

Human attention filters for single colors

Peng Sun^{a,b,1}, Charles Chubb^a, Charles E. Wright^a, and George Sperling^{a,c,1}

^aDepartment of Cognitive Sciences, University of California, Irvine, CA 92697; ^bDepartment of Psychology, New York University, New York, NY 10003; and ^cDepartment of Neurobiology and Behavior, University of California, Irvine, CA 92697

Contributed by George Sperling, September 6, 2016 (sent for review June 3, 2016; reviewed by Julian Hochberg and Richard M. Shiffrin)

The visual images in the eyes contain much more information than the brain can process. An important selection mechanism is feature-based attention (FBA). FBA is best described by attention filters that specify precisely the extent to which items containing attended features are selectively processed and the extent to which items that do not contain the attended features are attenuated. The centroid-judgment paradigm enables quick, precise measurements of such human perceptual attention filters, analogous to transmission measurements of photographic color filters. Subjects use a mouse to locate the centroid—the center of gravity—of a briefly displayed cloud of dots and receive precise feedback. A subset of dots is distinguished by some characteristic, such as a different color, and subjects judge the centroid of only the distinguished subset (e.g., dots of a particular color). The analysis efficiently determines the precise weight in the judged centroid of dots of every color in the display (i.e., the attention filter for the particular attended color in that context). We report 32 attention filters for single colors. Attention filters that discriminate one saturated hue from among seven other equiluminant distractor hues are extraordinarily selective, achieving attended/unattended weight ratios >20:1. Attention filters for selecting a color that differs in saturation or lightness from distractors are much less selective than attention filters for hue (given equal discriminability of the colors), and their filter selectivities are proportional to the discriminability distance of neighboring colors, whereas in the same range hue attention-filter selectivity is virtually independent of discriminability.

feature-based attention | color discrimination | selective attention | centroid estimation

The visual world on our planet is incredibly complex. The main mechanism for selecting a subset of the environmental visual information to process is orienting our head and body and pointing our eyes toward the most significant locations. Even then, each retinal image contains enormously more information than the brain can process. To further constrain the flow of information, perceptual attention processes select locations in space, intervals in time, objects that contain particular features, and ultimately particular complex objects for further processing. Here, we focus on attention to a particular color, an instance of feature-based attention (FBA). It is well established that FBA operates broadly across space, heightening sensitivity to the attended feature even at locations that are irrelevant to the task at hand (1–4). We exploit this global property of FBA in a centroid paradigm to derive human attention filters for color. Just as the physical description of a color filter describes the relative transmission of the filter for each wavelength of light, a color-attention filter describes the relative effectiveness with which each color in the retinal input ultimately influences performance.

Using the Centroid Paradigm to Derive Attention Filters

In a centroid task, a subject views a briefly presented stimulus containing a cloud of dots. The task is to then mouse-click the apparent center of the dot cloud. Drew et al. (5) first used a centroid task to estimate attention filters. The procedure was substantially refined by Sun et al. (6) into the centroid paradigm used here. Subjects attempt to locate the centroid (center of gravity) of only those items that have a to-be-attended feature (targets) and attempt to ignore the nontargets (distractors). In the present experiments, the targets are three dots of the to-be-attended color,

and the distractors are 21 dots, three each of seven other colors (Fig. 1A). Insofar as subjects can perfectly attend to targets and completely ignore distractors, and can perfectly compute the centroids of the attended targets,* the judged centroid location would be the centroid of the targets plus random response error (noise) that is independent of the distractors' colors and locations. In reality, however, reported centroids are influenced to some extent by every item in the stimulus, reflecting the subject's inability to selectively attend exclusively to targets and to completely ignore distractors. By analyzing judged centroid locations from 100 or so trials, the relative weight of each color's contribution to the centroid judgment—the attention filter—can be inferred. Because a trial takes 3–4 s, an attention filter can be roughly estimated in 5–7 min. Of course, more trials would give greater precision.

The inference process is illustrated by the model of selective attention illustrated in Fig. 1B particularized for the present experiments. On each trial, a stimulus containing three dots of each of eight colors is presented. The attention instruction is to assign equal weight to dots of the to-be-attended color (three targets) and zero weight to the remainder (21 distractors). Let the eight colors be represented as C_i , $i = \{1, 2, \dots, 8\}$, and let the target color be C_k . Then, the target attention filter ϕ_k is $\phi_k(i) = 1$ if $i = k$ and $\phi_k(i) = 0$ if $i \neq k$ (i.e., by convention, the total attention weight is 1.0).

The model's representation of a human subject attending to color C_k has an attention filter $f_k(i)$ that causes the model's predicted centroids to best match the observed centroids. The optimally predictive model filter $f_k(i)$ is obtained by simple linear regression (*Materials and Methods*).

The model's optimally predictive filter $f_k(i)$ is called the observed attention filter. It typically is a very good predictor of a subject's observed centroid judgments.† Therefore, we say for short that $f_k(i)$ is the subject's attention filter for attending to color C_k in that context.

Significance

The eyes present the brain much more information than it could possibly process. One important way to prioritize information is by selective attention to features, processing only items containing the attended features and blocking others (i.e., forming an attention filter). Here we demonstrate an extremely efficient paradigm and a powerful analysis to quantitatively measure, as accurately as one might measure physical color filters, 32 such human attention filters for single colors. These data are an essential basis for a theory of attention to color. The centroid paradigm itself, because it quickly and quantitatively characterizes basic attention processes, has numerous applications.

Author contributions: P.S., C.C., C.E.W., and G.S. designed research, performed research, analyzed data, and wrote the paper.

Reviewers: J.H., Columbia University; and R.M.S., Indiana University.

The authors declare no conflict of interest.

¹To whom correspondence may be addressed. Email: peng.sun@nyu.edu or sperling@uci.edu.

*A more realistic model of a subject's imperfect centroid computation improves the prediction of the observed centroid judgments, but we find that it does not improve the estimation of the subject's attention filter, and so we omit that complication.

†See the first footnote.




Fig. 1. The centroid paradigm: procedure and analysis. (A) Example of a subsection of a display that shows the centroid of orange targets among distractors of different colors. (B) Model of the attention-weighted centroid computation. The input image consists of dots of different colors. Filter weights shows the filter weights induced by attention instructions (e.g., attending to orange) when performing the task. The filter is derived such that it best predicts the subject's centroid judgment. The centroid computation computes the center of gravity (x,y) of the dot locations, weighting each location by the filter weight for the color of the dot at that location. The output x,y is the model's prediction for that subject's judged centroid for this particular input image.

Overview

We measured attention filters for single colors. The significant perceptual attributes of color are hue, saturation, and lightness. The attributes of isolated colors on a neutral background are fully described by the relative stimulation of long-, medium-, and short-wavelength-sensitive retinal cones, here represented as a 3D space (Fig. 2A). In this representation, hue is represented as a circle in an equiluminant plane with two orthogonal axes, red-green (R-G) and blue-yellow (B-Y); the other (nonorthogonal) axis is black-white (Bl-Wh).

In different blocks of trials, we tested the four sets of eight colors embedded in the 3D cone color space shown in Fig. 2A, which also shows an example stimulus from each set. Using the centroid paradigm, attention filters were derived for each of 32 different attended colors, each in the presence of the seven distractor colors from its set. Fig. 3B shows the results: 32 measured selective-attention filters for each of five subjects. Attention filters for colors on the hue circle are remarkably precise, similar for all subjects, and much more selective than attention filters along the saturation axes (R-G and B-Y) or lightness axis (Bl-Wh).

To determine whether the better hue-attention filters were due to stimuli on the hue circle being more discriminable than the stimuli on the axes, Exp. 2 measured the discriminability distances between the different stimuli in terms of the number of just-noticeable differences (JNDs) separating them. Filter selectivity was closely related to JND distances for stimuli along the three axes but not for hue stimuli. In Exp. 3, colors were arranged around the hue circle but with the JND spacing of the axis colors. The JND arrangement along the hue circle was enormously more advantageous than the same JND spacing along the Bl-Wh axis, confirming Exp. 2: In the color-attention system, attention filters for a specific hue are far more effective and follow different rules than attention filters that select for a color embedded within one of the other dimensions of color (saturation and lightness).

General Procedures

There were eight attention conditions for each set of stimuli, each corresponding to attending to (i.e., judging the centroid of dots of) just one the eight different target colors and ignoring the others. Each attended color was tested in separate counterbalanced blocks of trials. Randomly chosen stimuli from the same set were used in all eight attention conditions within that set. Because the stimuli were chosen from the same urn for all eight different attention conditions within a set, any statistically reliable differences in results (i.e., differences between the observed attention filters) are due entirely to the different attention instructions. The trial procedure is illustrated in Fig. 2 C-F (see *Materials and Methods* for details).

Exp. 1: Attention Filters for Four Sets of Eight Single Colors

Fig. 3A shows the attention filter achieved by a naive subject, S1, for a typical target color in each of the four stimulus sets. The illustrated

target color was arbitrarily chosen as the third (of eight) colors in each stimulus set. For the hue set, S1's attention filter quite closely approximates the target attention filter, but S1's attention filters are far from the target filters for the other stimulus sets.

Sun et al. (6) introduced five quantitative measures to evaluate attention filters, two of which are used here. The "selectivity ratio" is defined as the filter weight of the target color divided by the mean weight of all of the distractor colors. A selectivity ratio of 20 for a hue filter means that a dot of the target color has 20 times the influence on the centroid judgment as a distractor-colored dot. Efficiency is determined by giving a perfect centroid computer the subject's achieved attention filter, showing the computer the same stimuli as the subject except that some stimulus dots are randomly removed. The fraction of remaining stimulus dots that the centroid computer needs to perform with the same accuracy as the subject is the subject's "efficiency."[‡] An efficiency of 90% for the hue attention filter means that when 2 of 24 dots (8.3%) are randomly removed from the stimulus the perfect centroid computer is slightly more accurate than the subject, but when three dots (12.5%) are removed it does worse. Alternatively, 90% efficiency means that the average number of dots used for the centroid judgment is 2.7 (of 3). Clearly, subject S1's 20:1 selectivity and 90% efficiency describe an extremely good hue-attention filter.

Compared with attention filter for stimulus 3 in the hue set, attention filters for stimulus 3 in the R-G, B-Y, and Bl-Wh stimulus sets are significantly worse both in terms of selectivity and efficiency. Compared with the narrow tuning function (selectivity = 20) of S1's attention filter for stimulus 3 in the hue set, the tunings of filters for S1's R-G, B-Y, and Bl-Wh stimulus sets are quite broad, consistent with their relatively lower selectivity ratios: 3.8, 4.1, and 2.8 (Fig. 3). When the perfect centroid computer is handicapped with S1's attention filter, it can match the accuracy of S1's judged centroids by using only slightly more than half the stimulus dots (i.e., efficiencies of 55%, 61%, and 51% for the three axis stimulus sets).

Attention filters for all subjects and stimulus sets are shown in Fig. 3B. For the colors on the hue circle, all of the color-attention filters are sharply tuned. The attention filters for the three color sets on the three color axes are similar to each other and are much less selective and much more variable than attention filters for hue. Attention filters at the two ends of a color axis and the attention filters for colors closest to the gray background are the most selective filters.

Fig. 3C shows a summary of the attention filters' selectivities (Fig. 3C, *Top*) and efficiencies (Fig. 3C, *Bottom*), averaged across all subjects. Filter selectivities and efficiencies for colors lying on the hue circle are generally greater than those for colors on the axes. Again, for color-attention filters for colors on the three axes (R-G, B-Y, and Bl-Wh), selectivities and efficiencies are greater for colors at the two

[‡]See Sun et al. (6) for details.

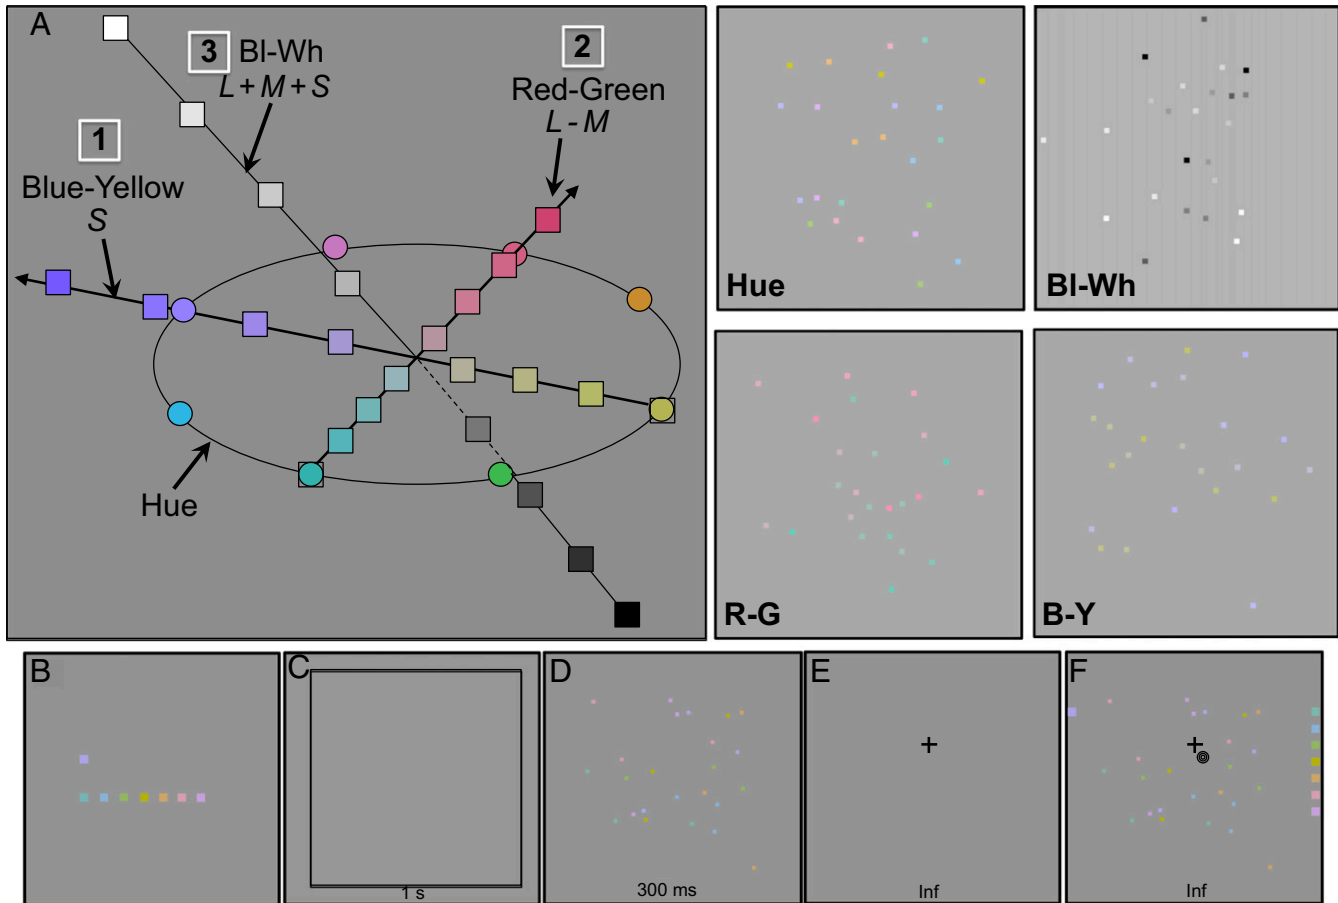


Fig. 2. A diagrammatic representation of the four sets of stimuli used in Exp. 1, sample stimuli, and the testing procedure. (A) The stimuli represented in a 3D space defined in terms of retinal cones. Circles: the eight hue-set stimuli that lie equally spaced on a circle in an equiluminant plane. Squares: 1, stimuli in the B–Y set that lie on a diameter defined by the short-wavelength (S, blue cone) sensitivity function; 2, the R–G-set stimuli that lie on a diameter perpendicular to the blue-cone axis corresponding to the long-wavelength (L, red) minus medium wavelength (M, green cone) sensitivity functions; 3, the achromatic set stimuli that lie on a line (L+M+S), commonly called the BI–Wh axis. Note that the achromatic line is not perpendicular to the equiluminant plane (the perpendicular axis to the equiluminant plane is the L+M axis, which is yellowish, not achromatic). Stimuli along each radius are spaced equally in cone space to use the maximum color range available on the monitor. An example stimulus from each of the four stimulus sets: hue, BI–Wh, R–G, and B–Y. (B) Example of an instruction display shown before a block session starts. Subjects are instructed to mouse-click the centroid of the three dots of the single color in the top row and to ignore the other dots (seven colors in the bottom row). (C–F) Stimulus sequence on each trial. (C) Fixation frame, 1 s. (D) Stimulus, 300 ms. (E) Response display shown until a response is made. (F) Feedback frame that shows the stimulus, the target centroid, the subject’s response, the target color (extreme left), and the to-be-ignored colors (extreme right). The feedback frame is displayed until the subject initiates the next trial.

ends of the three diagonal axes and for colors that are near the background gray. A possible explanation for the parametric properties of the color-attention filters is explored in the next two experiments.

Exp. 2: The Dependence of Color-Attention Filters on Target–Distractor Discriminability

Color-attention filters for some colors are more selective than for others. Is this simply because these target colors are further away from the distractor colors in some perceptual space? An objective measure of perceptual distance between two colors is the number of JNDs separating two colors in color space. To measure JND distances, the space along the (invisible) line in color space between each pair of neighboring stimulus colors was sampled at various points. At each sample point in color space two nearby colors were chosen and assigned to two stimulus dots. We measured the local JND (i.e., the minimum distance in color space between the nearby colors that enabled 82% discrimination of the color difference between the two dots).

Interpolation from the collection of local JNDs enables estimating the number of JNDs that separate two stimulus colors. JND distances were measured for all adjacent color pairs for

three of the original five subjects—two were unavailable because they had graduated. These JND measurements were an order of magnitude more time-consuming than the attention-filter measurements. The JND local isolation of a target color is defined as $1/(1/d_1 + 1/d_2)$, where d_i is the JND distance between it and its two adjacent (i.e., most confusable) colors. For brevity we omit “local” and refer simply to JND isolation. For the end colors on the three axes the JND isolation is simply the JND distance from the only adjacent color, conforming with the above definition.

Because the JND isolation measures the discriminability of the target color against its adjacent color(s), to make comparisons plausible we computed a new filter selectivity measure as the target weight divided by the mean weight of just-nearest distractors. We call this new measure local selectivity and compare it to JND isolation for each color. Fig. 4A shows the comparisons between color-attention filter local selectivity (target/adjacent distractor weight) and JND isolation. The mean local selectivity and the mean isolation ratio averaged for the three subjects are given in Fig. 4B.

Whereas Fig. 4A and B show local selectivity and JND isolation of color-attention filters as functions of target color, Fig. 4C shows directly the relationship between the two dependent variables, JND

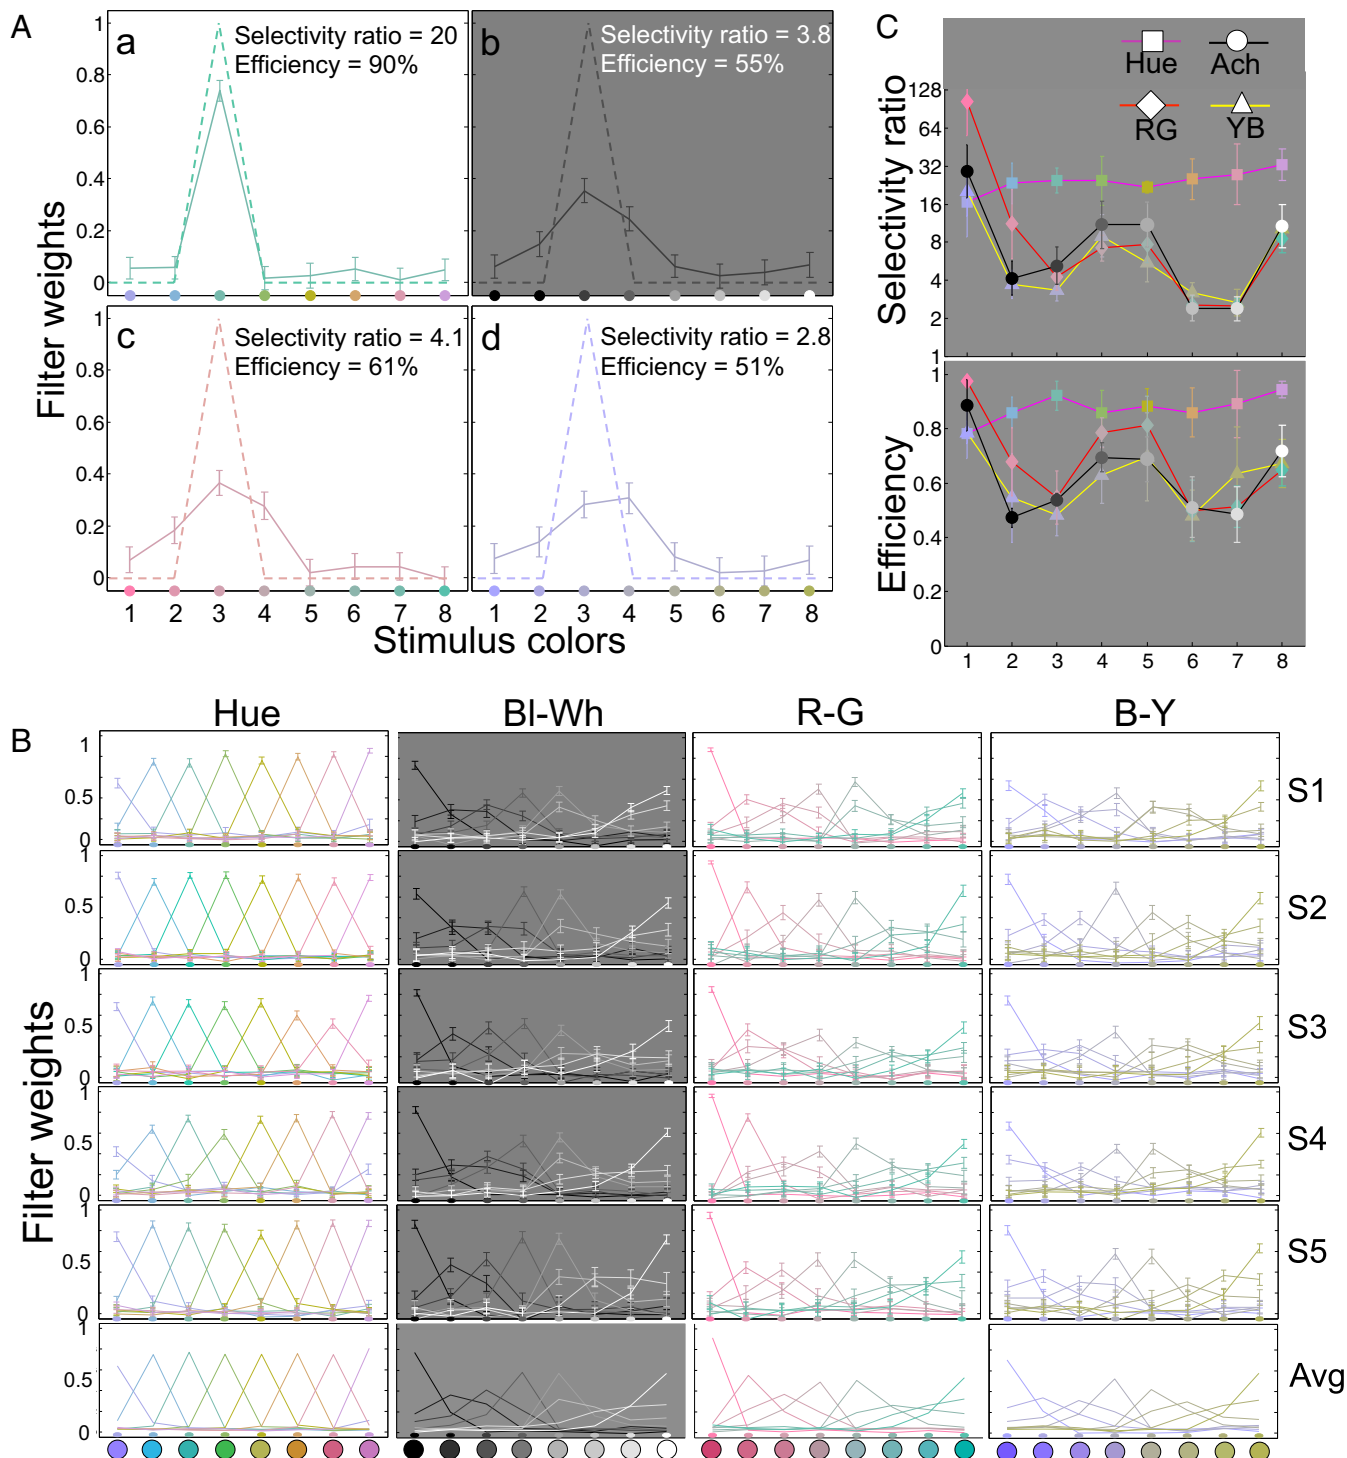


Fig. 3. Five subjects' attention filters, filter selectivity ratios, and filter efficiencies for 32 attention conditions. In all panels, the eight colored circles on the abscissa represent the eight colors in the particular set. (A) Four example target and achieved attention filters; each filter corresponds to a particular target color in one of the four stimulus sets: (a) hue set, (b) achromatic (BI-Wh) set, (c) R-G set, and (d) Y-B set. Dashed and solid curves are target and the subject's achieved filters. Filter curves are painted in the corresponding target colors. Error bars throughout Fig.3 represent 95% confidence intervals. (B) Attention filters of five subjects (separated in rows) for four sets of stimuli (separated in columns). Columns (from left to right) are hue set, BI-Wh set, R-G set, and Y-B set. Within each panel, each colored curve corresponds to the achieved attention filter for that particular target color. Filter curves for all of the colors within a stimulus set are shown in one panel although they are derived from data collected in separate blocks of trials. (C) Mean filter selectivities (target/distractor filter weights) and efficiencies (fraction of targets used) averaged across five subjects. The colors of points are the corresponding target colors; the line colors are arbitrary.

isolation and local selectivity. Fig. 4C makes it evident that whereas there is a strong positive correlation between local selectivity and JND isolation for axis colors, there is no such correlation for hue

colors, even in the same range of JND isolations. Exp. 3 further explores this difference between color-attention filters for hue and color-attention filters for the other dimensions of color.

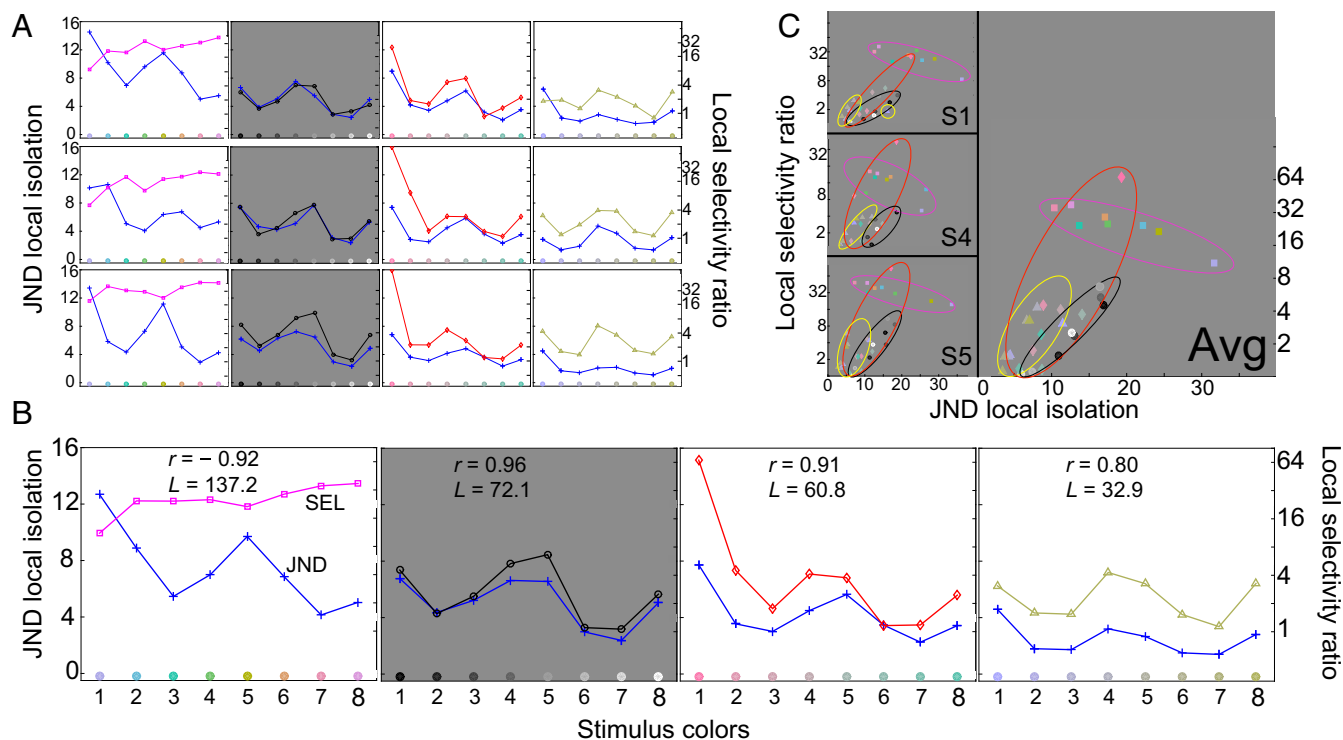


Fig. 4. Exp. 2: target attention filter local selectivity ratios and target JND isolation for 32 targets and three subjects. (A) Color attention filter local selectivity and color JND isolation (right and left ordinates) as a function of target color (abscissa) for three subjects (rows). In each panel, the blue curve (with data points indicated by +) plots JND isolation (see text for definition). Magenta, black, red, and yellow curves connect selectivity ratios for hue colors, contrasts, RG and BY colors, respectively. From left to right, subfigures are for hue, BI-Wh, RG, and YB color sets. (B) Average (over the three subjects) selectivity ratios and JND isolations. The correlations r between these two average measures are given in each panel. L is the average overall length in JNDs of the trajectory formed by each set of colors. (C) Scatter plots illustrating the relationship between target JND isolation (abscissas) and target local selectivity ratio (ordinates) for the three subjects and the average. The color of each point indicates the target color. Data points from the same stimulus set are grouped by ellipses colored as follows: hue, magenta; BI-Wh, black; RG, red; and BY, yellow.

Exp. 3: Color-Attention Filters for Spacing-Adjusted Hue Colors

One possible explanation for the advantage of color-attention filters for the saturated hue-circle colors versus the R-G, B-Y, and BI-Wh axis colors is that hue colors are more widely separated in perceptual color space. In Exp. 3, for each of the three subjects eight new equiluminant colors were chosen on the hue circle so that the JND distances between them were exactly the same as the JND distances between the eight stimuli on the BI-Wh axis. Selectivity ratios of the color attention filters for these eight spacing-adjusted colors are shown in Fig. 5.

The enormous attention-selectivity advantage of hue color filters over BI-Wh attention filters is still observable despite a JND spacing of the eight hue colors that is exactly equivalent to the spacing of the eight BI-Wh contrasts. As in Exp. 2, the selectivity of attention filters for specific hues seems to be virtually independent of the discriminability distance of the nearest neighbors, quite different from attention filters for specific BI-Wh contrasts. The special properties of attention filters for specific hues most likely reflect that greater ecological validity of hue (versus saturation or lightness) as a clue to the true nature of a surface that demands attention.

Discussion

The behavioral effects of FBA have most often been characterized in terms of either a reduced reaction time, when searching for attended features, or a reduced threshold for detecting and discriminating attended features (7, 8; see ref. 9 for a review). These measures reveal how much more effective FBA makes the attended feature compared with when the same feature is not attended.

However, other important aspects of FBA have been left largely unaddressed. For example, to what degree does FBA influence the effectiveness of irrelevant visual features? These aspects of FBA are likely to be especially important in natural settings in which irrelevant features abound.

Paradigmatic Efficiency. To probe feature attention quantitatively and efficiently we developed the centroid paradigm (5, 6). The centroid paradigm enables the simultaneous measurement of the relative weight of target items in the presence of other items in a neutral attention condition and in an attention condition when a particular target type is attended and the other item types are distractors. In both situations, the centroid paradigm provides the relative weights of all of the other item types relative to the target type(s), an enormous improvement in both the amount of information gained and in the paradigmatic efficiency of gaining this information compared with previous FBA paradigms. In principle, a human selective-attention filter for a hue can be determined as accurately as the typical physical specification of an optical color filter for a camera. Measuring 32 color-attention filters would have been impossible without such a highly efficient paradigm. Some interesting properties of FBA emerged as a result.

Extreme Selectivity, the Special Status of Hue. The principal result of the experiments is that attention to hue, as measured in the centroid paradigm, is extremely selective. In Exp. 1, selectivity (the weight ratio of the attended-hue dots divided by the average weight of the seven distractor-hue dots) exceeded 20:1 in most cases. In Exp. 2, the local selectivity (weight ratio of the attended hue to the most similar distractor hue) exceeded 10:1 for most

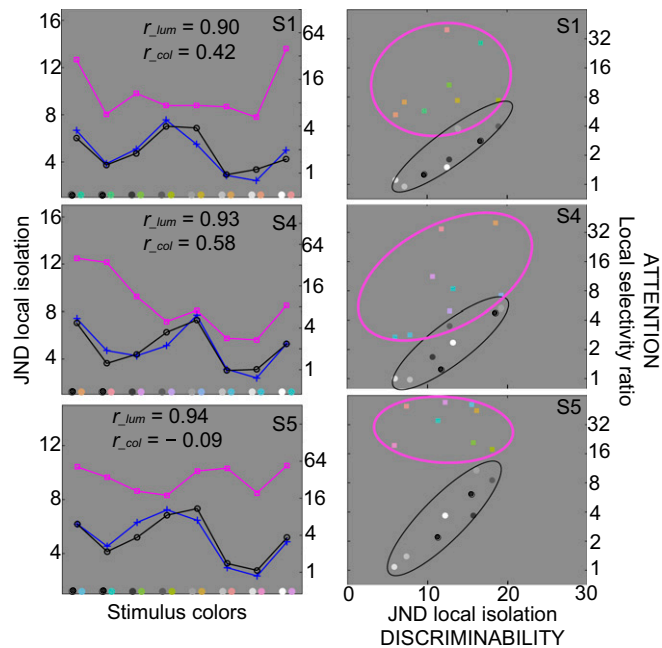


Fig. 5. Exp. 3: selectivity ratio and JND-isolation derived from attention filters for a BI-Wh color set and an equivalent JND-spacing-adjusted hue color set; data for three subjects. (*Left*) Black curves are JND isolation for the BI-Wh stimuli and also the JND isolation for hues spaced along the color circle to have exactly the same JND isolation as a function of target color. Blue curves are the selectivity ratio for BI-Wh targets, which is highly correlated with JND-isolation. Magenta curves are the selectivity ratio for the hue targets. (*Right*) Selectivity ratio versus JND isolation for the two equivalently spaced sets of color targets. The gray levels and colors of the points indicate the contrasts and colors of the target stimuli. Selectivity correlates strongly with JND isolation for BI-Wh targets but not for identically JND-spaced hue targets.

attended hues. When equiluminant hues are matched in discriminability distance to BI-Wh stimuli that differ in contrast, the average attention selectivity for hue is 9.6 times greater (averaged for the three tested subjects). As noted above, the great attention advantage of selectivity for a hue versus selectivity for a value within a BI-Wh, R-G, or B-Y axis probably is ecological: Because of reflectance differences caused by shadows, angle of lighting, and angle of view, hue is a more reliable cue to material identity than either lightness or saturation.

A second significant special property of attention to a hue is that attention selectivity for a hue is almost independent of the JND-discriminability distance of neighboring distractor hues, whereas in the same discriminability range of the BI-Wh, R-G, and B-Y stimulus sets attentional selectivity and JND distractor distance are highly correlated. When the JND distances between hues were reduced in Exp. 3, attention selectivity started to show enhanced dependencies on JND discriminability for two out of three subjects. It is likely that JND discriminability determines attentional selectivity for very closely spaced colors, but JND discriminability quickly becomes much less relevant when hues are more discriminable from each other. What determines attentional selectivity for hues is an important question but is beyond the scope of the current study.

Neural Enhancement of an Attended Feature. A dominant theory of FBA stems from the neural correlates of FBA (10). The response magnitude (impulses per second) of neurons in the middle temporal cortical area of macaque monkeys was measured for stimuli moving in attended, neutral, and unattended motion directions. When a monkey attended to the preferred motion direction for a neuron, the neuron's response magnitude increased by a factor of

about 1.15 versus neutral attention; attending to the antipreferred motion direction reduced neural responses by 0.94. It was proposed that FBA had a multiplicative (amplifying) effect on neural response magnitude. There are similar findings of such small (relative to the present experiments) but statistically significant effects of attention in the amplitude of responses in single neurons (11), EEG waveform amplitudes (12, 13), and on fMRI BOLD response (1).

Different Processing Levels: Attention to Color Precedes Attention to Motion. The multiplicative effect of attention on neurons that report direction of motion operates at a different brain processing level than the effect of attention to color. Neurons that are influenced directly by attention to color represent the inputs to motion-computing neurons. Blaser et al. (14) measured attention to a color (red versus green) in an ambiguous motion paradigm and also found a relatively small amplification factor (typically 1.3, range 1.2–2.2). The enormously smaller effect of attention to hue in an ambiguous motion paradigm versus a centroid task suggests that attention selectivity varies enormously between tasks that involve different brain processes. A critical difference between the centroid task and the ambiguous motion task of Blaser et al. (14) is that grouping of items may be very important in the centroid task (15, 16) but is irrelevant in the ambiguous motion task. That is, FBA seems to have an especially big influence on grouping processes.

Different Processing Levels (Multiple Object Tracking Versus Statistical Summary Representation) Depending on Number of Target Items.

Typically, a centroid judgment is assumed to be a statistical summary representation (SSR), that is, a statistic that accurately describes a property of a group of items even when there are so many items that the subject has accurate information only about few, if any, of the individual items (17). For example, in judging the centroid of 16 items, subjects achieve efficiencies of about 0.8 (6), which means that an ideal detector would have to know the precise location of $0.8 \times 16 = 12.8$ items to match the subjects' performances. This is three or four times the number of items that subjects can identify the locations of. However, when the number of items is small a different process comes into play. For example, in a search experiment, there typically is only one target item and the subject typically knows the location and other properties of the single target.

Inverso et al. (18) found that subjects are easily and accurately able to find the centroid of two vertical lines randomly placed among two horizontal distractors and, similarly, the centroid of two horizontal lines among two vertical distractors. Very surprisingly, subjects were unable to find the centroid of four, six, or eight verticals among an equal number of distractors, although they succeeded easily with black versus white attentional selection with the same stimulus sets. The point here is not the point of Inverso et al. (18), that there is something uniquely difficult about absolute vertical-versus-horizontal attentional selection, but rather that there is a different judgment process when the number of targets is less than four.

In various extensively studied tracking procedures in which a number of targets move randomly among randomly moving distractors, subjects typically can track about four targets (19, 20); they track fewer targets when displays are more complex (21). The number of targets in all of the present attention experiments was three. Therefore, we infer that the subjects' judgments probably did not exclusively rely on SSRs—if at all—but also involved their ability to keep track of the locations of three targets. This agrees with our subjects' introspective reports. Those queried said they were aware of the actual locations of the target dots. This does not change the significance of the color-attention filters that were observed; rather, it informs the processes by which the centroids were computed. Whether the same color-attention filters would be observed in paradigms that involved larger numbers of dots and required SSRs is a matter for future investigation.

Similar Processing of Dark and Light Target Items. It is well documented (22) that for the same contrast ratio magnitude of Weber contrasts darker-than-background stimuli (“dark”) are more discriminable than lighter-than-background stimuli (“light”). Figs. 4C and 5 show the measured JND separations between both adjacent-contrast light and adjacent-contrast dark spots and, for the same pairs of adjacent dot contrasts, the local attention selectivity in centroid judgments. Consider corresponding pairs of dark and light spots that are equally distant from the neutral background (e.g., $-1/4, +1/4$; $-1/2, +1/2$; etc.). Although it is difficult to see in reduced-size figures, in nearly all cases the dark point lies to the right of the corresponding light point, indicating a greater JND separation for adjacent dark versus light contrasts, and in nearly all cases the dark point in the graph lies above the corresponding white point, indicating greater attention selectivity for dark versus light. Therefore, in Figs. 4C and 5, within measurement error, both the dark and spots lie on the same diagonal line. That means the attention selectivity advantage (greater height in Figs. 4C and 5) of dark versus light spots is entirely due to their greater JND separation, not due to different parameters for attention to dark versus attention to light. Examining light and dark spots separately shows again that attention selectivity for light–dark contrasts is greatly dependent on their JND separation and is much worse than attention selectivity for hue, which is better and virtually independent of the target–distractor JND separation for precisely the same separations (Fig. 5).

A Threshold for Attention Filtering. An interesting aspect of the light–dark data is that the x intercept of the log selectivity versus JND isolation graphs for BI–Wh set of spots is greater than zero. A greater-than-zero x axis intercept means that, below a threshold number of JNDs (the intercept), the attention system is inoperative. In Fig. 4C, the attention threshold seems to be lower for the R–G (about five JNDs) and Y–B (about three JNDs) stimulus sets than for the BI–Wh stimulus set (attention threshold about seven JNDs).

FBA Acts on Location. Rather than directly assessing the perception of features, the measured effect of attention to color in the centroid paradigm as well as attention to color in motion paradigms is on the locations of the features. That is, both centroid and motion computations depend on location computations; features enter into these computations only indirectly insofar as they change the weights assigned to different locations. Similarly, Shih and Sperling (23) in a rapid stream of visual displays search experiment concluded that “selective attention to a particular size or color does not cause perceptual exclusion or admission of items containing that feature; it acts by guiding search processes to spatial locations that contain the to-be-attended features (ref. 23, p. 758).” Prioritizing locations is, in fact, the guiding principle in contemporary computational theories of visual search (e.g., refs. 24–26). To conclude: Although centroid, motion, and search tasks are used to measure FBA, counter to intuition, these tasks do not measure a direct effect of FBA on the perception of the attended features; they measure an indirect action of FBA in selectively weighting the locations of features.

Nonlocation Measurements of FBA. In paradigms other than centroid, ambiguous motion, and search, the attended feature does not necessarily involve location. Among the paradigms that measure effects of FBA other than its effect on location, the most prominent is adaptation of the attended feature, most commonly adaptation to a particular motion direction as revealed by the motion after-effect (e.g., ref. 27). Adaptation to grating orientation also varies depending on FBA (28). A useful demonstration of FBA occurs in repetition detection of digits that share color, spatial frequency, orientation, size, or combinations of these features (29). The extent to which and how the as-yet-unmeasured, implicit attention filters

activated by these many FBA paradigms relate to attention filters derived from centroid judgments is a critical unanswered question.

Materials and Methods

Apparatus. The experiment was conducted on an iMac intel computer running MATLAB with a Psychtoolbox package (30). The built-in 23-in., 60-Hz fresh rate, LED monitor with $1,920 \times 1,080$ resolution was used to display the stimuli. The luminance lookup table contained 256 gray levels generated by a standard calibration procedure. The mean luminance of the monitor was 52.1 cd/m^2 . Stimuli were viewed at a fixed distance of 60 cm.

Generating Test Colors. The eight equiluminant hue colors were derived using a minimum motion paradigm (31–33). Subjects adjusted the lightness of a color so that the contrast between the color and the monitor’s mean luminance produced no motion, only flicker. That is, all colors were made equiluminant to the mean luminance of the monitor according to the luminance sensitivity of motion system. The end points of the R–G, Y–B, and BI–Wh axes were the largest available contrasts available on the monitor. In addition to being in the equiluminant plane (as determined for each participant individually), the gamut of lights along the R–G axis is chosen to produce invariant activation of the S cones [as defined by the Stockman–Sharpe (34), 2° fundamental], and the Y–B axis is orthogonal (in cone-activation space) to the R–G axis. These constraints exhaust the three degrees of freedom of the color display monitor. The coordinates of the monitor screen were specified in the cone space. The eight BI–Wh contrasts were $-1, -0.75, -0.5, -0.25, 0.25, 0.5, 0.75,$ and 1 , where -1 is black and $+1$ is the most intense white, and the monitor background is 0 . The other half-axes R, G, Y, and B also were equally divided in physical cone space. The eight hues were chosen to lie in a circle in the 3D cone space that fell on an equi-luminant plane (29).

The Centroid Extraction Paradigm.

Subjects. Four naive subjects (S1–S4) were undergraduate psychology students who were unaware of the purpose of the study. The other subject, S5, was P.S. All subjects gave informed consent to participate in the study. All methods used were approved by the University of California, Irvine institutional review board.

Stimuli. The stimulus display was 512 by 512 pixels (visual angle 12.1°) centered in a display of $1,920 \times 1,080$ pixels. Dots were nine-pixel-wide squares (visual angle 0.21°). In any one condition, only one of the four sets of eight colors was tested. Every stimulus contained three dots of each of the eight colors. Target and distractor dot locations were drawn from two bivariate Gaussian distributions with different means but the same SDs of 3° . The means of the two dot distributions roved independently, both with SD of 0.7° . If a sampled dot location occurred within 10 pixels of another dot the location was resampled so that two dots never overlapped. The roving variation improves the statistical power of the paradigm and is imperceptible.

Procedure. Before the main experiments, subjects were trained to judge the centroid of 1, 3, 6, 12, or 24 dots, respectively. After 50 training trials on one dot and about 200 training trials on each of the other dot numbers, all with feedback, subjects were able to locate the centroid very accurately and were no longer improving. Each of the four conditions (hue, BI–Wh, RG, and BY) was tested in a blocked design. The target color was fixed within a block. Each block contained 100 test trials interleaved with 20 control trials in which only targets were shown. With the identification number given to each color in Fig. 3C colors were tested once in the order 7, 8, 2, 3, 1, 5, 4, 6 and then once again in the reverse order—an ABBA design.

Estimating Filter Weights and Filter Efficiencies. Let C_i be the color of the i th dot in a display and $g(C_i)$ be the weight assigned to it [implying $g(C_i)$ is dependent on C_i only]. The model estimates the mean location of the 24-dot array (its centroid) with each dot weighted by its corresponding $g(C_i)$. Then the model’s x, y response R_x, R_y is given by

$$R_x = \frac{\sum_{i=1}^{i=24} g(C_i) x_i}{\sum_{i=1}^{i=24} g(C_i)} + Q_x \quad \text{and} \quad R_y = \frac{\sum_{i=1}^{i=24} g(C_i) y_i}{\sum_{i=1}^{i=24} g(C_i)} + Q_y, \quad [1]$$

where Q_x and Q_y represent noise terms. Let X_k, Y_k be the sum of the locations of all three dots of color C_k ; then, Eq. 1 can be rewritten as

$$R_x = \sum_{k=1}^{k=8} f(C_k) X_k + Q_x \quad \text{and} \quad R_y = \sum_{k=1}^{k=8} f(C_k) Y_k + Q_y, \quad [2]$$

where $f(C_k)$ is equal to $g(C_k)/D$ for D the denominator in each of the two parts of Eq. 1. Eq. 2 describes a simple linear model in which $f(C_k)$ can be estimated by multiple linear regression.

Finding the Number of JNDs Between Target Colors. The JND is a measure of the color discriminability between nearby colors (here, colors lying on one of the three color axes or on the hue circle) (Fig. 2). To measure a single JND, two dots that were configured exactly like the stimulus dots in the attention experiments were shown for the same duration as in the attention experiments (300 ms) separated by 0.93° centered on the screen. One dot is the standard color around which the JND is measured, and the other is a variable color. The standard dot had a fixed color on a sequence of trials, and the variable dot took one of 31 (or 62) nearby colors on the color axis of that set. A staircase procedure was used to find the two-dot color discrimination threshold (JND) at each of the eight colors in a color set. JNDs were measured for the four sets of eight standard testing colors. Note that except for the end points there are two JNDs at each standard color, one in each direction on the axis or hue circle. JNDs were measured separately for each color set.

For each adjacent color pair, the continuum of the 33 colors (two adjacent colors and 31 colors between them) were shown to the subject before the staircase procedure began. In the test phase, subjects were to indicate which of the two dots in a pair was closer in color to a designated end of the color continuum. For example, to estimate the JND threshold for the color orange toward the color yellow, subjects were asked to indicate which dot was more yellow. There was feedback after each trial.

For each of the three diagonal axes, one additional JND threshold was measured around the background gray. In this case, one of the two dots had a color drawn from between one of the unsaturated axis colors (e.g., unsaturated red) and the background gray, and the other dot had a color drawn from between the other unsaturated color (e.g., unsaturated green) and the background gray. The subjects indicated which of the dots was closer to one of the unsaturated color (e.g., which one was more reddish or greenish). Two three-down, one-up staircases were interleaved within a block to measure the JND thresholds for a color difference. A staircase terminated after 16 reversals had occurred. The initial 10 trials in each block were discarded. The staircase data were fit with a Weibull function; the centering parameter was used as the threshold (i.e., 82% correct threshold).

The number of JNDs (i.e., perceptual distance) between two colors was estimated using the method described below. Other ways of calculating JND isolations and estimating perceptual distances were also tried but did not alter the conclusions.

Definition of the Stimulus Space. The stimuli in the present experiments are described in dimensions that are defined in terms of the human cone space, that is, in terms of the action spectra of the blue-, green-, and red-sensitive retinal cones. Or, as they are called by vision scientists, the short- (S), medium- (M), and long- (L) wavelength cones. To describe stimuli in terms of cone space, it was first necessary to measure the spectral characteristics of the monitor that presents the stimuli; this was done by standard methods.

We define the cone space in terms of the Stockman–Sharpe (34) 2° short-, medium-, and long-wavelength cone fundamentals. Each of S , M , and L is a function that maps the wavelengths 390, 395, 400, ..., 830 nm into \mathbb{R} describing a cone's relative wavelength sensitivity. We represent S , M , and L as column vectors of length 89. Then, let B be the three-column matrix in which the first column B_1 is equal to S rescaled to have norm 1, the second column B_2 is equal to $M + L$ orthogonalized relative to B_1 and rescaled to have norm 1, and the third column is $L - M$ orthogonalized relative to B_1 and B_2 and rescaled to have norm 1. B describes the stimulus space, three orthogonal dimensions (B - Y , R - G , and luminance) defined by cone sensitivities. (The S and $L - M$ dimensions are illustrated in Fig. 2A.)

A stimulus of color C with a wavelength spectrum Q (a column vector of length 89) is represented in the stimulus space by the 3D vector C_Q :

$$B^T Q = \begin{bmatrix} q_1 \\ q_2 \\ q_3 \end{bmatrix} = C_Q, \quad [3]$$

which is referred to informally in the text as a "color." The percept produced by Q depends entirely on C_Q . Any lights Q_1 and Q_2 for which $C_{Q_1} = C_{Q_2}$ will appear identical to human vision (such lights are called metamers) and the perceptual discriminability two colors C_{Q_1} and C_{Q_2} is a smoothly increasing function of the Euclidean distance between them; this will be important in measuring JNDs.

Calculating the Number of JNDs Separating Two Stimuli. To calculate the number of JNDs between two adjacent stimulus colors with spectra Q_1 and Q_2 , we measure one JND at each end of the line segment in \mathbb{R}^3 between C_{Q_1} and C_{Q_2} and then, as explained below, infer how many JNDs must be between them to fill distance covered by that line segment. Let a be the threshold Euclidean distance from C_{Q_1} toward C_{Q_2} ; let c be the threshold from C_{Q_2} toward C_{Q_1} ; and let $D = a + b + c$ be the total Euclidean distance between C_{Q_1} and C_{Q_2} . In all cases in the current study, $b > 0$, and $a \neq c$. We anticipate that human vision may be more sensitive to variations in some regions/directions of color space than in other regions/directions. However, JNDs were measured only between adjacent colors, so we can assume that such variations in sensitivity are gradual enough that discrimination thresholds between the pairs of adjacent colors in the current study undergo a smooth transition from a to c , with each successive JND equal to a fixed scalar α times the preceding one.

This implies that for n equal to the number of JNDs from C_{Q_1} to C_{Q_2} ,

$$c = a\alpha^{n-1}, \quad [4]$$

and in the case in which n is an integer,

$$D = a \sum_{k=0}^{n-1} \alpha^k. \quad [5]$$

Note that if b were equal to 0, then Eqs. 4 and 5 would imply that $\alpha = c/a$ and (as desired) $n = 2$.

From Eq. 4, it follows that

$$\frac{c}{a} = \alpha^{n-1}. \quad [6]$$

From Eq. 5 it follows that

$$\frac{D}{a} = \frac{\alpha^n - 1}{\alpha - 1}. \quad [7]$$

Eq. 7 generalizes Eq. 5 to noninteger values of n . Substituting Eq. 6 into Eq. 7 yields

$$\frac{D}{a} = \frac{\left(\frac{c}{a}\right)\alpha - 1}{\alpha - 1}, \quad [8]$$

implying that

$$\alpha = \frac{D - a}{D - c} = \frac{b + c}{b + a}. \quad [9]$$

Substituting Eq. 9 into Eq. 6 yields

$$\frac{c}{a} = \left(\frac{b + c}{b + a}\right)^{n-1}, \quad [10]$$

from which it follows that the number of JNDs between C_{Q_1} and C_{Q_2} is given by

$$n = 1 + \frac{\ln c - \ln a}{\ln(b + c) - \ln(b + a)}. \quad [11]$$

To reiterate: a is the Euclidean distance of the threshold from C_{Q_1} in the direction of C_{Q_2} , c is the Euclidean distance of the threshold from C_{Q_2} in the direction of C_{Q_1} , D is the Euclidean distance between C_{Q_1} and C_{Q_2} , and

$$b = D - a - c. \quad [12]$$

ACKNOWLEDGMENTS. This work was supported in part by National Science Foundation Grant BCS-0843897.

1. Serences JT, Boynton GM (2007) Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron* 55(2):301–312.
2. Zhang W, Luck SJ (2009) Feature-based attention modulates feedforward visual processing. *Nat Neurosci* 12(1):24–25.
3. Liu T, Mance I (2011) Constant spread of feature-based attention across the visual field. *Vision Res* 51(1):26–33.

4. White AL, Carrasco M (2011) Feature-based attention involuntarily and simultaneously improves visual performance across locations. *J Vision* 11(6):15.
5. Drew SA, Chubb CF, Sperling G (2010) Precise attention filters for Weber contrast derived from centroid estimation. *J Vision* 10(10):20.
6. Sun P, Chubb CF, Wright C, Sperling G (2016) The centroid paradigm: Quantifying feature-based attention in terms of attention filters. *Atten Percept Psychophys* 78(2):474–515.

7. Wang Y, Miller J, Liu T (2015) Suppression effects in feature-based attention. *J Vis* 15(5):15.
8. Störmer VS, Alvarez GA (2014) Feature-based attention elicits surround suppression in feature space. *Curr Biol* 24(17):1985–1988.
9. Carrasco M (2011) Visual attention: The past 25 years. *Vision Res* 51(13):1484–1525.
10. Treue S, Martinez Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399(6736):575–579.
11. Bichot NP, Rossi AF, Desimone R (2005) Parallel and serial neural mechanisms for visual search in macaque area V4. *Science* 308(5721):529–534.
12. Müller MM, et al. (2006) Feature-selective attention enhances color signals in early visual areas of the human brain. *Proc Natl Acad Sci USA* 103(38):14250–14254.
13. Andersen SK, Müller MM (2010) Behavioral performance follows the time course of neural facilitation and suppression during cued shifts of feature-selective attention. *Proc Natl Acad Sci USA* 107(31):13878–13882.
14. Blaser E, Sperling G, Lu ZL (1999) Measuring the amplification of attention. *Proc Natl Acad Sci USA* 96(20):11681–11686.
15. Scholl BJ (2001) Objects and attention: The state of the art. *Cognition* 80(1–2):1–46.
16. Halberda J, Sires SF, Feigenson L (2006) Multiple spatially overlapping sets can be enumerated in parallel. *Psychol Sci* 17(7):572–576.
17. Ariely D (2001) Seeing sets: representation by statistical properties. *Psychol Sci* 12(2):157–162.
18. Inverso M, Sun P, Chubb C, Wright CE, Sperling G (2016) Evidence against global attention filters selective for absolute bar-orientation in human vision. *Atten Percept Psychophys* 78(1):293–308.
19. Pylyshyn ZW, Storm RW (1988) Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spat Vis* 3(3):179–197.
20. Cavanagh P, Alvarez GA (2005) Tracking multiple targets with multifocal attention. *Trends Cogn Sci* 9(7):349–354.
21. Pylyshyn ZW (2004) Some puzzling findings in multiple object tracking: I. Tracking without keeping track of object identities. *Vision Cogn* 11(7):801–822.
22. Lu Z-L, Sperling G (2012) Black-White asymmetry in visual perception. *J Vision* 12(10):8.
23. Shih SI, Sperling G (2002) Measuring and modeling the trajectory of visual spatial attention. *Psychol Rev* 109(2):260–305.
24. Wolfe JM (1994) Guided Search 2.0 A revised model of visual search. *Psychon Bull Rev* 1(2):202–238.
25. Elazary L, Itti L (2010) A Bayesian model for efficient visual search and recognition. *Vision Res* 50(14):1338–1352.
26. Geisler WS, Cormack LK (2011) Models of overt attention. *Oxford Handbook of Eye Movements*, eds Livingsedge SP, Gilchrist ID, Everling S (Oxford Univ Press, New York), pp 439–454.
27. Boynton GM, Ciaramitaro VM, Arman AC (2006) Effects of feature-based attention on the motion aftereffect at remote locations. *Vision Res* 46(18):2968–2976.
28. Liu T, Larsson J, Carrasco M (2007) Feature-based attention modulates orientation-selective responses in human visual cortex. *Neuron* 55(2):313–323.
29. Sperling G, Wurst SA, Lu Z-L (1992) Using repetition detection to define and localize the processes of selective attention. *Attention and Performance XIV*, eds Meyer DE, Kornblum S (MIT Press, Cambridge, MA), pp 265–298.
30. Brainard DH (1997) The Psychophysics toolbox. *Spat Vis* 10(4):433–436.
31. Anstis S, Cavanagh P (1983) A minimum motion technique for judging equiluminance. *Color Vision: Psychophysics and Physiology*, eds Mollon JD, Sharpe LT (Academic, London), pp 155–166.
32. Lu Z-L, Sperling G (2001) Sensitive calibration and measurement procedures based on the amplification principle in motion perception. *Vision Res* 41(18):2355–2374.
33. Herrera C, et al. (2013) How do the S-, M- and L-cones contribute to motion luminance assessed using minimum motion? *J Vision* 9:1021.
34. Stockman A, Sharpe LT (2000) The spectral sensitivities of the middle- and long-wavelength-sensitive cones derived from measurements in observers of known genotype. *Vision Res* 40(13):1711–1737.