)

An inspection of whether (the later) Experiment 1 with its extensive training protocol led to higher filter accuracy than that found in Experiment 2 here.

As mentioned, we sought to know if performance in experiment 1 was higher than performance in the current experiment, experiment 2. We compared Fidelities, or filter accuracies, for Experiment
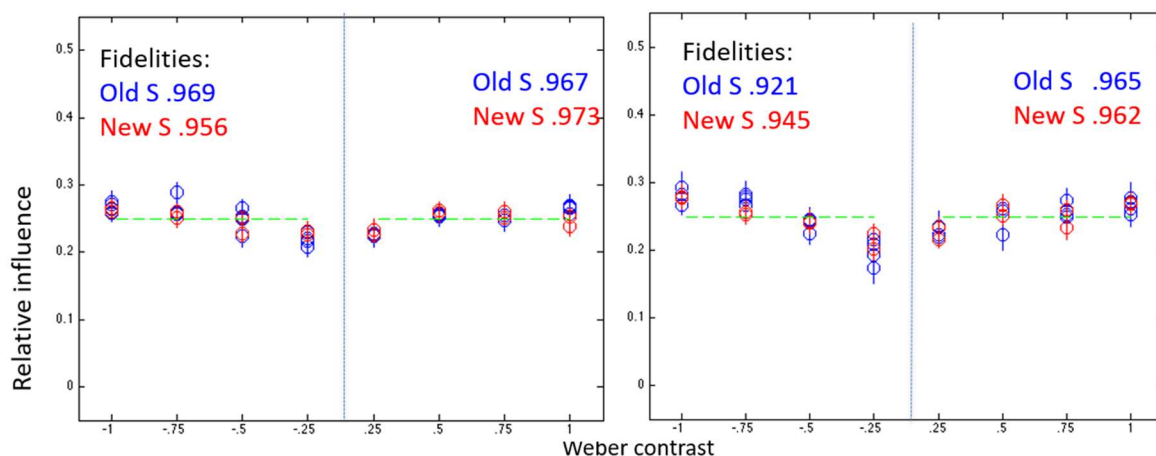
1 in the four practiced subjects that completed both Experiment 1 and 2, with the two subjects from Experiment 1 that had never performed centroids. Fidelity differences for Dark-only L set and Bright-only L set displays never exceeded .025. Thus, the new subjects performed

Figure 1.14. Comparing uniform attention filter accuracies (from Dark-only L set and Bright-only L set display conditions) of two naïve subjects (filters in red symbols) with the four old subjects (filters in blue

38

symbols) that were practiced from experiment 2. Left panel shows 8-dot display filters; right panel shows 16-dot display filters.

highly similarly to old subjects despite having run thousands of less trials. In fact, in two out of four Uniform attention Dark-only L set and Bright-only L set conditions we see slightly lower Fidelities for the four veteran subjects than seen for the new subjects. (Fig. 1.13) Graded filter Fidelities for Dark-only L set and Light-only L set displays were very slightly higher here for old subjects than new subjects, but differences in Fidelities never exceeded 0.13 and old subjects show Fidelities higher than new subjects by only an average of 0.083 across conditions. (Fig. 1.14)
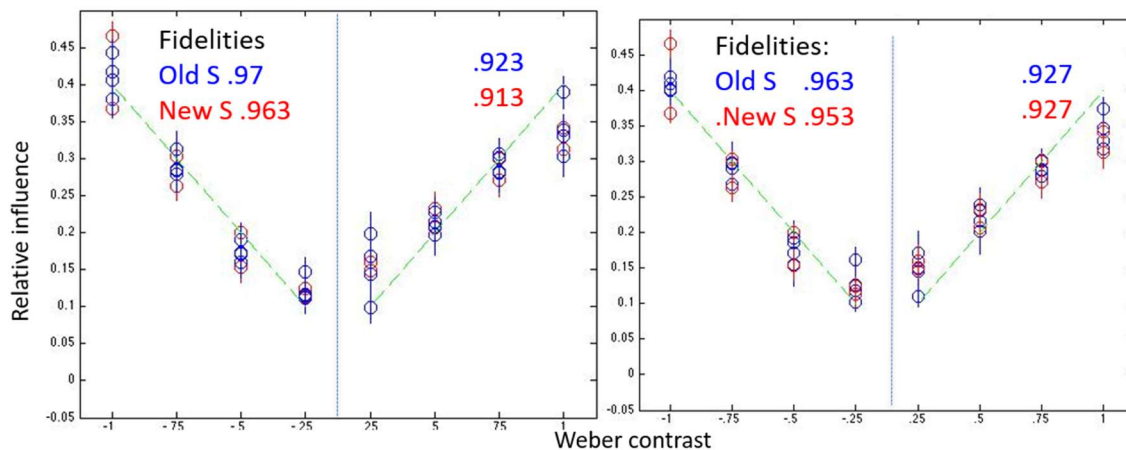
Figure 1.15. Comparing graded attention filter accuracy (from Dark-only L set and Bright-only L set display conditions) of two naïve subjects (filters in red symbols) with the four old subjects (filters in blue symbols) that were practiced from experiment 2. Left panel shows 8-dot display filters; right panel shows 16-dot display filters.

Looking at the Attend-to-Dark and Attend-to-Bright conditions, we see a meaningful filter Fidelity difference between those of new subjects and the veteran subjects. Attend-to-bright condition: Fidelity in the Uniform attention condition was 0.913 for the old subjects but only 0.782 for the new subjects. (Fig. 1.15) However, it is worth mentioning that these old subjects were just as new to the Attend-to-Dark and Attend-to-Light conditions as the new subjects were. In the Graded attention condition, the old subjects had an average Fidelity of 0.916 which was somewhat higher

than the average new subject Fidelity of 0.855. However, in the Attend-to-Dark condition, differences were negligible. (Fig. 1.16)
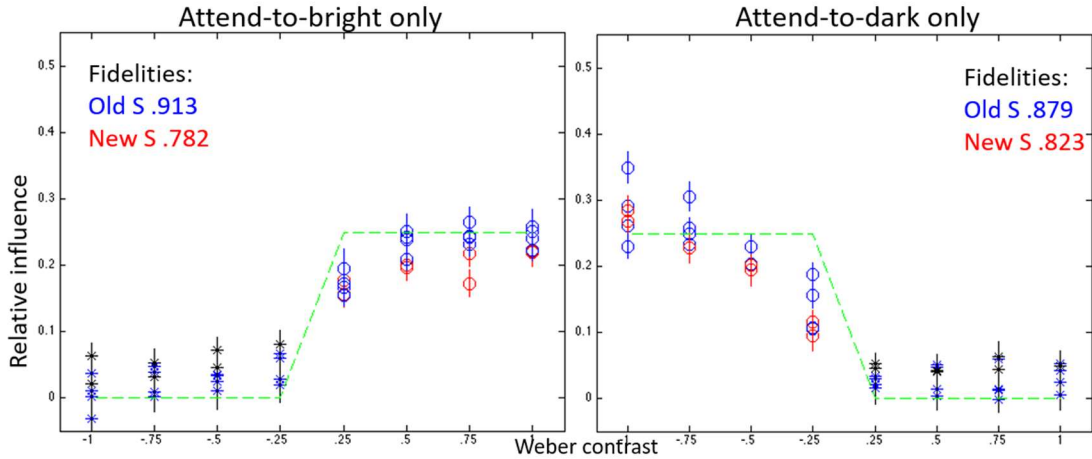
Figure 1.15. Comparing uniform attention filter accuracy of two naïve subjects (filters in red symbols) with the four old subjects (filters in blue symbols) that were practiced from experiment 2. Full-axis displays were used: Left panel shows Attend-to-bright filters; right panel shows Attend-to-Dark filters. Error bars indicate 95% confidence intervals.
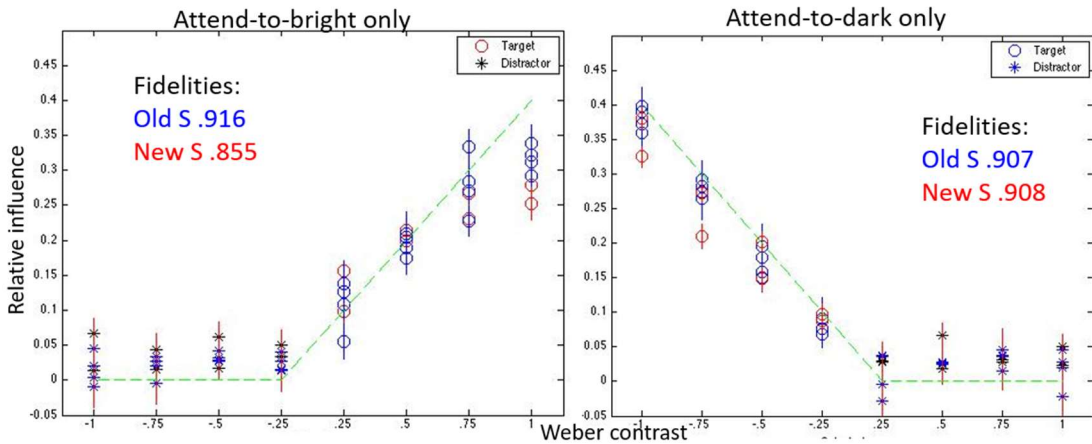
Figure 1.16. Comparing graded attention filter accuracy of two naïve subjects (filters in red symbols) with the four old subjects (filters in blue symbols) those that were practiced from experiment 2. Full-axis displays were used: Left panel shows Attend-to-bright filters; right panel shows Attend-to-Dark filters. Error bars indicate 95% confidence intervals.

The question of whether or not subjects could through sufficient training achieve target graded filters in the Full-axis display condition was the original motivation of the experiment.

Interestingly, both old practiced subjects from the initial experimental protocol of Experiment 2, as well as new subjects that only ran Experiment 1, showed lower Graded filter Filter-fidelities after Experiment 1's extensive protocol and training compared to Experiment 2's. Ironically, it was the Uniform filter which was less emphasized in these investigations wherein a slightly higher yet not significantly different Fidelity was seen in Experiment 1 with its new extensive protocol compared to that for Experiment 2: the new subjects had slightly higher Fidelity in Experiment 2, whereas the old subjects had unappreciably higher Fidelity. (Fig. 1.17)
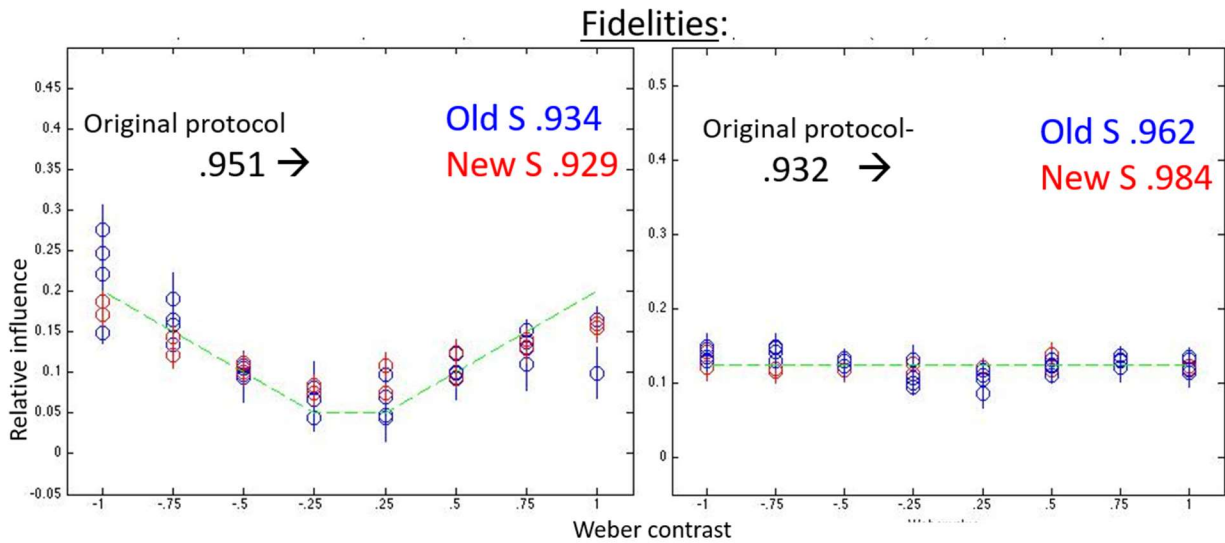
Figure 1.17. Comparing uniform attention filter accuracy of two naïve subjects (filters in red) with the four old subjects (filters in blue) that were well-practiced from experiment 2. Full-axis displays were used: Left panel shows Graded attention filters; right panel shows Uniform filters. Error bars indicate 95% confidence intervals.
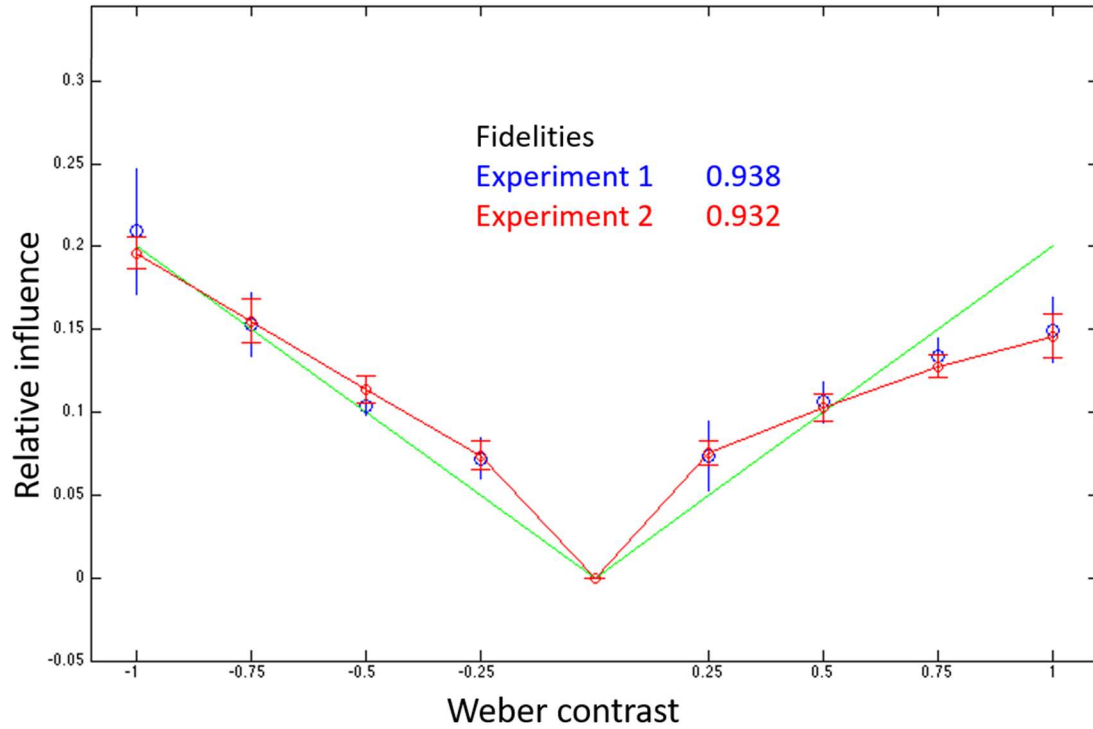
Figure 1.18. Average graded attention filters and Fidelities, or filter accuracies, (in matching colors) for Full-axis displays from Experiment 2 (red circles and capped error bars) which occurred before Experiment 1 (blue circles). Experiment 1 had an extensive training regimen and four out of six subjects had completed Experiment 2 first. (Experiment 1: n = 6; Experiment 2: n = 8) Error bars show 95% confidence intervals.

Finally, directly comparing Fidelities from the extensive training of the more reent Experiment 1 with those from Experiment 2, we see that in the graded attention condition, average Full-axis Filter-fidelities are virtually the same for Experiment 1 and 2 although error bars overall seem slightly larger for Experiment 1 indicating more variability among subject attention filters. (Fig. 1.18)
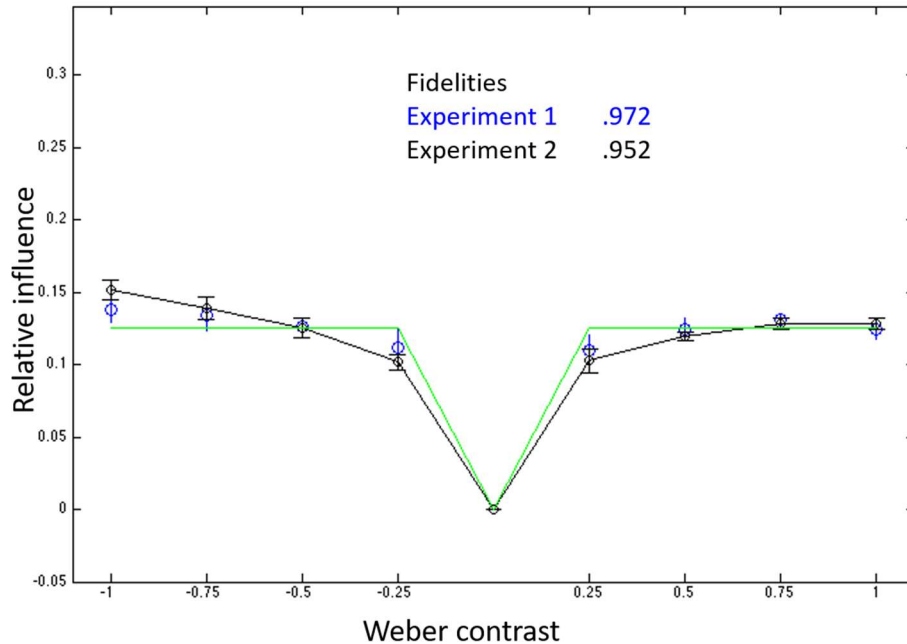
Figure 1.19. Average uniform attention filters and Fidelities, or filter accuracies, in matching colors for Full-axis displays from Experiment 2 (red circles and capped error bars) which occurred before Experiment 1 (blue circles). Experiment 1 had an extensive training regimen and four out of six subjects had completed Experiment 2 first. (Experiment 1: n = 6; Experiment 2: n = 8) Error bars show 95% confidence intervals.

Comparing average Uniform attention filter fidelities for Experiments 1 and 2, we see that training en masse with only challenging stimuli in the original protocol of Experiment 2 did not lead to performance appreciably worse that that found in Experiment 1 despite its protocol with a criterion of significant improvement and subsequent asymptote before graduating to the next level of difficulty. (Uniform Fidelity: 0.952 vs Graded Fidelity: 0.932) (Fig. 1.19.) Hence, significant difference in Filter-fidelity was not seen and this arguably shows the four old subjects in the massed training study of Experiment 2 saturated improvement in Filter-fidelity for the Graded-attention condition before beginning Experiment 1.

Discussion

Subjects in Experiment 2 achieved performance levels comparable to that from Experiment 1 despite being run on only challenging conditions. In fact, despite extensive practice the four subjects that ran both experiments showed a lower rather than higher average Filter-fidelity in the

43

second experiment despite its systematic protocol of training in the graded attention condition. On the other hand, they showed a slight improvement in the uniform attention condition that Drew et al.'s (2010) subjects were able to achieve with little training.

It was not necessary to impose Experiment 1's systematic training that upon completion of performance criteria graduated to conditions with increased Weber contrasts: all conditions in this experiment had 8 contrasts. Unlike this, the other experiment began with graded attention training on two dots of two different contrasts. Further, all displays had either 8 or 16 dots and none of the single dot or dot pair displays from the graduated protocol in the other experiment occurred here.

Although Drew's subjects (2008) did not achieve Graded filters for Full-axis displays of both polarities possibly due to running only a minimal number of blocks, subjects showed here they achieved this with non-systematic training on blocks of relatively complex stimuli always with high numbers of eight contrasts and 8 or 16 dots. Note that Subject 8 performed well yet achieved their filter within 300 trials, which was far less than the other subjects. Filters for Dark-only L set and Light-only L set displays were particularly quickly achieved. However, thousands of trials here may not have improved accuracy much. Hence, numerous easier steps of the graduated protocol were subsumed by this challenging condition that ostensibly had a steep learning curve. Experiments 1 and 2 together reveal performance levels in the Dark-only and Light-only conditions are reached quickly and saturate more rapidly than expected.

Training on displays of only one luminance polarity at a time may have helped with achieving accurate filters for both polarities simultaneously. The previous chapter found human vision affords access to filters which weight attention in proportion to contrast. It turns out these attention filters were attained with reasonable accuracy easily particularly for stimuli of only a single polarity, but performance did not improve with increased training. The foregoing clarifies

that vision can easily and automatically apply feature-based attention with a large set of different attention weights within a feature dimension like contrast. However, it saturates quickly and further training may not be beneficial.

Chapter 2

Do centroid and mean orientation judgments have access to the same attention

filters? A Comparison of Performance Across Two Summary Statistic

Representations.

Abstract

This study sought to answer whether human vision affords access to the same Weber grayscale-selective attention filters for estimating the centroid, or mean location, of a cloud of oriented bars as it does for estimating the mean orientation (MO) of the bars. In each task, three participants were tested in three different attention conditions: Attend-Light (give equal weight to all bars with positive Weber contrast and weight 0 to all bars with negative Weber contrast), Attend-Dark (give equal weight to all bars with negative Weber contrast and weight 0 to all bars with positive Weber contrast), and Attend-All (give equal weight to all bars in the display). In the centroid task, participants achieved attention filters which matched target filters fairly well in all three attention conditions. In the MO task, filter accuracy was mixed: particularly, in the Attend-Dark and Attend-Bright conditions, filters were often a slightly better match to a graded target which weights in proportion to contrast intensity, and the most intensely bright or dark target dots were often given more weight than they were in the centroid task. Importantly, for all subjects in all attention conditions, Efficiency (the minimum possible proportion of bars included in computing the response assuming an otherwise ideal observer) was more than 100% higher in the centroid task than in the MO task. The pattern of lower Efficiencies combined with giving most weight to salient contrasts in violation of target filters is suggestive of a subsampling strategy used in the MO task, but not the centroid task.

Introduction

Chapter 1 shows that human vision affords access to a variety of grayscale-selective attention filters that can be applied to the visual input prior to estimating the centroid of the filtered stimulus. We would like to conclude that these experimentally observed attention filters are available to human vision for a wide range of purposes. After all, feature-based attention is a general capability. For example, we can deploy feature-based attention to find a book whose jacket we know is a certain shade of gray, to determine whether all the black elements in a painting are contained within a trapezoidal region, or to judge whether the white items in the field of view are, on average, smaller in size or larger than the red items.

However, all the attention filters reported in Ch. 1 were derived using a centroid estimation task; thus, it is possible that the results of Ch. 1 do not generalize to other visual tasks that may involve different neural resources. The current chapter investigates this issue by comparing the grayscale filters participants can achieve in the centroid task in response to three different attention instructions calling either for attention to brighter than background bars, 2) attention to darker than background bars, or 3) attention to both at once to the to the filters they can achieve in a mean orientation (MO) estimation task.

To facilitate meaningful comparison between the two tasks across different attention conditions, the same types of displays comprising oriented bars varying in grayscale on a gray background were used in both tasks and across all attention conditions.

Methods

Observers

Three unpaid volunteers participated in this experiment. One was female, two were male. All methods were approved by the UC Irvine Institutional Review Board.

Every condition of the experiment used these stimuli in displays: 16 bars of length: 0.72° width: 0.045°, two bars of each Weber contrast: ± 0.25, ±0.5, ±0.75, ± -1.0.

Bar locations were drawn from a circular bivariate normal distribution with standard deviation 1.8° and mean location that varied randomly across trials.

Bar orientations were random, drawn from a wrapped Gaussian distribution with standard deviation of 22.5° around a mean angle that varied randomly across trials.

Bars were always located more than half a bar length from another. The position generating algorithm checked to ensure that bar positions were regenerated if positions overlapped.
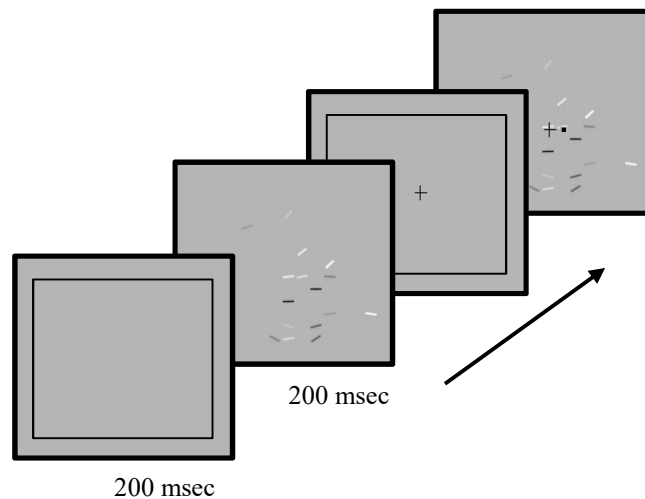
200 msec

200 msec

Figure 2.1. Example sequence in a trial of the centroid task. The trial began with a thin black square cue frame lasting 200 msec, followed by the stimulus which was displayed for 200 ms. After it disappeared, it was replaced by a cursor, which the subject moved with a mouse to click their estimated centroid location. Feedback with the subject's response, the correct response marked with a dot, and the stimuli appeared. The subject was able to study feedback as long as desired; they pressed a key to initiate the next trial.

Tasks

Centroid task

Each trial, the stimulus appeared on-screen for 200 ms, after which it disappeared and the cursor on a blank screen replaced it. The subject moved the cursor with a mouse and indicated their centroid judgment with a mouse-click. Following each trial, feedback was shown indicating the stimulus, the correct response location and the observer's response location. The feedback frame was available for the subject to review for as long as he/she wished. The following trial did not begin until after the subject initiated it with a key-press.

Mean orientation (MO) estimation task

Each trial, the stimulus appeared for 300 ms. After the stimulus disappeared, a large probe bar appeared which was white half the trials and black the other half. (Fig. 3.2) The subject rotated (counter-)clockwise with two different key presses to adjust their estimate of the mean orientation. After indicating their answer with a third key-press, feedback appeared showing the trial stimulus, the correct mean orientation, and the observer's inputted mean orientation.
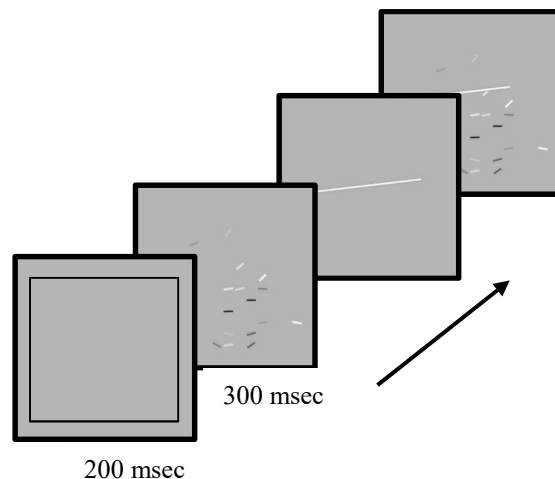
300 msec

200 msec

Figure 2.2. Example sequence in a trial of the MO task. The trial began with a square thin black cue frame lasting 200 msec, followed by the bars stimulus which was displayed for 200 ms. After it disappeared, it was replaced by a probe bar, which the subject manipulated with two different key presses to input the response. Feedback with the subject's response, the correct response, and the stimuli appeared. The subject was able to study feedback as long as desired; they pressed a key to initiate the next trial.

Attention conditions

In each of the centroid task and the MO task, each participant was tested in three attention conditions. In each of the following conditions, the participant strove to give equal weight to a particular set of *target bars*. The target bars were

1. All bars in the display in the ***Attend-all*** condition.

2. Just the bars brighter than the background in the ***Attend-light*** condition.

3. Just the bars darker than the background in the ***Attend-dark*** condition.

The same method of generating stimuli was used across both tasks and all attention conditions to ensure performance differences were due only to differences in top-down attention and the demands imposed by the centroid vs. MO tasks, not differences in stimuli.

Number of trials run—

Attend-dark condition

In the MO task, subject 1 ran 900 trials; subject 2 ran 900 trials; and subject 3 ran 1,000 trials. In the centroid task, subject 1 ran 400 trials; subject 2 ran 300 trials; and subject 3 ran 200 trials.

Attend-light condition

In the MO task, subject 1 ran 800 trials; subject 2 ran 700 trials; and subject 3 ran 700 trials. In the centroid task, subject 1 ran 400 trials; subject 2 ran 400 trials; and subject 3 ran 500 trials.

Attend-all condition

In the MO task, subject 1 ran 1,500 trials; subject 2 ran 1,700 trials; and subject 3 ran 1,900 trials.

In the centroid task, subject 1 ran 400 trials; subject 2 ran 400 trials; and subject 3 ran 500 trials.

Results

Attention filters in the centroid task were derived by fitting data to the linear regression model

given by equations 3a and 3b (Chapter 1, page 20).

Attention filters in the orientation judgment task were derived using a Metropolis-Hastings

Markov chain Monte Carlo simulation to estimate the joint posterior density according to a Von

Mises distribution. A uniform prior distribution was used which began with a random guess at

parameter vector V for which S1 = V.

Then for candidate C, $\qquad Rn = \dfrac{ln\,(\lambda\,(C)\,)}{ln(\lambda\,(S_{n-1}\,)}$ and

if Rn ≥ 1, set Sn = C.

Otherwise, set Sn = C with probability Rn; Sn-1 with probability 1 – Rn.

On the nth iteration of the Monte Carlo process, a parameter vector C is selected near Sn-

1. By selecting a sampling window of 1,000 with $S_{last\ 1000}$ as the matrix whose columns are the

1,000 most recent parameter vectors added to the list, the method can efficiently converge on the

joint posterior density. Results were stable after 20,000 iterations; we collected 100,000 iterations.

To test the null hypothesis H_0 that the same sensitivity functions were applied in both tasks for

each subject, a likelihood test

$$\text{Log-likelihood } L \ = \frac{L_0}{L} = \frac{L(\text{maximum})_{\text{Restricted model}}}{L(\text{maximum})_{\text{Full model}}}$$

was computed (Hoel, 1950) to determine the ratio of the log-likelihood of the restricted model

fitting data from the two tasks together assuming the same attention filter was used, to the log-

likelihood of the full model that fit data for each task separately under the assumption the two tasks have different filters. It is common to write this as: $D = -2 \log L$ which is distributed as a $\chi 2$ variable for the full model with $N-1$ degrees of freedom, $N = 16 =$ (Full-axis number of Weber contrasts) x (number of tasks). If the restricted model describes the true state of the world, the number of degrees of freedom equals the number of free parameters in the full model minus the number of free parameters in the restricted model, or 8. From the computed test value D we were able to arrive at a p-value that determined whether we rejected the null-hypothesis H_0 at $\alpha = 0.05$ that the centroid and MO tasks shared the same attention filter.

Computing Efficiency. On a given trial in a particular attention condition in the centroid task, our model assumes that the subject applies an attention filter f to the items in the stimulus cloud, computes the centroid of the filtered cloud, and produces a response that deviates from this centroid due to random error E that can be estimated from residual deviation of the responses predicted by the model from the actual responses of the subject. We use the following process to gauge the subject's Efficiency in the centroid task: For a given probability p, on each simulated trial of an experiment using precisely the same stimuli as were presented to the subject, (1) items are removed independently from the display with probability p, (2) the attention filter f achieved by the subject is applied to the remaining items, (3) the simulated response is taken to be the centroid of the decimated and filtered stimulus, and (4) the deviation D_p of this simulated response from the response predicted for the subject (using the undecimated stimulus cloud) is computed. We use simulations to estimate the probability $p*$ for which the mean value of D_{p*} is equal to the estimated random error E for the subject. That is, we find the decimation probability $p*$ that produces the same amount of random error as that observed for the subject. Then the subject's Efficiency is

taken to be 1- *p**. This provides a lower bound on the mean proportion of items that the subject

would need to include in his/her centroid computation in order to achieve residual error as small

as *E*. (Note that it is considered a lower bound because the model assigns all error to unregistered

dots.)

 The process used to compute Efficiency in the MO task is precisely analogous. We conduct

simulations to find the probability *p** with which items must be decimated from the display in

order for the mean orientation of the remaining items, filtered by the subject's attention filter f to

deviate from the predicted response by an amount equal to the observed error in the subject's

responses.
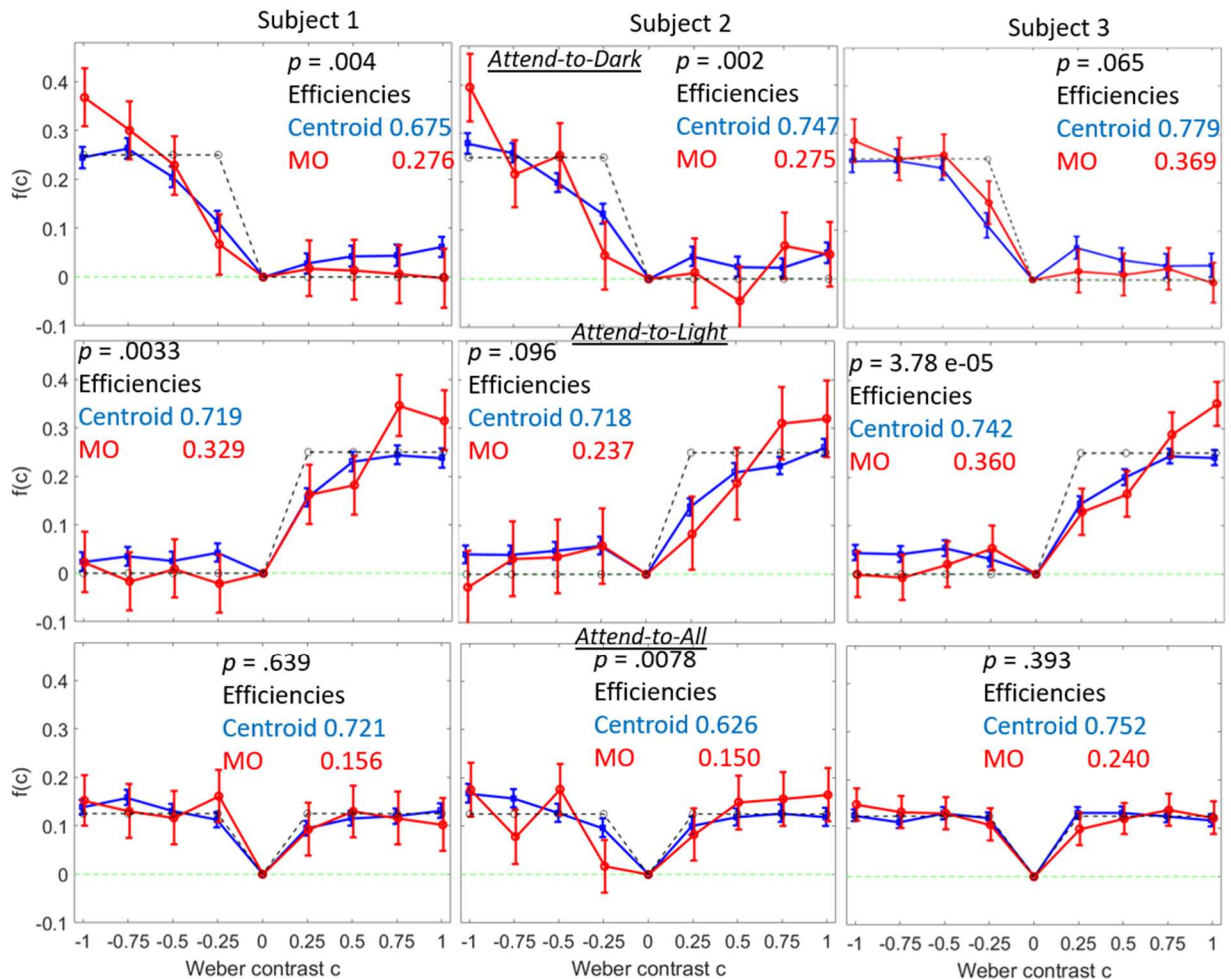
Figure 2.3. Each column contains results for a different subject. The top row of panels show results for the Attend-to-Dark condition in both the Centroid task (blue) and the mean orientation estimation task (red). The middle row of panels show results for the Attend-to-Light condition in both the Centroid task (blue) and the MO task (red). The bottom row of panels show results for the Attend-to-All condition in both the Centroid task (blue) and the MO task (red). Error bars show 95% Bayesian credible intervals. Black dotted lines show target filters.

Individual differences. The Efficiency achieved by a given subject in a given attention condition reflects both (1) random error in the subject's responses as well as (2) model failure (i.e., deviations of predicted from observed responses that are systematic but not captured by the model used to generate predictions). Subject 3 was by far the most practiced of the three subjects having run hundreds more trials in variations of the MO task. It is therefore not surprising that Subject 3's Efficiency was higher than the Efficiencies of the other two subjects in all attention conditions.

Are subjects using the same attention filters for Centroid and MO tasks?

Attend-all condition

The attention filters achieved in the Attend-All condition are similar for the Centroid and MO tasks. The likelihood-ratio test of filter equality (for Centroid and MO tasks) showed two out of three subjects had a non-significant difference between attention filters for the two tasks. The exception (Subject 2) showed MO filter shape for dark Weber contrasts that, unlike the other subjects, was jagged rather than flat. This is not a clear pattern and may not even be reliable enough to predict future performance. Error bars were consistently larger for MO than centroid.

Attend-to-Light condition

For two of three subjects, the likelihood ratio test of attention filter equality for Centroid and MO tasks yielded $p < 0.01$. Subject 3's filter showed strong significance with vanishing small $p = 0.000038$. Subject 1's filter showed non-significant difference $p = 0.09$.

Attend-to-Dark condition

Like Attend-to-Light, the likelihood ratio test revealed two of three subjects had significant difference between filters with $p < 0.01$. The exception was Subject 3 whose filter difference was not quite at significance with $p = 0.065$.

Difference in distractor weights across tasks

In the Centroid task, distractors always received non-zero weight. In the MO task, (with the marginal exception of distractor of Weber contrast -0.25 in the Attend-to-Light condition for Subject 3) distractor weights did not differ significantly from 0. (Subject 3 in Attend-to-Light condition showed lower bound of error bar for f(-0.25) = .00755 ≈ 0

Dramatic Efficiency differences

For all subjects in all attention conditions, performance in the Centroid task was more than twice as efficient as performance in the MO task. Specifically, the average ratios of Centroid-task-Efficiency to MO-task-Efficiency were 2.44 in the Attend-to-Light condition and 2.43 in the Attend-Dark condition; strikingly, in the Attend-to-All condition, the average ratio of Centroid-task-Efficiency to MO-task-Efficiency jumped to 3.97.

This last observation reflects an interesting difference in the pattern of results for the Centroid task vs. the MO Task: In the Centroid task, Efficiencies were similar across all attention conditions; by contrast, in the MO task, Efficiencies decreased sharply for all subjects in the Attend-to-All condition in comparison to the Attend-to-Dark and Attend-to-Light conditions. The implications of this difference will be explored in the discussion section.

A pattern of different attention filters for MO vs. Centroid. Centroid filters were generally accurate and often approximate fits to Uniform attention target-filters in all attention conditions. Conversely, overall MO filters in the Attend-to-Dark and Attend-to-Light conditions were better fits to *graded* target-filters with weight in proportionate to contrast intensity. (Table 2.1 & Fig.

2.3) In the Attend-to-All conditions, MO filters like Centroid filters showed Uniform attention weighting for all subjects for Subject 2.

However, the most striking and important filter difference between the MO task and Centroid task however is that the high contrast Webers c: 0.75 and 1.0 were overweighted in the MO task compared to the Centroid task: in Attend-to-Dark, Weber c: -1 were significantly elevated in MO task compared to Centroid with the exception of Subject 3. In Attend-to-Light, Webers c: 0.75 and 1.0 were elevated in MO task compared to Centroid. In two out of three subjects, significant difference was seen at one of the two contrasts. The exception Subject 2 had MO and Centroid filter weights at Weber c: 0.75 that showed only slight overlap in error bars.

Filter-fidelities, or filter accuracies, show MO filters were overall more graded to contrast intensity than Centroid filters in Attend-to-Dark and Attend-to-Light conditions.

In the Atttend-to-All condition, MO filters matched target filters with similar Filter-fidelities as Centroid filters. The exception Subject 2 showed an unclear filter pattern. MO filter error bars were larger than those from Centroid filters.

Overall, MO filters for Attend-to-Dark and Attend-to-Light conditions, those requiring attention selective to polarity, according to Filter-Fidelities, more closely matched Graded target filters than Uniform target filters. This no doubt is partly due to the overweighted extreme contrasts c: +0.75 and ±1.0. Mean filter weights in the *MO* task in five of six attend-one-polarity conditions (Three subjects x two selective attention conditions = six conditions) were at least slightly more closely matched with target filters whose weights are proportionate to contrast intensity instead of matching Uniform weight target filters as instructed.

Mean MO fidelities in the Attend-to-Dark and Attend-to-Light conditions. (Table 2.1):

a) *Attend-to-Light—* Fidelity_assuming_*uniform* filter$_{MO}$ = 0.838 <

Fidelity_assuming_*graded* filter$_{MO}$ = 0.923.

b) *Attend-to-Dark—* Fidelity_assuming_uniform filter$_{MO,}$ = 0.829 <

Fidelity_assuming_graded filter$_{MO}$ = 0.919. The lone exception, Subject's filter in

Attend-to-Dark had a Fidelity_assuming_*uniform* filter$_{MO,}$ = 0.908 >

Fidelity_assuming_*graded* filter$_{MO}$ =0.866. (As detailed later, Subject 3 may have

performed better due to being the most experienced dedicated subject that ran many

MO trials outside of this dissertation's documentation.)

Conversely, filter weights for the Centroid task in five of six attend-one-polarity conditions (three

subjects x two selective attention conditions = six) at least slightly more closely matched target

filters applying uniform weight to all contrasts rather than graded weight that increased in

proportion to contrast intensity.

Mean centroid Fidelities in the Attend-to-Dark and Attend-to-Light conditions.

a) *Attend-to-light—* Fidelity_assuming_*uniform* filter$_{centroid}$ = 0.877 >

Fidelity_assuming_*graded* filter$_{centroid}$ = 0.832.

b) *Attend-to-dark—* Fidelity_assuming_*uniform* filter$_{centroid}$ = 0.858 >

Fidelity_assuming_*graded* filter$_{centroid}$ = 0.848. The sole exception Subject 2's Attend-to-Dark

filter had a Fidelity_assuming_*uniform* filter$_{centroid,}$ = 0.864 ≈ Fidelity_assuming_*graded*

filter$_{centroid}$ = 0.865.

	Mean orientation task					Centroid task				
	Attend-to-Dark condition									
	Subject 1	Subject 2	Subject 3	Average		Subject 1	Subject 2	Subject 3	Average	Target Filter:
Uniform	0.799	0.759	*0.908*	0.829		*0.848*	*0.853*	0.864	*0.858*	**Uniform**

Graded	*0.947*	*0.868*	0.866	*0.919*		0.838	0.834	*0.865*	0.848	Graded
	Attend-to-Light condition									
Uniform	0.853	0.806	0.828	0.838		*0.896*	*0.87*	*0.862*	*0.877*	**Uniform**
Graded	*0.892*	*0.9*	*0.923*	*0.923*		0.83	0.828	0.834	0.832	Graded
	Attend-to-All condition									
Uniform	*0.934*	0.833	*0.953*	*0.907*		*0.947*	*0.929*	*0.98*	*0.963*	**Uniform**
Graded	0.804	*0.846*	0.853	0.835		0.851	0.863	0.807	0.844	Graded

Table 2.1. Fidelities, or filter accuracies, are given indicating match to a target *uniform* filter (first row of each attention condition) or a target *graded* filter (second row for each attention condition) for the mean orientation estimation and centroid tasks in each of the attention conditions *Attend-to-Dark*, *Attend-to-Bright*, and *Attend-to-All*. Bolded target filter and fidelity indicate the target filter with which the observed filter for the task was more closely aligned.

Discussion

The central finding of the current study is that Centroid task performance is much more efficient than MO task performance. One might wonder whether the MO task Efficiencies reported here are consistent with previous results. The answer is yes; other recent studies have also found low subject Efficiencies for mean estimates of visual orientations.

For example, low Efficiencies were found during a centroid task that required selected either vertical or horizontal bars that were randomly scattered (Inverso, et al., 2014). Studies finding that orientation can be effectively integrated frequently used bars whose locations and orientations remained unchanged throughout the study (Chubb & Talevich, 2003) or use a search task with high orientation contrast between targets and distractors to increase the ease with which target oriented line segments can easily induce pop-out (Wolfe and Cave, 1990).

Solomon et al. (2016) used displays comprising eight target bars situated at equal angles around a circle centered at the subject's fixation point, and the task was to judge whether the mean orientation of the bars was rotated clockwise or anti-clockwise relative to a probe bar. They compared Efficiency in this task across drastically different display times. For 100 ms displays, they observed a mean Efficiency of 0.25. Strikingly for very long (1.7 sec.) displays, mean Efficiency went up only slightly to 0.35. Hence, our MO Efficiencies for set size of eight items are actually not particularly low. Our Efficiencies are also similar to the $\sqrt{8}$ estimated in Dakin's (2001) studies which popularized the use of effective set size to quantify performance. Like Solomon, Dakin used displays of oriented Gabor textures whose mean the subject indicated as either clockwise or anti-clockwise of vertical. The displayed number of Gabor elements was either 4, 16, 256, 512 or 1024 and lasted for 100 ms. In three experiments, the stimuli across trials fixed only one of these three components and varied the other two: density of texture patch, numerosity of elements within the texture, and radius of the patch. Dakin concluded that effective set size was dependent on numerosity and not spatial arrangement.
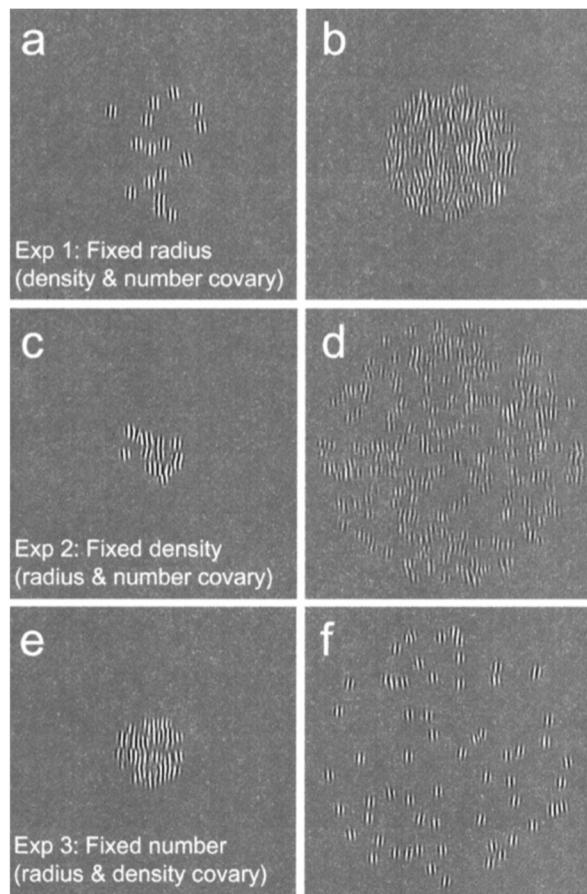
Figure 2.4. Examples of stimuli used by Dakin (2001) in three different experiments that each fixed one of these three stimulus parameters, respectively: (Experiment 1) texture patch radius (panel a); (Experiment 2) density (panel c); or (Experiment 3) numerosity (panel e) and covaried the other two- (Experiment 1, panel b); (Experiment 2, panel d); (Experiment 3, panel f).

In what ways does Centroid estimation differ from MO estimation?

The results for the Centroid task vs. MO task differ in the following intriguing way: In the

Centroid task, Efficiencies were similar across all attention conditions; however, in the MO task,

Efficiencies decreased sharply for all subjects in the Attend-All condition vs. the Attend-Dark

and Attend-Light conditions. This pattern of results is consistent with the idea that the subject

uses a "guided subsampling" strategy to perform the MO task vs. a "filtering" strategy to

perform the Centroid task. We conjecture that indeed there is a fundamental difference in the strategies used in the Centroid vs. the MO tasks.

Under the filtering strategy, the subject applies an attention filter to all the items in the display and produces her response by computing the centroid of the distribution of filter output across the entire display. Efficiency is lowered by noise that enters this process due to various factors, potentially including (i) imperfect registration of bar gray-scales and locations, (ii) centroid computation error, and (iii) motor error in producing the mouse-click response. Importantly, however, under the filter strategy, all bars in the display are included in the Centroid computation; the relative influence of bars of different gray-scales is controlled entirely by the attention filter deployed by the subject. Under this strategy, Efficiency should be equal across all attention conditions.

Under the guided subsampling strategy, the subject selects a "decision sample" comprising a fixed number of bars from the stimulus display and uses only the bars in this sample to estimate the response. Crucially, we hypothesize that the subject can select the bars in the decision sample based on their gray-scales. Thus, in the Attend-Dark and Attend-Light conditions, the subject can improve his performance by including only target bars in the decision sample. By contrast, in the Attend-All condition, guided subsampling offers no improvement in accuracy over sampling randomly from the stimulus. For this reason, if the subject is using a guided subsampling strategy, we expect Efficiency to decrease by 50% in the Attend-All condition vs. the Attend-Light and Attend-Dark conditions. The current data show this pattern for the MO task but not for the Centroid task. Suppose for illustration that the subject can only select four bars out of the eight target bars in the Attend-Dark condition in which case their Efficiency

61

might be around 0.25. In the Attend-All condition, there are twice as many targets, or 16 target bars. Thus, if the subject is using a guided subsampling strategy, we expect Efficiency to decrease by 50% in the Attend-All condition vs. the Attend-Light and Attend-Dark conditions. The current data show this pattern for the MO task but not for the Centroid task.

The idea that the subject uses a guided subsampling strategy to perform the MO task vs. a filtering strategy to perform the Centroid task is also supported by another feature of the data: the finding that distractor bars exert zero weight on responses in the Attend-Light and Attend-Dark conditions in the MO task but nonzero weight in the Centroid task. If subjects are using the guided subsampling strategy to perform the MO task, then by hypothesis, the subset of items used by the subject to compute his response in either the Attend-Dark or Attend-Light condition contains only target bars, implying that distractor bars exert no influence on responses. By contrast, under the filtering strategy, it is natural to expect distractors to exert non-zero influence on responses for the following reason: the filters achievable by the subject will be limited by the preattentive mechanisms in human vision and the ways in which these mechanisms can be combined. Given these hardware limitations, it is *a priori* unlikely that the optimal filter achievable by the subject will assign zero weight to distractor bars.

How should lower distractor weights seen for MOE compared to centroid be interpreted?

Filter weights by convention were constrained to sum to 1. (This incidentally proffers the intuition that they can be thought of as percentages of total influence weight.) Some readers may assume the lower distractor weights for MO compared to Centroid merely follow as a consequence of our convention that filters sum to 1 and the fact filter weights for extreme contrasts were often higher for MO compared to centroid. This would clearly be correct if target Weber contrasts $c < |\pm 1|$ never had lower weights. However, in five out of six cases (with exception of Subject 3, Attend-

to-Light) subjects in the polarity-selective conditions showed subsequent reduced weight at one or more of the non-extreme Weber target contrasts $C = \pm\{0.25, 0.5, 0.75\}$, but distractor filter weights (not including the error bars) were *still* lower for MO than centroid at virtually every distractor contrast as detailed. Hence, this counter argument is derailed which strengthens the argument that MO distractor weights are arguably *not* differentiable from zero.

If indeed subjects are using a guided subsampling strategy to perform the MO task, then as the current results suggest their decision samples contain very few bars. If we imagine that subjects extract the mean orientation of the decision sample errorlessly, then (given that they never include distractor bars in their decision samples), their decision samples include around $0.31 \times 8 = 2.48$ bars in the Attend-Light and Attend-Dark conditions and $0.18 \times 16 = 2.91$ bars in the Attend-All condition. Of course, we expect that the estimates of mean orientation that subjects derive from their decision samples to include some error. If so, then the actual number of bars in the decision samples will be slightly higher than the estimates given above.

In summary, the current results suggest that subjects are using very different strategies to produce their responses in the Centroid vs. MO tasks. In Centroid task, it seems likely that subjects are using a filtering strategy: i.e., they are applying condition-specific attention filters to the stimulus cloud and computing the centroid of the distribution of filter output across space. By contrast, in the MO task, it seems likely that subjects are using a guided subsampling strategy: i.e., they are basing their responses on a relatively small decision sample of items selected exclusively from the target bars in the display.

The MO attention filter is consistent with attentional capture by salient items. This may explain the overweighting of extreme contrasts in the MO task relative to the centroid task. It is not hard to imagine why in MO estimations the Weber contrasts given the most weight might be

black and white bars if indeed a subsampling strategy is being employed. If only a small subset of bar orientations are incorporated in MO estimations, the most easily detected and natural to incorporate could be salient black or dark bar orientations. McGowan and Chubb (1996) found that when subjects were only told to look toward dispersed dots as a whole, they looked toward the center-of-gravity with very small error. When half the dots were a tenth as intense as the other half, the more intense dots were weighted 30-80% more.

Further evidence that it is easiest to select the most intense contrasts: Sun et al. (2016) found in a centroid task with stimuli comprising the eight Webers of MO displays that attention filters selective for either only extreme black or extreme white contrasts were best-aligned with targets calling for all weight to be given to the single contrast and weight zero to the other seven dot contrasts. However, in the six other attention conditions in which the single target contrast were non-extreme contrasts, weight for the single target contrast was lower and much more weight was given to distractor contrasts. (Fig. 3.6) Sun et al's study is evidence that it is easiest to select extreme black and extreme white contrasts from displays of the Full-set of contrasts than it is to select any of the other contrasts. Hence, if feature-based attention is unable to integrate more than a few orientations of the target luminance polarity, it stands to reason that orientations most consistently selected may be those of the contrast most easily selected. (For an extended discussion with converging evidence from other chapters, cf. General discussion.)

One natural prediction of the Guided Sub-sampling hypothesis is that when the number of targets is reduced to no more than the effective set size we found for the given observer, the observer should more accurately achieve the Uniform target filter. The reasoning for this claim is that if subjects are able to sample all the items required by the target filter, they will not need

to resort to a more manageable strategy of primarily weighting easy to select items that deviate from the target filter.

High Efficiencies and performance in the centroid task may be due to its "special" status in attracting attention rapidly and even without inhibition of return (Zhou et al., 1996). It is not a foregone conclusion that similar Efficiencies were found for bar centroids as dot centroids from Chapter 1 because bar centroids involve the additional stage of finding the mean position of each bar whereas dots are nearly points in space. Therefore, we might reasonably expect that the ability to accurately compute centroids arguably has ecological significance in allowing an animal to maximize an object's representation within its fovea, which may explain why McGowan et al.'s observers (1998) directed saccades toward the centroid when merely told to look toward stimulus dot configurations as a whole. Movement is constrained by physics to act through the center-of-mass as is vividly illustrated by the high jumper barely clearing the bar by using the "Fosbury flop" technique to elevate and rapidly rotate their back in a snake-like action. Hence, computing centroids is arguably important for not just visual detection, but anticipating movements and hence key to survival for both predator giving chase, its prey attempting to elude, as well as competing with conspecific rivals.

Chapter 3

Polarity-Specific Complex Attention Filters for Low-Contrast Items

Abstract

In estimating the centroids of clouds of dots of different grayscales, people can (1) give equal weight to all dots (Drew, et al., 2010) or (2) weight dots in proportion to contrast intensity (Yang, et al., 2013). They can also apply these same weighting rules selectively to light dots alone (ignoring dark dots) or dark dots alone (ignoring light dots). Here we asked: Can observers achieve an attention filter for the centroid task that weights such complex dots stimuli in inverse proportion to Weber contrast? Attention filters derived from separate blocks with either 8 or 16-dot displays of a single contrast-polarity showed high accuracy matches to target filters. However, filters derived from challenging identically-generated displays with dots of both polarities from blocks calling for either selective attention to 1) bright dots-only, 2) dark dots-only, or 3) attention to all dots showed lower accuracy in matching target filters but nonetheless weighted dots in inverse proportion to contrast magnitude. Critically, in the selective attention conditions the lowest contrast items of the target polarity were consistently given the most weight and the lowest contrast items of the distractor polarity were given weight nearly 0, indicating that the mechanism subserving this attention filter serves to heighten sensivity to low contrast items of a given polarity but is not sensitive to the other.

Introduction

A large body of research supports the claim that visual attention can be captured by items which have much higher levels of a feature property (such as luminance or contrast) than other items in the scene and that such items exert more influence on performance in a range of visual tasks (e.g., Borji, Sihite & Itti, 2013; Borji & Itti, 2013; Borji & Itti, 2012; Carrasco, 2011; Pomplun, 2006). From this perspective the finding, reported in Ch. 1 of this dissertation, that participants can achieve attention filters for grayscale that weight targets in proportion to contrast intensity seems natural. Perhaps slightly more surprising is the finding (Ch. 1 of this dissertation, Drew et al, 2010) that participants can also achieve attention filters that give roughly equal weight to items varying strongly in Weber contrast.

The project reported in the current chapter addresses the following question: Can attention filters give weight in inverse proportion to contrast intensity? This question is important for the following reason. We assume that all attention filters that can be achieved by our participants make use of the rapid, spatially parallel image transformations hardwired into the architecture of human vision. Many neurons in low-level primate vision produce responses that increase monotonically with stimulus contrast. For example, on-center and off-center neurons in both the magnocellular and parvocellular layers of the lateral geniculate nucleus produce responses that increase with the contrast of the stimuli used to excite them. It is easy to imagine how retinotopically organized arrays of such neurons might be recruited by top-down attention to achieve attention filters that give higher weight to high-contrast items than to low-contrast items. It is less obvious, however, what neurons exist in human vision that might be used to achieve attention filters that reverse this pattern to give higher weight to low-contrast versus high-contrast items. If participants can achieve such attention filters, this will imply the

existence in human vision of one or more preattentive mechanisms (i.e., one or more image transformations realized by a retinotopically organized arrays of neurons) that assigns higher activation to low-contrast vs. high-contrast items.

Methods

If not explicitly stated otherwise, methods were identical to those in chapter 1's dot centroid experiment.

Observers

There were three subjects (1 female). Subjects 1 and 2 were well-practiced in centroid tasks from having participated in the experiments reported in chapters 1-2. Subject 3 was new. Thus, unlike subjects 1 and 2, subject 3 had not previously received training in the graded attention condition which is more challenging than the uniform attention condition making it a useful intermediate stepping stone en route to this experiment. Despite being naïve to centroids, subject 3 performed well compared to the other two subjects. All subjects had normal vision or vision corrected to normal. All methods used were approved by the UC Irvine Institutional Review Board.

Materials & display device

The same Viewsonic VGA CRT monitor, computer mouse and chin rest set-up from the last two experiments were used to perform the experiment.

Stimuli & trial procedure

Method of stimuli generation was the same as in Chapter 1. Stimulus dot Weber contrasts were drawn from the set {-1.0, -0.75, -0.5, -0.25, 0.25, 0.5, 0.75, 1.0} on the mean gray background (64.5 cd/m^2). Each square 7x7 pixel dot subtended 0.21° of visual angle. The

dots' positions were drawn from a circular bivariate Gaussian density with standard deviation of 100 pixels (3.01° of visual angle). The mean of the Gaussian cloud was also subject to a random perturbation with a circular, bivariate normal distribution with mean equal to the center of the display with standard deviation of 30 pixels (.904° of visual angle). The stimulus field within which dots appeared was 512x512 pixels. This region subtended 14.94° of visual angle at the viewing distance of 61 cm, and it was circumscribed by a thin, square, black frame centered in an otherwise gray display.

Also identical were the sequence of events each trial (see Ch. 1, Fig. 1.1). Before each trial, the participant fixated the center of the stimulus field and pressed a button to initiate the trial, which was marked by the onset of the cue frame. The stimulus was then displayed for 200 ms, during which time the border disappeared. After the dots disappeared, the box reappeared along with a small, black, cross-shaped mouse-cursor in the center of the box. The participant used the mouse to move the cursor to click on the response location. Visual feedback was then given; the feedback display showed the dots from the current trial, a bull's-eye composed of three concentric black rings indicating the correct response location, and a small dot indicating the location selected by the participant. The participant then pressed space bar to proceed to the next trial; this allowed him/her to study the feedback as long as desired.

Experimental conditions

A given experimental condition is marked by the *Weber-contrast set* used to generate the dots in a given stimulus cloud, the *number* of dots used in the stimulus clouds and the *target filter* used in that condition. Three Weber-contrast sets were used: the "Light-Only" set comprising Weber contrasts 0.25, 0.5, 0.75 and 1.0; the "Dark-Only" set comprising Weber

contrasts -0.25, -0.5, -0.75 and -1.0; and the "Full" set comprising all 8 Weber contrasts in the union of the light-only and dark-only sets. In the conditions using the Dark-only and Light-only Weber-contrast sets, the target attention filter f that was used to give feedback gave weight to each dot in inverse proportion to the absolute value of the dot's Weber contrast. Specifically (using the convention that target filter values sum to 1), for conditions using the light-only Weber-contrast set, the target filter assigned $f(0.25) = 0.4$, $f(0.5) = 0.3$, $f(0.75) = 0.2$, and $f(1.0) = 0.1$; for conditions using the dark-only Weber-contrast set, the target filter assigned $f(-0.25) = 0.4$, $f(-0.5) = 0.3$, $f(-0.75) = 0.2$, and $f(-1.0) = 0.1$.

All stimuli that used the full set of Weber contrasts included 2 dots of each Weber contrast. Three different attention conditions used Full-set stimuli:

1. Attend-to-all-inverse-graded: the target function f assigned f(-0.25) = 0.4, f(-0.5) = 0.3, f(-0.75) = 0.2, and f(-1.0) = 0.1.

2. Attend-to-dark-inverse-graded: f(-0.25) = 0.4, f(-0.5) = 0.3, f(-0.75) = 0.2, and f(-1.0) = 0.1; f(0.25) = f(0.5) = f(0.75) = f(1.0) = 0.

3. Attend-to-light-inverse-graded: f(-0.25) = f(-0.5) = f(0.75) = f(-1.0) = 0; f(0.25) = 0.4, f(0.5) = 0.3, f(0.75) = 0.2, f(1.0) = 0.1

Every subject was tested in the different conditions in the same order. The purpose of this aspect of the design was to use the easier conditions as training for the more-challenging conditions. The outline below reflects the order in which conditions were performed. Also shown is the number of trials performed in each condition by each of the three subjects.

1. Dark-only Weber-contrast set

a. 8-dot displays (2 dots of each Weber contrast)

    i. HY: 400 trials

    ii. MR: 500 trials

    iii. EB: 300 trials

b. 16-dot displays (4 dots of each Weber contrast)

    i. HY: 600 trials

    ii. MR: 1,200 trials

    iii. EB: 300 trials

2. Light-only Weber-contrast set

a. 8-dot displays (2 dots of each Weber contrast)

    i. HY: 600 trials

    ii. MR: 400 trials

    iii. EB: 300 trials

b. 16-dot displays (4 dots of each Weber contrast)

    i. HY: 600 trials

    ii. MR: 2,100 trials

    iii. EB: 300 trials

3. Full Weber-contrast set (all displays contained 2 dots of each Weber contrast)

a. Attend-to-all-inverse-graded

    i. HY: 800 trials

    ii. MR: 1000 trials

    iii. EB: 800 trials

b. Attend-light-inverse-graded

i. HY: 300 trials

ii. MR: 300 trials

iii. EB: 600 trials

c. Attend-to-dark-inverse-graded

i. HY: 600 trials

ii. MR: 500 trials

iii. EB: 700 trials
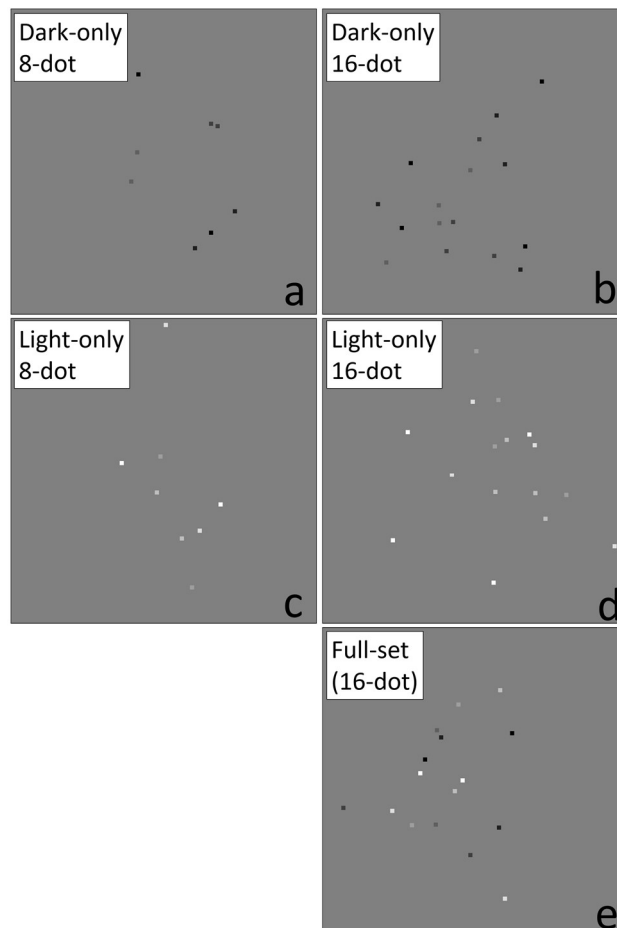

Figure 3.1. Example trial dots centroid stimuli in each of the different experimental conditions tested in separate blocks using the inverse contrast graded weighting rule. Panels on first row show stimuli from the Dark-only dot conditions, with either 8-dot displays **(a)** or 16-dot displays **(b)**. Panels on second row show stimuli from the Bright-only dot conditions, with either 8-dot

displays **(c)** or 16-dot displays **(d)**. Third row panel **(e)** shows example stimuli for conditions with 16 dots comprising Full-set Weber contrasts used in the study. 3 different attention conditions had Full-set displays: 1) Attend-All, 2) Attend-to-Light, and 3) Attend-to-Dark.

Results

The results for the Dark-only and Light-only Weber-contrast displays are shown in Fig. 3.2. Each panel contains two groups of 3 numbers. In each case, the uppermost group gives the Efficiencies, or minimum proportion of set items required by an otherwise ideal observer to achieve identical performance, listed for subjects S1, S2 and S3 in the given condition; and the lower group gives the attention filter Fidelities, which measure accuracy of observed filter-to- target filter match. (A perfect match has a Fidelity of 1.)

Overall patterns across conditions: in accordance with targets, attention filters gave high weight to low contrast dots that was in fact overweighted relative to target; they gave relatively low weight to salient extreme contrast dots of Weber $\pm$ 1.0. (Figs. 3.2- 3.4) Attention-filters achieved in the Attend-to-Light and Attend-to-Dark conditions are highly effective at excluding distractors, even those of the very lowest contrast. Note that indeed the distractor Weber –(+) .25 was not given more weight than other distractor contrasts. (Fig. 3.4) Inverse-graded filters deviated from the target filters in being concave instead of linear. (Fig. 3.2) This indicates that within a luminance polarity sharpest discrimination occurs between lowest Weber contrasts +(-) 0.25 and +(-)0.50 which then declines between each successively higher contrast level. As with graded filters reported in Chapter 1, attention filters achieved for dark contrasts are overall more aligned with target filters than attention filters achieved for bright contrasts. (Fig. 3.2)

Performance differences across conditions: attention filters achieved with single-polarity displays showed the highest performance (as reflected both by Filter-Fidelity and Efficiency) and were comparable to the Graded and Uniform filters reported in Chapter 1. (Fig. 3.2.) The lone exception

73

was the attention filter achieved by Subject 1 with the Light-only displays: attention filter shape for 8-dot Light-only displays was markedly flatter than the filters achieved by the other two subjects. However, even in this case, weights for higher contrast levels were lower in accordance with target filter. Subject 1's 16-dot Light-only filter was more accurate but still flatter than the others; importantly, it was consistent with the others in giving less weight to higher contrasts. (Importantly, Subject 1 participated in every experiment in this dissertation, but was never high performing.)

Full-set displays with the decidedly greater challenge of attending to both luminance polarities simultaneously instead of only one showed the most marked concave shapes. Although the filter shapes were more patently non-linear compared to filters from a single display, they overall showed a monotonic decrease in weights for higher contrasts. Fidelities, or filter accuracies, for Full-set displays fell below Fidelities for Uniform and Graded filters. (Chapter 1) Unlike filters for Dark-only and Bright-only conditions, filters for Full-set conditions did not overweight low Weber c : +(-) 0.25.

Comparing Full-set display filters: conditions with selective attention to polarity vs. Attend-All condition. Filters from Attend-to-Dark and Attend-to-Light conditions underweighted middle contrasts c: + (-) 0.50 and +(-) 0.75, but filters from Attend-to-All conditions did not. (Fig. 3.4) Notably, filters in the Attend-All condition showed weight above target at contrast c: +0.75 for two out of three subjects and at c: -0.75 for one out of three subjects. Conversely, attention filters for Attend-to-Dark and Attend-to-Light conditions did not show weight above target other than at extreme contrasts c: ±1.0. Overall, Attend-to-Dark and Attend-to-Light filters showed more monotonic grading inverse to contrast than did Attend-All filters.
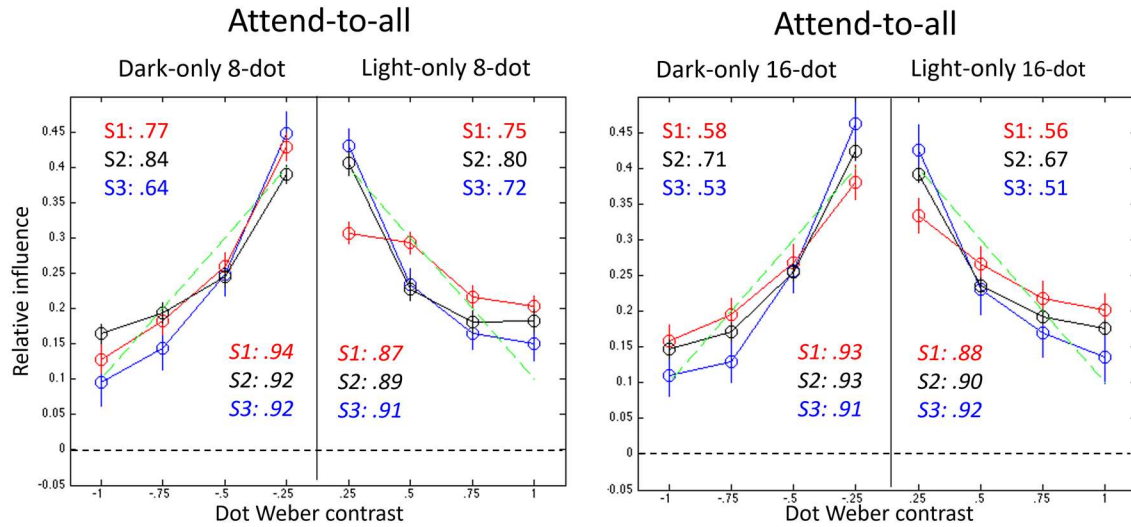
Figure 3.2 Attention filters for the four conditions with displays of one contrast polarity: (Dark / Bright x 8 / 16 dots). (Vertical black line separate data from different display conditions.) Green dashed lines indicate target filters. Circles connected by lines in red, black, and blue indicate attention filters for each of three subjects. Error bars indicate 95% confidence intervals. Numbers along top of boxes indicate Efficiencies, or estimated proportion of set items used to achieve the subject's filter of the same color. Italicized numbers along bottom of boxes represent Fidelities (filter accuracies).
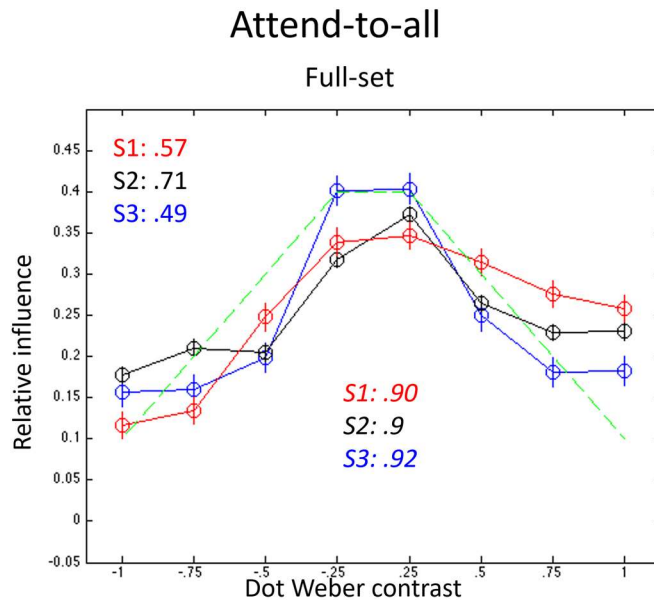
Figure 3.3. Attention filters for the Full-set condition comprising both contrast polarities, attention given to all. Green dashed lines indicate target filter. Circles connected by lines in red, black, and blue indicate attention filters for each of three subjects. Error bars indicate 95% confidence intervals. Numbers along top of boxes indicate Efficiencies, or estimated proportion of set items used to achieve the subject's filter of the same color. Italicized numbers along bottom of boxes represent Fidelities (filter accuracies).
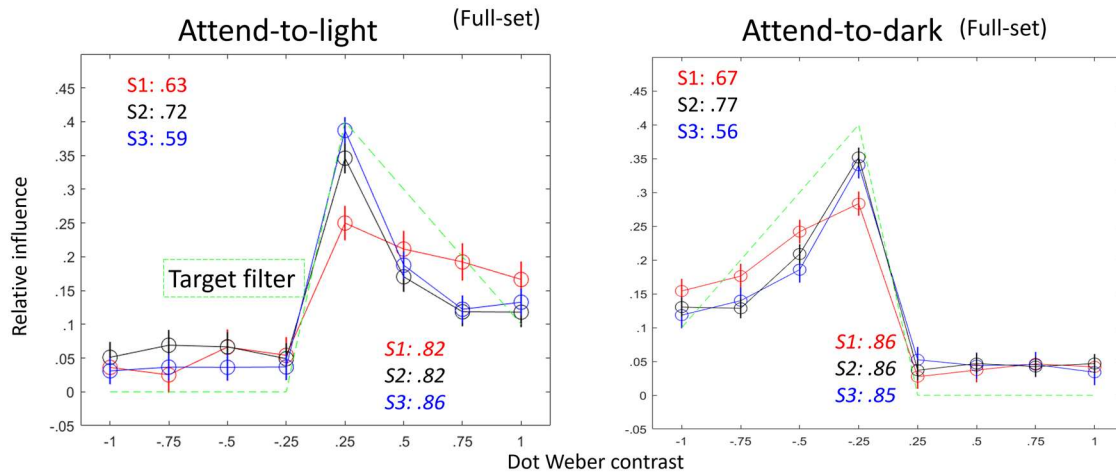
Figure 3.4. Attention filters for Full-set displays comprising both contrast polarities and selective attention: Attend-Bright only (left) and Attend-Dark (right). Green dashed lines indicate target filters. Circles in red, black and blue connected by lines indicate attention filters for each of three subjects. Error bars indicate 95% confidence intervals. Numbers along top of boxes indicate Efficiencies, or estimated proportion of set items used to achieve the subject's filter of the same color. Italicized numbers along bottom of boxes represent Fidelities (filter accuracies).
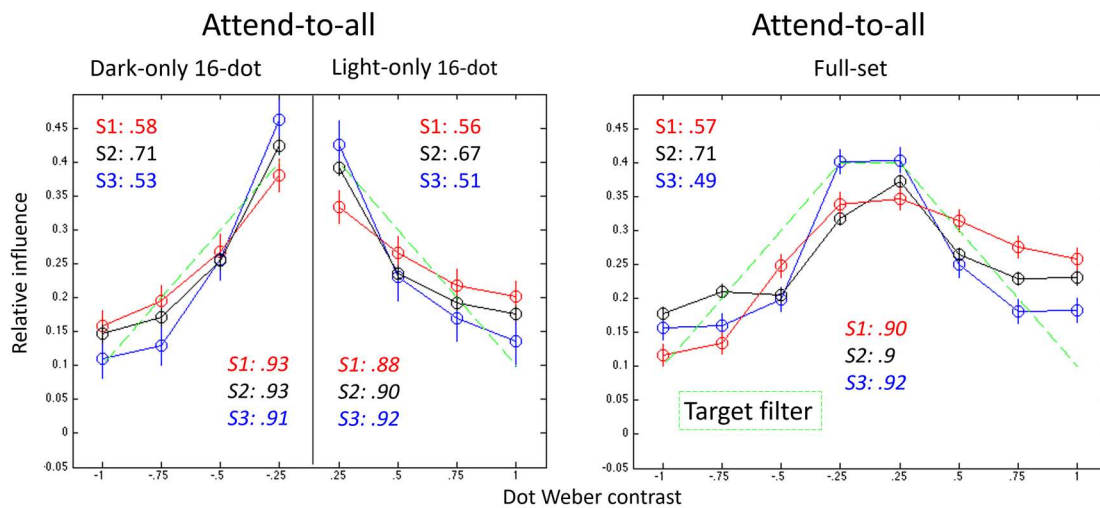
Figure 3.5. Comparing the loss of 16-dot display filter accuracy from attending to only dark dots (left panel, left side), or in another condition, only bright dots (left panel, right side), with attending simultaneously to the Full-set of dark and bright Weber contrasts (right panel). Green dashed lines indicate target filters. Circles in red, black and blue connected by lines indicate attention filters for each of three subjects. Error bars indicate 95% confidence intervals. Numbers along top of boxes indicate Efficiencies, or estimated proportion of set items used to

76

achieve the subject's filter of the same color. Italicized numbers along bottom of boxes represent Fidelities (filter accuracies).

Discussion

In all of the conditions in the current study, the participant strove to give weight to dots of different gray-scales graded in inverse proportion to the absolute value of dot Weber contrast. The main finding of this chapter is that participants achieve successful performance in most variants of this task. Attention filters were fairly good match to target filters with sharp discrimination for the lowest contrast c: ±0.25, and Efficiencies are only slightly lower than those characterizing the filters achieved in the Attend-to-all-graded conditions in Chapter 1.

Performance differences reflected by attention filters: All inverse-graded filters had concave shapes due to overweighting relative to target-weight for high contrast c: ±1.0, relatively weaker discrimination for high contrasts c: +(-) 0.75, +(-) 1.0, and at least slight underweighting relative to target-weight for middle contrasts c: ±0.50, ±0.75. (Cf. General discussion on concave filter shapes at dissertation's end.) However, each condition had characteristic attention filters indicating different patterns of performance and sensitivity compared to those from other conditions. Attention filters from Attend-Dark and Attend-Bright displays showed overweighting relative to target-weights for low contrasts c: ±0.25 indicating strongly heightened salience for the most prioritized yet low intensity contrast. Full-set Attend-one-polarity filters gave low but non-zero weight to distractors. Recall that attention filter weights are constrained by our Centroid model convention to sum to 1. Weights were lower for non-salient middle Weber contrasts c: ± 0.50 and ± 0.75 compared to Dark-only and Bright-only conditions. (Figs. 3.2, 3.4) Evidently the attention cost of filtering opposite polarity distractors was decreased weight to the middle two contrasts.

Full-set Attend-all filters showed effective discrimination at the lowest contrast c: ±0.25. Inverse-graded filters in all conditions had relatively low discrimination for the two highest contrasts c: +(-) {0.75, 1.0}, but Attend-all filter weights also had reduced discrimination noticeable in flatter slopes and smaller weight differences at contrasts c: +(-) {0.50, 0.75, 1.0}.

We can determine the attention cost of attending to two polarities simultaneously compared to attending to each polarity one at a time by comparing 16-dot display filters, either from the Full-set Attend-all condition or from the combined filters from Dark-only and Light-only conditions. (Fig. 3.5) A comparison of filters for the two credible subjects S2 and S3 suggests that attending to two polarities simultaneously reduces weight for priority low contrasts c: ±0.25 and raises weight for extreme contrasts c: ±1.0 that should receive the least weight. Interestingly, Efficiencies were not reduced.

Performance differences reflected by Efficiencies: Inverse-graded conditions were somewhat more challenging than the graded conditions tested in Chapter 1. With the 8-dot Dark-only and Light-only displays, participants achieved Efficiencies that ranged between 0.64 and 0.8 (mean = 0.75); these Efficiencies are comparable to those for the filters achieved in the Attend-to-all-equal and Attend-to-all-graded conditions in Chapter 1 in which Efficiencies ranged between 0.67 and 0.91. (mean = 0.8).

For all participants, with both the Light-only and Dark-only displays, when dot number was increased from 8 to 16 dots, Efficiency decreased. The average size of this decrease is 21%, and the six individual decreases range from 15% and 30%. By contrast, in Chapter 1, in the Attend-to-all-graded conditions with both the Light-only and Dark-only displays, when dot number was increased from 8 to 16 dots, Efficiency decreased slightly less: the average size of the decrease was 14%, and the twelve individual decreases ranged from 4% to 23%.

Importantly, each subject had 16-dot Efficiencies that were similar across-conditions, including the challenging Attend-Dark, Attend-Light, and Attend-All conditions. Although these differences were small, they were highly consistent: this pattern was also seen in the graded attention conditions in 5 of 6 subjects, but not Uniform filters (Chapter 1). Efficiency reflects centroid accuracy but not necessarily filter accuracy, suggesting Efficiency differences were more a function of item set numerosity than attention condition. This suggests inverse graded attention filters but not uniform filters may begin to decline in effectiveness beginning with 16-dot and larger set sizes. Importantly, because each subject had similar Efficiencies for all 16-dot conditions regardless of difficulty, it can be argued that performance in each of the attention conditions using inverse grading are similarly credible and differ only in filter accuracy.

The accuracy of Full-Set Attend-All filters is lower in the inverse-graded attention condition than the graded attention condition. Although Filter-fidelities, or filter accuracies, are around 0.91 for both attention conditions, the match to target shape is patently worse for inverse-graded filters, which are strikingly concave rather than linear in shape. In contrast, Graded filters are linear and a clear match-to-target on the negative (dark) axis. Efficiency in the Graded attention condition averaged 0.72 for 6 subjects (Chapter 1) but in the inverse-graded condition averaged only 0.57 for three subjects.

Thus, the current results suggest that the inverse-graded conditions tested in this chapter are more difficult than the proportionally-graded conditions tested in Chapter 1; however, this difference is relatively slight. For example, compare the current Efficiencies to those achieved in the mean orientation task (Chapter 2). The mean Efficiency (across three participants) for the three variants of the centroid task tested with these displays was 0.72; by contrast, the mean Efficiency for the mean orientation task in the same attention conditions was 0.27 (63% lower).

Participants are using a filtering strategy (not a guided subsampling strategy) in the current tasks. We argued in Chapter 2 that participants used a different (guided subsampling) strategy to perform the mean orientation task that was fundamentally different from the (filtering) strategy they used to perform the centroid task. The primary basis for this conclusion was the precipitous drop in Efficiency achieved in the Attend-to-all condition in the mean orientation task from the Attend-to-light and Attend-to-dark conditions. It should be noted that no such dramatic drop in Efficiency is present in the results for the inverse-graded attention conditions tested in the current chapter. The mean Efficiency for the Attend-to-light and Attend-to-dark conditions with the Full-set displays is 0.66; in comparison, the mean Efficiency for the Attend-to-all condition is 0.59 (only 11% lower). The current results thus suggest that participants are indeed using a filtering strategy to perform the inverse-graded conditions tested in the current Chapter.

The reported performance despite brief displays of 8 and 16-dot sets with too much visual information to remember requires the existence of a preattentive mechanism whose peak tuning matches the peak of the attention filter. The current results thus imply that human vision possesses a mechanism selective for low-contrast items in the visual input.

In particular, they show the existence of separate, preattentive mechanisms selective for low-contrast light and low-contrast dark scene elements. Any filtering strategy of the sort achieved by our participants in the current tasks requires the existence of one or more preattentive mechanisms that can be recruited to construct the filter achieved in the task. The attention filters achieved by subjects S2 and S3 (and to a lesser extent by participant S1) in all conditions show a marked concavity: instead of giving weight inversely proportional to dot Weber contrast, these filters tend to give sharply increasing weight to the dots with Weber contrasts ±0.25. This suggests

that the primary mechanisms recruited for the tasks in this chapter are sharply tuned for low contrast scene elements.

That there exist separate "low-contrast-dark" and "low-contrast-light" selective mechanisms is suggested by the attention filters achieved in the Attend-to-light and Attend-to-dark conditions with Full-set stimuli. In the Attend-to-light condition, for example, participants achieve attention filters that are from 5 to 12 times more sensitive to dots of Weber contrast 0.25 than to distractor dots of Weber contrast -0.25. The results are analogous for the Attend-to-dark condition. The effectiveness with which participants are able to ignore low-contrast distractor dots in these two conditions while simultaneously achieving high, sharply tuned sensitivity to low-contrast target dots implicates separate mechanisms tuned to low-contrast elements of opposite polarity. An argument can be made that the cleanest estimates currently available of the sensitivity functions characterizing the low-contrast-light and low-contrast-dark mechanisms come from the results of Sun et al., 2016. This study tested each of five participants in 32 different task conditions. Eight of these conditions used stimuli that consisted of clouds comprising 24 dots, three each of the same eight Weber contrasts ±0.25, ±0.5, ±0.75, and ±1.0 displayed in the Full-Set condition here. In different attention conditions, the task was to mouse-click the centroid of the three dots of a single target grayscale, ignoring dots of all the other grayscales. The average attention filters achieved in each of these conditions are plotted in Fig. 3.6. Note in particular the attention filters achieved for target Weber contrasts 0.25 and -0.25. These two attention filters closely resemble the attention filters achieved in the Attend-to-light and Attend-to-dark conditions of the current study with the Full-set displays.
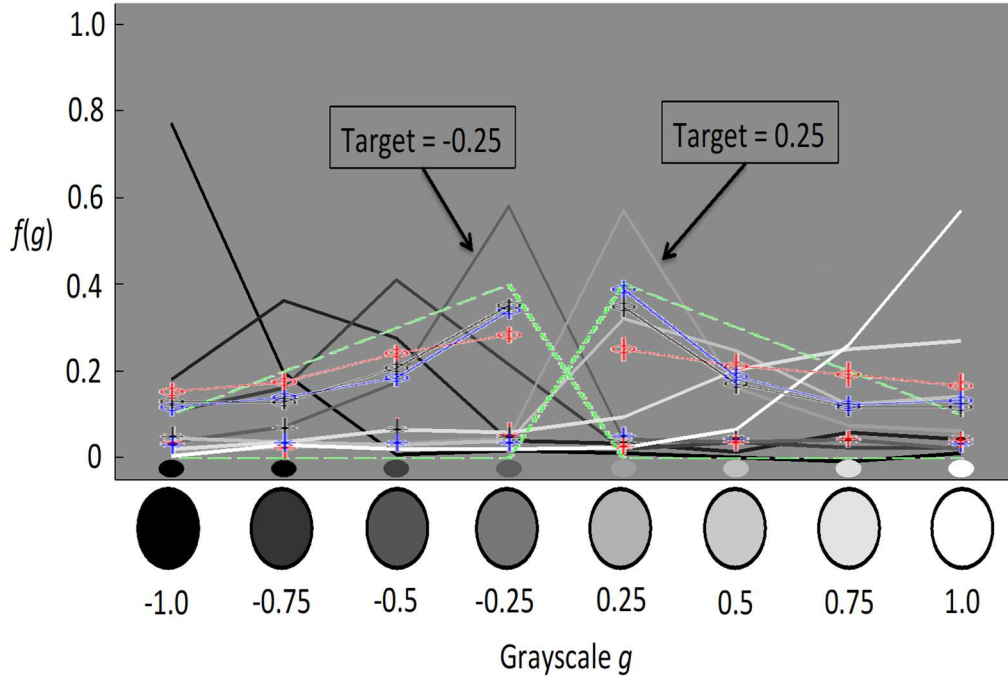
Figure 3.6. Superimposed attention filters for single contrasts by Sun, et al. (2016) drawn in color of target dot contrast (for each of eight attention conditions using the same stimuli) with inverse-graded attention filters for Full-set display, Attend-Dark and Attend-Bright conditions. Circles (target polarity dots) and asterisks (distractor polarity dots) in red (S1), black (S2) and blue (S3) connected by lines indicate attention filters for each of three subjects. (Abscissa: Weber contrasts present in all displays: {-1, -.75, -.5, -.25, .25, .50, .75, 1} Error bars indicate 95% confidence intervals.   Green dashed lines indicate target filters for Attend-Dark and Attend-Bright which cross near bottom-center of figure.

It should also be noted that in accordance with target filters, the attention filters achieved in all conditions succeeded in giving non-zero weight to all non-distractor Weber contrasts. In fact, in most cases non-distractor dots of Weber contrasts ±1.0 received weight greater than the weight specified by the target filter.   This property of the attention filters achieved in the current chapter is unsurprising. As shown in Chapter 1, participants easily achieve attention filters for grayscale that give weight approximately equal to all target Weber contrasts. This suggests that participants possess a "grayscale-invariant" mechanism that responds equally to dots of all grayscale contrast.

Approximating an inverse graded filter selective to polarity by combining other filters.

Proof of principle for a linear combinations theory of contrast: an important hypothesis concerning feature-based attention proposes that the only attention filters human observers can achieve are linear combinations of the sensitivity functions that characterize the preattentive mechanisms in human vision. In other words, all sensitivity functions are either elemental functions or composite functions of these elemental functions. A corollary of this hypothesis is that complex filters like the inverse-graded filter which give a range of weights for a variety of feature-levels might comprise separate independent filters. For instance, if we appropriately combine the output of the grayscale-invariant mechanism shown in chapter 1 with the output of the single grayscale selective mechanism discovered by Sun et al., 2016, we may be able to achieve the inverse weighted filters in this chapter. By taking a weighted sum of a Uniform filter and the attention filter for c: - 0.25 from Sun et al., we arrive upon an approximation to the average inverse-graded attention filter in the Attend-to-Dark condition for the three subjects. (Fig. 3.7) This is proof of principle for the linear combinations theory of FBA.

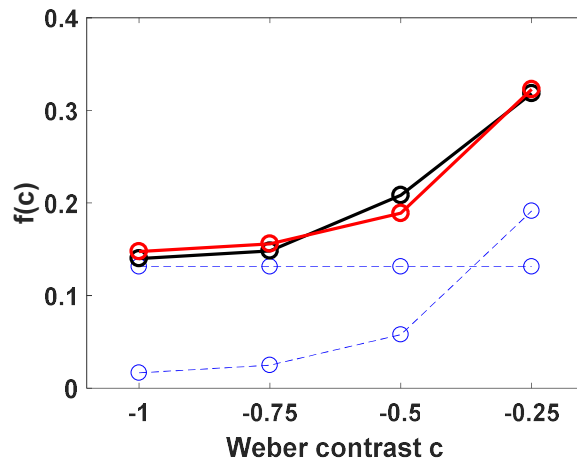
Figure 3.7. Computing the weighted sum of 1) the single contrast filter for c: -0.25 from Sun et al., 2016 (lower blue curve) and 2) the Uniform target-filter for all contrasts (horizontal upper blue line) yields a fair approximation (red) of the mean filter of the three subjects of the Attend-to-dark inverse-graded condition (black). (Bright distractor polarity weights omitted for clarity.)

Sun, et al.'s (2016) filters may also suggest that inverse graded filters will not match the linear shape of targets. Listed in order from most aligned with the target weight to least were: f(c) <≅ 0.8 for the extreme dark contrast c : -1; f(c) ≅ .6 for extreme bright contrasts c: 1, 0.25, and -0.25. The contrasts c: ±0.50, ±0.75 were even less aligned with targets.

When target contrast was middle contrast W =0.75, W = 1 actually had higher weight; when target contrast was the other middle contrast W = 0.5, W = 0.25 had higher weight. These bright filters for W= 0.75-only and W = 0.25-only did not just overweight distractors and underweight targets, they showed high and low contrast distractors closest in contrast to target contrasts were treated more like targets than targets were! This suggests that filters which require giving less weight to the extreme contrast than the middle contrasts may not be highly accurate.
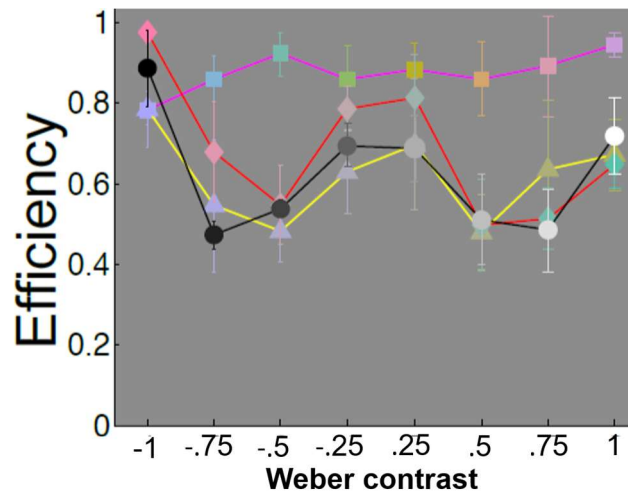
Figure 3.8. Efficiency, or the proportion of total dots an otherwise ideal detector would need to yield observed performance (on ordinate) in a 24-dot centroid task with selective attention for a single Weber contrast (abscissa) amidst distractor dots of the other seven Weber contrasts (3 dots of all contrasts). Attention filters for extreme black dots had close to perfect Efficiencies. Filters for extreme white dots had similar Efficiencies as filters for low contrast dots. Lowest Efficiencies were seen by filters for close-to-mean contrasts ±0.5, ±0.75.

Efficiencies for the 3 subjects ranged from 0.59 to 0.72 for Attend-to-Light and 0.56 to 0.72 for Attend-to-Dark. These Efficiencies are similar or slightly lower than the range of Efficiencies with mean = 0.7 seen by Sun et al. in their two conditions, Attend-to-Weber c: 0.25 and Attend-to-Weber c: -0.25. (Fig. 3.7.)

The relative Efficiencies also recapitulate the relative accuracies of these filters in matching target filters: the highest Efficiencies were shown for extreme target contrasts Webers = $\pm$ 1.0; Efficiencies for low target contrasts Webers = $\pm$ - 0.25 were lower; and Efficiencies for target contrasts Webers = $\pm$- 0.5, $\pm$ 0.75 (near set mean Weber c: 0.625) had the lowest Efficiencies. Both Sun's filters and the inverse graded filters require ignoring one polarity and giving maximal weight to a low contrast Weber and as described may have a basis in common attention filters, which would entail the similar Efficiencies found.

With proper training, how much might inverse-graded filters have improved? Remarkably, one of the subjects was naïve to the centroid task. As seen in Chapter 1, inverse-graded attention filters are achieved rapidly but may saturate rapidly. It behooves us to consider the space of top-down filters that can be imposed on bottom-up attention to contrast. (Cf. General discussion at dissertation's end for causes of concave filter shapes.)

The neurophysiological substrate of the low-contrast-light and low-contrast-dark mechanisms. The evidence for attention filters that assign greater weight to items in a scene with lower contrast intensity and hence natural salience call to question the mechanism responsible. Schiller and Lee (1991) ablated V4 in monkeys and found a deficit in the speed and accuracy with which they made saccades to an oddball low contrast circle that appeared in an array of 8 iso-eccentric positions with uniformly more intense circles. However, the V4 lesioned monkeys did not show a deficit in making saccades to an oddball circle when instead it had greater contrast than the distractors or

even when a low contrast circle appeared alone without distractors. Hence, V4 is a candidate region for mechanisms of feature-based attention for targets low in contrast or features like motion speed, etc. Clarification of substrate is a worthwhile newer direction for future research.

General Discussion

Conjectured Mechanisms Subserving Attention Filters Throughout Chapters.

Disclamer: The following is the culmination of taking intuition and evidence down their course. It is equally plausible that unrelated mechanisms instead are responsible.

An account of the concave shape of graded filters tying-in evidence from other chapters.

The concave filter shape reflected the fact inverse-graded filters gave sharply discriminating weights between the lowest contrast and the next highest, but the sensitivity seemed to dull with each increase in contrast. In particular, filters gave at least target weight to the lowest contrast, underweighted at least one of the middle two contrasts, but characteristically overweighed the highest contrast. The most consistent of these features was that extreme contrasts were overweighted. The reasons are unclear, but it is a worthy question to ask. To this end, consider the following: we demonstrated earlier that the average inverse-graded attention filter in the Attend-to-Dark condition could be well-approximated by averaging the target Uniform filter from Chapter 1 with Sun et al.'s attention filter for Weber c: -0.25-only (Fig. 3.6). However, this does not explain why graded and inverse-graded filters are not symmetric in shape. We will attempt to provide a plausible account here with converging evidence from other chapters which overall support the claim that Graded filters assign weight based on salience of contrast intensity and hence weight increases roughly linearly; inverse-graded filters heighten salience for lower contrasts that are relatively less salient and hence they suppress naturally salient high contrasts

leading to poor discrimination of greater contrast levels. A vivid analogy is that you run faster down a mountain than up it, but of course this is not because you have two different "modes" of running depending on whether going uphill or downhill; it is simply due to the constant presence of gravity.

Both Graded and *uniform* attention filters (Chapter 1: figs. 1.3, 1.5) overall showed a pattern of higher relative influence for higher contrasts than for lower contrasts, albeit differences were less significant for bright targets. This was noticeable in the filter slopes of Dark-only, Bright-only, and Attend-to-Dark and Attend-to-Dark conditions. Recall in particular that for dark contrasts, Graded attention filters often matched their targets better than Uniform filters did. Uniform filters also had higher weights for dark targets of high contrast than those of low contrast. This demonstrates patently that certainly for dark targets, subjects gave higher weights to higher contrasts across both Graded and uniform attention conditions. For bright targets, the filter *slope* is suggestive of or at least consistent with a possible trend of increasing weight for higher contrasts. This is more compelling in the Bright-only L attention filter for displays of *eight Weber contrasts* wherein the slope continues rising throughout nearly every contrast level. (Chapter 1: Experiment 2) The filter slope sign remains positive through all filter weights from Weber c: {0.125 - 0.875}. The only slope that is perfectly flat is the tiny segment of the filter connecting Weber c: {0.875 - 1.0}.

Consider evidence from the mean orientation task (Chapter 2): recall that the MO task's attention filters in both the Attend-to-Light and Attend-to-Dark conditions showed higher weights for extreme contrasts than for lower contrasts. The very low Efficiencies support the guided subsampling hypothesis that only a subset of targets are actually used to estimate mean

orientations. It is plausible that the extreme contrasts were more salient and easier to register than lower contrast bars and were more likely to influence judgments in mean orientation estimations. Hence, as performance deteriorates in more challenging attention conditions such as Full-set Attend-All, we expect to see more weight for higher contrast items and less weight for low contrast items. This is what we observed.

Finally, consider the evidence that low level-vision is organized around retinotopic contrast sensitive center-surround cells which fire with frequency in proportion to contrast intensity. It is not hard to imagine that all other attention biasing being equal, their increased rate of input might influence subsequent visual responses within their receptive fields by a winner-take-all competition.

Why did Mean Orientation estimation attention filters deviate from Uniform targets in being graded with high weight for extreme contrasts?

In addition to the arguments in Chapter 2, recall that in Chapter 1, Experiment 1 had across Uniform attention conditions filters for Attend-to-Dark and Attend-to-Light displays that were more graded than were filters for Dark-only, Light-only, and Full-set displays. Moreover, Uniform attention filters for Attend-to-Dark and Attend-to-Light were highly similar to Graded attention filters, especially the Attend-to-Dark filter. Particularly for the MO task, we also observed graded filters in the Attend-to-Dark and Attend-to-Light conditions.

Secondly, in Chapter 1 Uniform attention filters for Dark-only and Light-only displays of 8-dots were often less graded than were filters for 16-dot displays. (Figs. 1.2 & 1.3.) It overall appears that challenging conditions and decreased performance often predict filters which give relatively more weight to salient, intense stimuli. Indeed, the MO task showed much lower Efficiencies than

did Centroid task. Hence, in the selective attention conditions filters for the MO task were more graded than they were for the Centroid task.

Conclusion

In Chapter 1, we found that participants can achieve two distinct types of attention filters: attention filters that (1) give equal weight to all dots regardless of their Weber contrasts, and (2) give weight to dots graded in proportion to contrast intensity. Human vision affords access to three attention filters for centroid estimation: (1) an attention filter that assigns equal weight to all gray-scales in the stimulus, (2) an attention filter that assigns weight near 0 to all negative Weber contrasts and weight to dots of positive polarity graded in proportion to Weber contrast, and (3) an attention filter that assigns weight near 0 to all positive Weber contrasts and weight to dots of negative polarity graded in proportion to the absolute value of Weber contrast. In Chapter 2, we sought to determine whether human vision affords access to the same grayscale-selective attention filters for estimating the centroid as it does for estimating the mean orientation (MO) of a group of bars varying in Weber contrast. In the centroid task, participants achieved attention filters that matched target filters fairly well. In the MO task, filter accuracy was slightly worse. Efficiencies were more than 100% higher in the centroid task than in the MO task. The pattern of Efficiencies across the two tasks and the different attention conditions suggested that participants were using a very different strategy in the MO task than they were in the centroid task. Specifically, the centroid task results suggested that in each condition, participants were applying a filter to the input and computing the centroid of the filter output; by contrast, in the MO task, the results suggested that participants were selecting a sample comprising a fixed number of bars from the target set on each trial and basing their responses on the mean orientation of that sample. In Chapter 3, we found

89

that participants can achieve attention filters which give graded weight inverse to dot contrast intensity, which shows preattentive vision can dampen saliency of the extreme white or black and sharply tune to the dim low contrast item. However, attention filters achieved deviated from the linearly graded target functions in being overly concaved in shape: the relative weight given to the dimmest display items was too high. This suggests that the inverse-graded filters achieved made use of a preattentive mechanism that is sharply tuned to dim items, with responses dropping off rapidly with increasing contrast intensity. Attention filters also show mechanisms subserving these attention filters may be responsive only to a single contrast polarity. Overall, in addition to the discovered mechanisms, other constraints to filters likely include mechanisms of black-white asymmetry and the attention capture of intense black and whites. Taken together, prediction of untested attention filters is plausibly partly predictable.

References

Baldassi, S., & Verghese, P. (2005). Attention to locations and features: Different top-down modulation of detector weights. *Journal of Vision*, *5*(6), 7-7.

Borji, A., & Itti, L. (2012, June). Exploiting local and global patch rarities for saliency detection. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on* (pp. 478-485). IEEE.

Borji, A., & Itti, L. (2013). State-of-the-art in visual attention modeling. *IEEE transactions on pattern analysis and machine intelligence*, *35*(1), 185-207.

Borji, A., Sihite, D. N., & Itti, L. (2013). What stands out in a scene? A study of human explicit saliency judgment. *Vision research*, *91*, 62-77.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision research*, *51*(13), 1484-1525.

Chubb, C. F., & Talevich, J. (2001, March). Attentional control of texture judgments. In *INVESTIGATIVE OPHTHALMOLOGY & VISUAL SCIENCE* (Vol. 42, No. 4, pp. S864-S864). 9650 ROCKVILLE PIKE, BETHESDA, MD 20814-3998 USA: ASSOC RESEARCH VISION OPHTHALMOLOGY INC.

Dakin, S. C. (2001). Information limit on the spatial integration of local orientation signals. *JOSA A*, *18*(5), 1016-1026.

Drew, S. A., Chubb, C. F., Ehrlich, T., Rubin, T., & Sperling, G. (2008). Binary versus graded filters for selectively attending to dots of different contrasts. *Journal of Vision*, *8*(6), 881.

Drew, S. A., Chubb, C. F., & Sperling, G. (2010). Precise attention filters for Weber contrast derived from centroid estimations. *Journal of vision*, *10*(10), 20-20.

Haenny, P. E., & Schiller, P. H. (1988). State dependent activity in monkey visual cortex. *Experimental Brain Research*, *69*(2), 225-244.

Hoel, P. G., Port, S. C., & Stone, C. J. (1971). Introduction to statistical theory.

Inverso, M., Sun, P., Chubb, C., Wright, C. E., & Sperling, G. (2016). Evidence against global attention filters selective for absolute bar-orientation in human vision. *Attention, Perception, & Psychophysics*, *78*(1), 293-308.

Jin, J., Wang, Y., Lashgari, R., Swadlow, H. A., & Alonso, J. M. (2011). Faster thalamocortical processing for dark than light visual targets. *The Journal of Neuroscience*, *31*(48), 17471-17479.

Kremkow, J., Jin, J., Komban, S. J., Wang, Y., Lashgari, R., Li, X., ... & Alonso, J. M. (2014). Neuronal nonlinearity explains greater visual spatial resolution for darks than lights. *Proceedings of the National Academy of Sciences*, *111*(8), 3170-3175.

Lankheet, M. J., & Verstraten, F. A. (1995). Attentional modulation of adaptation to two-component transparent motion. *Vision research*, *35*(10), 1401-1412.

Ling, S., Liu, T., & Carrasco, M. (2009). How spatial and feature-based attention affect the gain and tuning of population responses. *Vision research*,*49*(10), 1194-1204.

Liu, T., Larsson, J., & Carrasco, M. (2007). Feature-based attention modulates orientation-selective responses in human visual cortex. *Neuron*,*55*(2), 313-323.

Lu, Z. L., & Sperling, G. (2012). Black–white asymmetry in visual perception.*Journal of vision*, *12*(10), 8-8.

Martinez-Trujillo, J. C., & Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology*, *14*(9), 744-751.

Maunsell, J. H., Sclar, G., Nealey, T. A., & DePriest, D. D. (1991). Extraretinal representations in area V4 in the macaque monkey. *Visual neuroscience*, *7*(06), 561-573.

McGowan, J. W., Kowler, E., Sharma, A., & Chubb, C. (1998). Saccadic localization of random dot targets. *Vision research*, *38*(6), 895-909.

Müller, M. M., Andersen, S., Trujillo, N. J., Valdes-Sosa, P., Malinowski, P., & Hillyard, S. A. (2006). Feature-selective attention enhances color signals in early visual areas of the human brain. *Proceedings of the National Academy of Sciences*, *103*(38), 14250-14254.

Pomplun, M. (2006). Saccadic selectivity in complex visual search displays. *Vision research*, *46*(12), 1886-1900.

Saenz, M., Buraĉas, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision research*, *43*(6), 629-637.

Serences, J. T., & Boynton, G. M. (2007). Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron*, *55*(2), 301-312.

Schiller, P. H., & Lee, K. (1991). The role of the primate extrastriate area V4 in vision. *Science*, *251*(4998), 1251-1254.

Shih, S. I., & Sperling, G. (1996). Is there feature-based attentional selection in visual search?. *Journal of Experimental Psychology: Human Perception and Performance*, *22*(3), 758.

Solomon, J. A., May, K. A., & Tyler, C. W. (2016). Inefficiency of orientation averaging: Evidence for hybrid serial/parallel temporal integration. *Journal of vision*, *16*(1), 13-13.

Sun, P., Chubb, C., Wright, C. E., & Sperling, G. (2016). The centroid paradigm: Quantifying feature-based attention in terms of attention filters. *Attention, Perception, & Psychophysics*, *78(2)*, 474-515.

Sun, P., Chubb, C., Wright, C. E., & Sperling, G. (2016). Human attention filters for single colors. *Proceedings of the National Academy of Sciences*, *113*(43), E6712-E6720.

Treue, S., & Trujillo, J. C. M. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, *399*(6736), 575-579.

Yeh, C. I., Xing, D., & Shapley, R. M. (2009). "Black" responses dominate macaque primary visual cortex V1. *The Journal of Neuroscience*, *29*(38), 11753-11760.

Zhou, X., Chu, H., Li, X., & Zhan, Y. (2006). Center of mass attracts attention. *Neuroreport*, *17*(1), 85-88.

Appendix A

I.

1. Let N be the number of different Weber contrasts c in a fully-loaded display. L_X (L_Y) be the $N_{trials} \times (N+2)$ matrix whose $(i,j)^{th}$ entry is (a) the sum of the x-locations (y-locations) of all dots of Weber contrast w_j presented on fully-loaded trial i for $j \leq N$, and (b) is 1 (0) for $j=N+1$, or (c) is 0 (1) for $j=N+2$.

2. Let R_X (R_Y) be the column vector of length N_{trials} whose i^{th} entry is the x-coordinate (y-coordinate) of the participant's response on fully-loaded trial i.

Then

1. Form the $2N_{trials} \times (N+2)$ matrix M by appending the matrix L_Y to the bottom of L_X.

2. Form the vector R of length $2N_{trials}$ by appending R_Y to R_X.

3. Perform linear regression to derive the weights W minimizing $SS_{Residual} = ||MW\text{-}R||^2$.

Then, writing $n_j, j = 1,2,\ldots,N$ for the number of items of type j in each fully-loaded stimulus cloud, D (Data-Drivenness) is estimated by

$$D = \sum_{j=1}^{N} W_j \, n_j \qquad (4)$$

For $j=1,2,\ldots,N$, $f(j)$ is estimated by

$$f(j) = \frac{W_j}{\sum_{i=1}^{N} W_i} \qquad (5)$$

σ^2 is estimated by taking

$$\sigma^2 = \frac{SS_{residual}}{df}, \quad df = 2N_{trials} - (N+2) \qquad (6)$$

(where model parameters $x_{default}$, $y_{default}$, and D absorb 3 degrees of freedom and the attention filter f absorbs only N-1 additional degrees of freedom because it is constrained to sum to 1). Finally, $x_{default}$ and $y_{default}$ are estimated by taking

$$x_{default} = \frac{W(N+1)}{1-D} \quad \text{and} \quad y_{default} = \frac{W(N+2)}{1-D} \quad (7)$$

A formal justification of this modeling method and a description of the methods used to estimate confidence intervals for model parameters is provided in Sun et al. (2016).

Quantifying performance. Note that actual judged centroid location R_X, R_Y can deviate from target performance (Eq. (2)) for several reasons. We quantify these deviations using:

1. *Imperfect Data-Drivenness*. Data-Drivenness is a Bayesian parameter that describes how much the participant's response is driven by the stimulus or by a preconception of where centroids most often occur ($x_{default}$, $y_{default}$). The Data-Drivenness D of the participant's responses can be less than 1.

2. *Filter mismatch*. The attention filter f achieved by the participant can deviate from the target filter f_{targ}. This will cause the responses of the participant to deviate systematically from the correct responses. To quantify the degree to which the responses of the participant are immune from this sort of error, we use a descriptor of the participant's attention filter f, called *Filter-fidelity*. Filter-fidelity is a measure of how far away from the ideal filter (ϕ) the achieved filter (f) is. Specifically,

Filter-fidelity $= 1 - \frac{||f-\phi||}{||f_{Worst}-\phi||}$ (5)

96

where f_{Worst} is a worst possible attention filter for the condition with target filter ϕ, i.e., f_{Worst} is an attention filter that maximizes the expected response error. Specifically, in each of the conditions tested in the current study, f_{Worst} is obtained by choosing a Weber contrast c_{min} for which $\phi(c_{min})$ is minimal, setting $f_{Worst}(c_{min})=1$ and $f_{Worst}(c)=0$ for all c other than c_{min}. Notice that if the participant perfectly matches this target filter, *Fidelity* will be 1 whereas if the participant achieves a minimally effective filter, *Fidelity* will be 0.

3. *Random noise*. The standard deviation σ of the random variables *Noise$_X$* and *Noise$_Y$* is nonzero. This will cause the responses of the participant to deviate randomly from the correct responses. Although σ itself could be used to gauge the amount of random error corrupting the participant's responses, this model parameter is difficult to interpret because it depends on several factors (such as the size of stimulus display clouds) that are likely to vary across different experiments. To facilitate comparison of results across experiments, we use a descriptor called *Efficiency* to quantify immunity to random error. Efficiency is the greatest lower bound on the proportion of display items that the model needs to yield the performance accuracy of the participant. If the trial-to-trial random errors *Noise$_X$* and *Noise$_Y$* were due solely to missing (i.e., failing to include) some of the items in the display in computing the centroid, then the model would need to miss a proportion $p = 1 -$ Efficiency of the items on each trial in order for *Noise$_X$* and *Noise$_Y$* to have standard deviation σ. Thus, if the participant were to attain an Efficiency of 0.75, this would imply that he/she is including, on average, in his/her centroid computation at least three quarters of the items in the stimulus display. It should be emphasized, however, that if some of the random noise corrupting responses were due to some other error sources such as (i) early perceptual noise or (ii) instability in the centroid computation or (iii) motor noise, then the

actual proportion of items included in the centroid computation itself would be higher than

Efficiency.