Available online at www.sciencedirect.com
SCIENCE DIRECT®
Vision Research xxx (2005) xxx–xxx
Vision Research
www.elsevier.com/locate/visres

Quantifying object salience by equating distractor effects

Liqiang Huang, Harold Pashler *

Department of Psychology, University of California, San Diego, 9500 Gillman Dr., La Jolla, CA 92093, United States

Received 6 October 2003; received in revised form 22 October 2004

Abstract

It is commonly believed that objects viewed in certain contexts may be more or less salient. Measurements of salience have usually relied on asking observers "How much does this object stand out against the background?". In this study, we measured the salience of objects by assessing the distraction they produce for subjects searching for a different, pre-specified target. Distraction was measured through response times, but changes in response times were not assumed to be a linear measure of distracting potency. The analysis rested on measuring the effects of varying disparities—in size, luminance, and both—between a target object, a key distractor, and other background items. Our results indicate: (1) object salience defined by luminance or size disparity is determined by the ratio between its defining feature value and the corresponding feature value of background items; this finding is congenial to Weber's law for discrimination thresholds. (2) If we define salience as the logarithm of a feature value ratio, then salience increases approximately as fast due to increase in area as due to increase in luminance. (3) The sum of salience arising from object-background disparity in both size and luminance is larger than their vector sum (orthogonal vectors), but smaller than their scalar sum.

1. Introduction

1.1. Background: salience and search

The term salience refers to how much an object stands out from the scene in which it appears. It is a basic psychological concept referred to by early writers on attention (Titchener, 1908); attention, it has long been noted, is relatively likely to be drawn to relatively salient objects. Salience has also become a central concept in the computational analysis of vision (for a recent model, see Parkhurst, Law, & Niebur, 2002), with applications to problems such as the development of efficient image-compression schemes. In the study of visual search, the guided search model (Wolfe, 1994) posits that the brain computes salience within each perceptual dimension and sums these salience signals into an "activation map", with total salience determining sampling rate. In Lu and Sperling's (1995) proposed architecture for motion perception, salience is jointly determined by bottom-up and top-down factors, with what they term "third-order motion" being computed within a salience map. Some people hypothesize that salience, when represented in a binary form, mediates figure-ground segregation (Lu & Sperling, 1995).

Given the importance of salience to diverse aspects of vision science and its applications, good measures of salience are clearly desirable. Some standard psychophysical discrimination measures provide a means of quantifying salience, but only within a certain range. To consider another example, in a visual search task where the target differs from other objects in one dimension (so-called "singleton search"), if the target–distractor difference is small, the time taken by observers to find the target increases with the number of items. The increasing response time for each extra item ("search slope") presumably reflects the salience of the target against the distractors. However, when the difference be-

* Corresponding author.
E-mail addresses: liqiang@psy.ucsd.edu (L. Huang), hpashler@ucsd.edu (H. Pashler).

0042-6989/$ - see front matter © 2005 Published by Elsevier Ltd.
doi:10.1016/j.visres.2005.01.013

tween the target and background items is reasonably large, the slope is essentially zero (often termed "pop-out"; for a recent review, see Wolfe, 1998). Thus, while search slope can provide an objective measure of singleton salience, once the salience exceeds a certain point, the method cannot be used.

The measurement of suprathreshold salience differences has been examined in various ways, most notably by Nothdurft (1993a, 1993b, 2000). In one of his studies, Northdurft briefly presented two unique items against an identical array of distractors (i.e., a homogeneous background). Subjects were required either to judge which of the two was more salient, or to modify a feature value to make them equally salient. The results show that a more salient target can be found more efficiently (Nothdurft, 1993b) and that increasing feature contrast in an additional dimension makes a singleton more salient (Nothdurft, 2000). He also showed how subjectively assessed salience varied across several different dimensions.

1.2. Basic measurement and basic rationale

The present work takes a slightly different approach to the problem of measuring suprathreshold salience. We did not ask our subjects to make judgments about salience, nor did we use the word "salience" in their instructions. Rather, we sought to base our measure of salience on the behavioral effects of salience differences, namely the tendency of salient objects to draw attention even when they are nominally task-irrelevant (a tendency that seems to be implicit in the concept of salience as it has been discussed over the years).

The method for measuring salience that will be offered here is based on a visual search task. In two-thirds of the trials, we presented two singletons against a homogeneous background (with a total of 20 elements in each display). One of the singletons (the *target*) was both brighter and larger than the background items; it was this target that the subjects searched for. The other singleton (termed the *key distractor*) was brighter than the background items *or* larger than the background items, but not both.[1] A key distractor brighter than the background items will be referred to as a *luminance key distractor* (LKD). A key distractor larger than the background items will be referred to as a *size key distractor* (SKD) (see Fig. 1 for two sample displays). One-third of the trials had a target and a size key distractor (T + SKD); one-third had a target and a luminance key distractor (T + LKD); the remaining one-third had one target and no key distractor (T). The 18

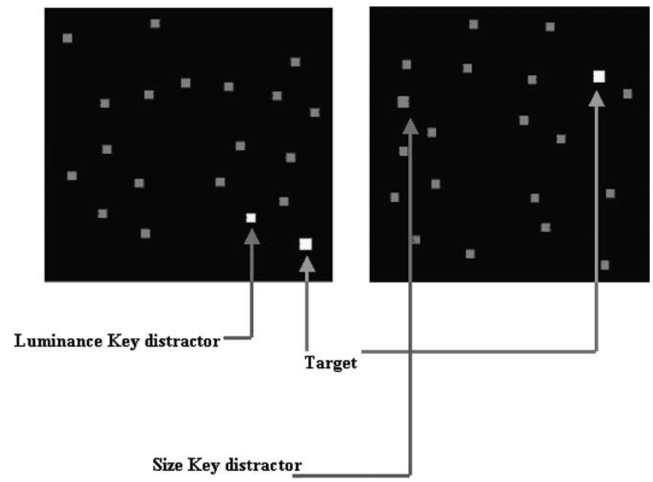

Fig. 1. Examples of displays used in Experiment 1. We presented two singletons against an array of identical items. One of the singletons (the *target*) was both brighter and larger than the background items; it was this target that the subjects searched for. The other (termed the *key distractor*) was brighter than the background items *or* larger than the background items, but not both. A key distractor brighter than the background items we refer to as a *luminance key distractor.* A key distractor larger than the background items we refer to as a *size key distractor*. We calculated the salience of key distractors by how disruptive they were to the task of finding the target.

or 19 identical items comprising the rest of the display will be referred to as *background distractors*.

The method that will be described here derives conclusions about salience from measurements of the distracting effects of size and luminance key distractors. Distraction is measured by comparing average response times (RTs) between trials in which the key distractors are present, and trials with only background distractors [2]. Thus, the situation described yields two distraction measures: RT(T + SKD) − RT(T), and RT(T + LKD) − RT(T).

We will assume here that a more salient key distractor will produce a greater distraction effect. This could occur for a variety of reasons. For example, subjects may sometimes detect a key distractor before the target, and more salient targets would be misdetected more frequently. Alternatively, more salient distractors might "hold" attention for a longer period of time (even if the frequency of misdetection is no higher), thus increasing RTs. Third, when the key distractors are not very salient, the observer may rely mainly on bottom-up sig-

[1] In Experiment 4, investigating the combination of salience from different dimensions, the key distractors did sometimes differ from the background distractors in both dimensions (composite key distractors).

[2] One might wonder about the role of similarity in these experiments. Salience is determined by what we define as the feature difference and the task relevance of a dimension; similarity affects a different (later) processing. We tried to make the key distractor always very different from the target in order to keep the role of similarity at a minimum. Results of Experiment 3 confirm that a less similar but more salient key distractor is indeed more distractive than one more similar and less salient. (The distractor most similar to the target evinces no unique advantage in the distraction effect curve.)

nals, and when the key distractors become more and more salient, observers have to rely on more and more top-down (and effortful) strategies and the response will therefore be slowed down.

We will not assume that the magnitude of the distraction (i.e., the two differences referred to above) offer any kind of direct measure, e.g., a linear function, of the underlying salience of the distractor. Given the complexity of mechanisms that might mediate the distracting effect, as noted above, any such assumption would obviously be unwarranted. Furthermore, in general, one cannot assume that the salience of an object is solely dependent on the physical differences between the element and other surrounding elements. The "top-down" weighting of different dimensions (reflecting the specific task instructions that the observer has been given) might also contribute to salience. One might think that if such assumptions cannot be made, then measuring the distraction produced by introducing the two different types of distractors (LKD and SKD) cannot be of much help in measuring salience.

Nonetheless, in an appropriate experimental design, it is possible to make some fairly strong inference about salience based on distraction differences. To see the logic that will be employed here, the reader should refer to Fig. 2. This figure illustrates the featural values of the four elements in the situation just described (SKD, LKD, T, and background distractors). The four points represent a square in 2-dimensional feature space. The basic strategy used to derive conclusions about salience differences is as follows. The feature values of luminance are held constant throughout the experiment. The values of size, on the other hand, are adjusted from one block of trials to the next. (Fig. 3 shows hypothetical feature values for the four element types in three possible different blocks of trials. Note that the two bottom points move up or down together between blocks, i.e., the size of the normal and luminance key distractor is always the same.)

In the nutshell, our logic relies on the fact that when the experiment is arranged as described in the preceding sentence (where the features of the normal distractor are yoked to the features of the two key distractors), then

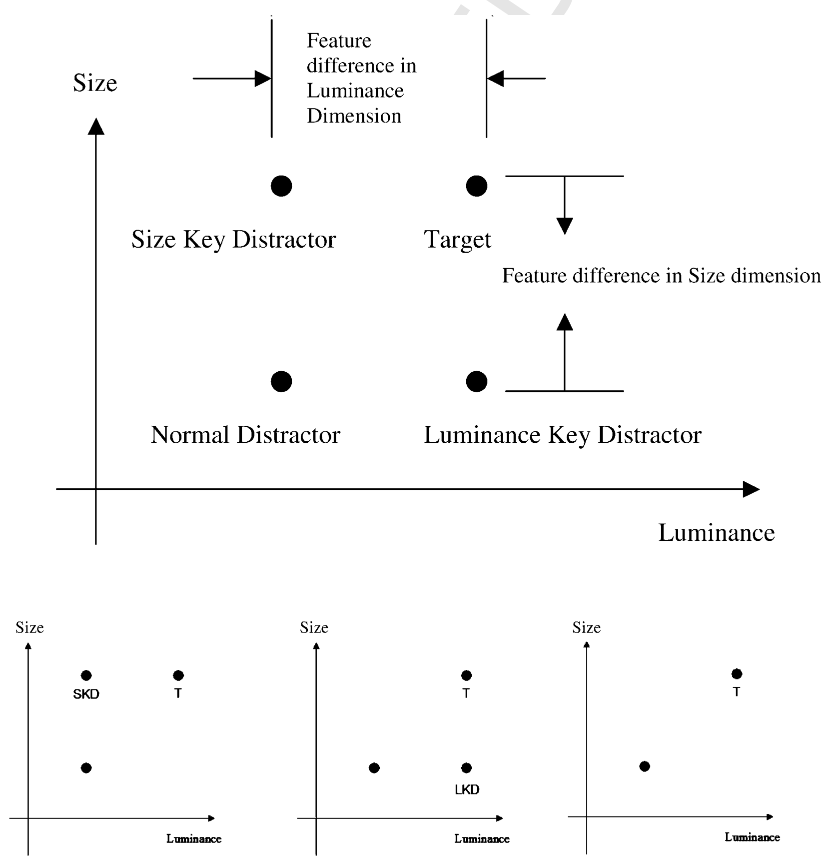

Fig. 2. (Top) In Experiment 1, we made the value of the defining feature (size or luminance) of each key distractor the same as the corresponding value of the target. In this way, the feature difference determining the task relevances of the two dimensions (the difference between the background distractors and the target), and the feature difference determining the salience of the key distractors (the difference between the background distractors and the key distractors) were equal, therefore, the distraction effect as we defined above was directly related to the feature difference values of the two dimensions. (Bottom) Three types of trials: SKD + T, LKD + T, T.

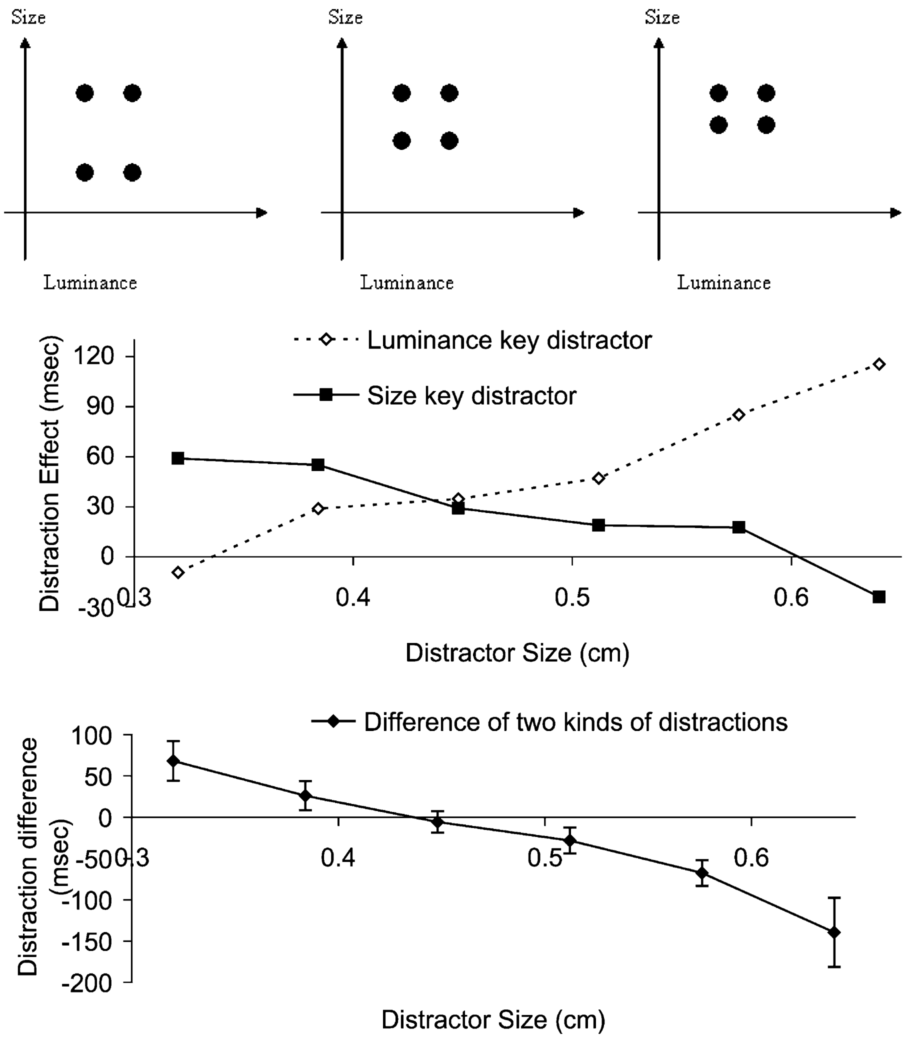

Fig. 3. Methods and results of Experiment 1: (top) In Experiment 1, for different blocks, however, when the size of background distractors increased, the size feature difference decreased and the luminance feature difference kept constant; (middle) distraction effects of two kinds of trials: luminance key distractor, size key distractor and (bottom) the difference between distraction effects in the two dimensions of luminance and size. The size distractor is most distractive at the left (with small background distractors) and decreases gradually to the right. The luminance distractor's effect increases gradually from left to right and is finally much greater than the size distractor's effect. The Error bars represent 95% confidence interval of the difference between distraction effects of size and luminance key distractors.

the following conclusion must also hold: *the salience of the disparity between the background distractors and the key distractors in one dimension equals the salience of the disparity between the background and key distractors in the other dimension, if and only if average response times are the same for SKD and LKD trials*. The reader may question whether the italicized sentence can really hold in any general sense, given the fact that latencies appear to be a fairly arbitrary measure of processing difficulty, and given the heterogeneity of mechanisms whereby salient distractors might increase latencies (as noted above). Fortunately, the statement can be derived from quite weak assumptions, whose reasonableness seems inherent in most prior discussions of salience and search (as well as being consistent with various results in the literature).

1.2.1. Assumptions

The assumptions needed to warrant the italicized statement above are presented informally in the text, and more formally in Appendix A. First, we assume that whether any two objects differ on only one dimension, or on more than one, the perceived difference between the objects can always be represented with a single scalar value. Take, for example, the difference between two orientations, vertical and horizontal, and the difference between two colors, red and green. While there is no meaningful way to compare these two physical differences, since they involve different dimensions, we assume that the sense of disparity they produce in an observer can be represented as a point on a common scale. Thus, the perceived overall difference between red and green is larger than the perceived overall difference between two identical lines, one tilted 45° and one tilted 46°. We will assume very little about the unknown feature difference function: merely that associated with each pair of objects in a display is a subjective feature difference value, and that this value is a monotonically increasing function of feature disparities within each dimension.

Second, we assume that salience of an element in a display basically depends on the feature difference values that characterize each of the differences between the element and the others in the display. However, as noted earlier, there are also a variety of top-down factors that may affect salience. For the purpose of assessing salience, we will only consider the relative importance of two different feature dimensions: specifically, luminance and size. These dimensions need not be weighted identically. If one dimension is weighted more highly, then this dimension will contribute more to salience. This top-down factor will be designated the "task relevance of a dimension". The impact of this task relevance is widely assumed in models of visual search (e.g., Wolfe, 1994) and empirical evidence for its effects has been reported (Mueller, Heller, & Ziegler, 1995).

Third, we assume that when the subject searches for a target that differs on two dimensions from the distractors, then the task relevance for each dimension is a monotonically increasing function of the feature difference for that dimension. That is, the more useful the input from a given dimension, the greater the weighting that dimension will be assigned. For example, when an observer performs a task that depends upon luminance information, luminance information in the stimuli will be assigned a higher task relevance than, say, information about orientation.

To return to the visual search task that is the focus of the present paper, consider what happens as we vary the physical magnitude of the size differences (Fig. 3). When the size difference increases, so do the corresponding feature difference values. There will be some value of the physical size difference where the feature difference will equal the feature difference associated with the fixed luminance difference. When this is the case, the size difference and the luminance difference will be equally useful in the subject's task of discriminating targets from background distractors (and also for discriminating targets from either of the possible key distractors). Thus, from the final assumption mentioned above, we can infer that the weighting for the two dimensions should be the same. From this, and the very weak assumptions that we make about a monotonic relationship between distractor salience and observed distraction cost, we can infer that for this size difference (and only for this size difference), the two observed differences $RT(T + SKD) - RT(T)$, and $RT(T + LKD) - RT(T)$ would be equal. In addition, conversely, one can infer that when the RT differences for the two key distractors are the same, the underlying salience must be the same.

A further inference, namely that this size value will be the *only* point where we will observe equal RT differences, can be justified as follows. We have assumed that, in the task of searching for the target singleton, the task relevance of any feature dimension is decided by the feature difference between the target and the background distractors (distractors of the homogeneous background).[3] In our case, if one of the dimensions in question—luminance or size—had a greater feature difference (as for target/ background distractor difference) than the other, then it would have a larger task relevance. This rationale is intuitive: there is no reason to spend more resource collecting information from one dimension if richer information is available from the other dimension.

It should be kept in mind that all the conclusions described here apply only because when the size is changed between blocks, both the key luminance distractor and the background distractor change in size together. Otherwise, the task relevance for size could well be larger than the task relevance for luminance even when the feature difference of luminance key distractor is greater than size key distractor. (That will happen if we only make the feature difference between target and background distractor arise mostly from differences in the size dimension.) Graphically, this means that in Figs. 2 and 3, the fours dots are constrained to lie at the corners of a rectangle.

Our goal in Experiment 1 was to find two pairs of feature values, one in luminance and one in size, with equal feature differences. These feature values were then used to calibrate the other experiments to make sure that the task relevances of the two dimensions were equal.

Let us summarize the points introduced so far

1. The salience of a key distractor is reflected in its distraction effect.
2. The salience of a key distractor is jointly decided by the feature difference between it and background distractors and by the task relevance of the dimension in which it differs from background distractors.

[3] One might think that not only the target, but also the key distractors will affect the task relevances. To us it seems reasonable that the underlying search mechanism will be optimized to distinguish between the target and the majority of distractors. The key distractor should therefore play very little role. In addition, in all experiments we tried to make the presence of the key distractors of the two different dimensions equal and unpredictable, so that, even if the presence of key distractors does have some effect on the task relevance, it should not have caused any systematic bias in the measurement.

3. The task relevances of size and luminance is increased when the feature differences (in size and luminance dimensions, respectively) between the target and background distractors are increased.

These three points seem intuitively reasonable and not overly restrictive, as well as being generally consistent with current theories and models of visual search. Points two and three are further supported by studies on priming (Mueller et al., 1995). As will be seen below, they are also supported by the results of our experiments.

1.3. Outline of the study

Throughout Experiment 1, the target's luminance and size both remained constant. Background distractors' luminance also remained constant, but their size changed from block to block. Two key distractors in each block were each defined with one feature of the target and one feature of background distractors (see Fig. 2). When the size of background distractors changed, the feature difference in the size dimension changed; the two bottom dots shown in Fig. 3 moved together up or down.

According to the rationale discussed above (and detailed in Appendix A, the distraction effect of a certain key distractor reflects its salience, and also the task relevance of its defining dimension. The task relevances were constant throughout any one block, since the target and background distractors' properties were constant. For different blocks, however, when the size of background distractors increased, the size feature difference decreased and the luminance feature difference kept constant (see Fig. 3). Therefore, the size task relevance decreased and the luminance task relevance increased. Thus, we could expect the salience (and distraction effect) of the size key distractor to decrease and the salience (and distraction effect) of the luminance key distractor to increase. Such was our prediction for the data pattern, and as we will record below, the prediction was exactly confirmed.

The significant yield from Experiment 1 is a pair of size values and a pair of luminance values with equal feature differences. Applying these values makes the luminance key distractor and size key distractor equally distractive. At this point the task relevances of the two dimensions are equal.

The well-known Weber's Law (Weber, 1834) states that in some perceptual stages, for a person to distinguish between two feature values, the specific values themselves are not important, but only the ratio between them. In Experiment 2 we tested the applicability of this law to our Experiment 1 results. That is to say, we tested whether the feature difference between two specific feature values (in other words, the function relating salience to the difference between two feature values) can be simplified into a function of only the values' ratio. We found that the law does basically apply: increasing or decreasing together the target and distractors' feature values while keeping the ratio between them the same had no significant effect on the feature difference and salience.

Finding in Experiment 2 that the key determinant of salience is not specific features values themselves, but rather the ratios between the paired values, we further speculated that the salience (or feature difference) in the two feature dimensions under study (luminance and size) might increase according to the same function, though possibly at different rates. Using a constant to compensate for the difference of rate, the ratios might be directly linked to feature difference and salience with a single function across the two dimensions. In Experiment 3, we varied the feature difference of the key distractors, measuring the distraction effect for several different values in both luminance and size dimensions. We found that the luminance and size curves basically match each other.

Experiment 4 addressed the question of how salience in a single object is combined from more than one dimension of feature difference. We found that when we add a small amount of salience from one dimension to salience of the other, it does increase the overall salience, but at a discount—that is, only by a certain portion of the added amount. Nothdurft (2000), pursuing the same question by a different method, has reached a similar conclusion.

2. General method

2.1. Subjects

Subjects were from the University of California, San Diego. All had normal or corrected-to-normal vision. There were 18 subjects in Experiment 1, 9 subjects in Experiment 2a, 9 subjects in Experiment 2b, 14 subjects in Experiment 3, and 24 subjects in Experiment 4.

2.2. Apparatus

Stimuli were presented on a high-resolution MAG DX-15T color monitor. Responses were recorded from two adjacent keys on a standard keyboard. The subjects viewed the displays from a distance of about 60 cm.

2.3. Stimuli

Two example displays are shown in Fig. 1. The subjects searched for one target (0.768 cm × 0.768 cm, 17.9 cd/m^2) among 20 items. In Experiments 3–4, background distractors were squares measuring

0.448 cm × 0.448 cm with luminance 5.31 cd/m^2. In 1/3 of the trials, the target was presented among 19 background distractors (making 20 items in all). In the other 2/3 of the trials, there were 18 background distractors and 1 key distractor. The properties of key distractors and background distractors are given below for each experiment. All the items were randomly located in a 19 cm × 19 cm region. The background was black (<0.2 cd/m^2). There was one small red dot on the left edge or right edge of each item. The location of the red dot on the target (left edge or right edge) decided the response key.

2.4. Procedure

Each trial began with a small green fixation cross presented in the center of the screen. The subject was instructed to fixate the cross, which remained present for 400 ms. The cross was followed by a short blank interval (400 ms). That was followed by the display, which remained until the subject responded. In all the experiments, once the display appeared, the subject found the target, decided whether the target's red dot was on the left or right edge of the target, and responded by pressing one of two adjacent keys ('j': left side; 'k': right side) with two fingers of the right hand. This forced-choice discrimination task is used instead of having observers report the presence or absence of the target in order to reduce the variability involved in the "yes/no" decision. Subjects were not instructed to keep fixation on the center through the whole trial. Eye movement after display presence is in fact very common. Subjects were told to respond as rapidly and accurately as possible. A positive or negative sound was played to provide feedback on the accuracy of each response. Each subject performed 13 blocks of 80 trials each, with the first block excluded as practice. Different block conditions, when they existed in one experiment, were counterbalanced across subjects.

In all experiments, RTs greater than 5000 ms or smaller than 100 ms were excluded from the above RT analysis (and in all experiments of this study). Trials excluded were less than 1%.

3. Experiment 1

The purpose of Experiment 1 was to equalize the task relevances of the two feature dimensions. There were 6 block conditions, each with one size of background distractor: 0.32 cm, 0.384 cm, 0.448 cm, 0.512 cm, 0.576 cm, 0.64 cm. Background distractors always had luminance 5.31 cd/m^2. The target always measured 0.768 cm square and had luminance 17.9 cd/m^2. In each block, there were two kinds of key distractor: the size distractor had the size of the target but the luminance of background distractors; the luminance distractor had the luminance of the target but the size of background distractors. Each of these appeared in 1/3 of the trials, and in the remaining 1/3 there were only background distractors.

The mean RT of Experiment 1 was 754 ms. The results of Experiment 1 are given in Fig. 3. The distraction effect of size and luminance key distractors are given in the top panel. The size key distractor was very distractive when the background distractors were small; it became less so as background distractor size increased. The luminance key distractor followed the opposite trend. This result fits our prediction given above. The bottom panel shows the difference between the distraction effects of the two types of key distractor. This difference is positive for trials with small background distractors, gradually decreasing, and becoming negative for trials with large background distractors.

It should be mentioned that the results of Experiment 1 (and all the following experiments) include trials with or without gaze shifts, since eye movements were not controlled. This undoubtedly introduces some noise into the situation, but we believe it will not systematically bias our result since all assumptions made for the rationale is true whether the gaze shifts occurs or not.

The error rate in this experiment (as in all experiments of this study) was very low (and generally consistent with the effects manifest in RT measurements) and so we have omitted a detailed description of the error rate as unnecessary and irrelevant.

It should be mentioned that when we increased the size of background distractors, the feature difference in the luminance dimension remained constant. The observed increase of the distraction effect of the luminance key distractor therefore provides reliable evidence for the gradual shifting of task relevances that we predicted.

The point we looked for in the data was where the size and luminance key distractors were equally distractive: at that point the two feature dimensions were weighted equally. Our estimation of the background distractor size at this point is 0.44 ± 0.03 cm. The best estimation our monitor resolution allowed was 0.448 cm. In our later experiments we applied this feature setting to calibrate the task relevances of the two feature dimensions: target = 0.768 cm, 17.9 cd/m^2; background distractor = 5.31 cd/m^2, 0.448 cm.

4. Experiment 2

Experiment 1 identified two feature value pairs, one from each feature dimension, whose feature differences were equal to each other. But it told us little in general about how feature difference is computed from feature values. Weber's Law, if applied here, would predict that salience increases as the logarithm of the target–distractor ratio of feature values in each dimension. Experiment 2 tested this hypothesis.

In Experiment 2, we made all the ratios of target feature value to distractor feature value match the ratios of feature values we obtained in Experiment 1. In Experiment 2a, luminance values remained constant (target: 17.9 cd/m^2, distractor, 5.31 cd/m^2). Size values changed, through three conditions, but the target–distractor size ratio was always the same: 0.768 cm vs. 0.448 cm; 1.536 cm vs. 0.896 cm; 0.384 cm vs. 0.224 cm. In Experiment 2b, size values remained constant (target: 0.768 cm, distractor, 0.448 cm). Luminance values changed, through 3 conditions, but the target–distractor luminance ratio was always the same: 17.9 cd/m^2 vs. 5.31 cd/m^2; 26.9 cd/m^2 vs. 7.97 cd/m^2; 11.9 cd/m^2 vs. 3.54 cd/m^2.

Proceeding from the hypothesis that the same ratio of feature values creates the same degree of salience, we expected the distraction effect to be the same for both dimensions in all conditions of this experiment. The mean RT of Experiment 2a and 2b was 757 ms and 744 ms, respectively. The results are given in Fig. 4. One may find it is hard to appreciate how well, in fact, they fit our prediction. Looking at the data of Experiment 1, however, we see that the distraction effect is very sensitive to differences in feature values; when any feature value differed 20% from the optimal match point, the resulting difference in the distraction effect was about 40 ms—a difference larger than effects we observe in Experiment 2. So it is at least safe to say that, regardless of the specific feature values, the same ratios of feature values create *roughly* the same amount of salience with admittedly some slight non-linearity in this function.

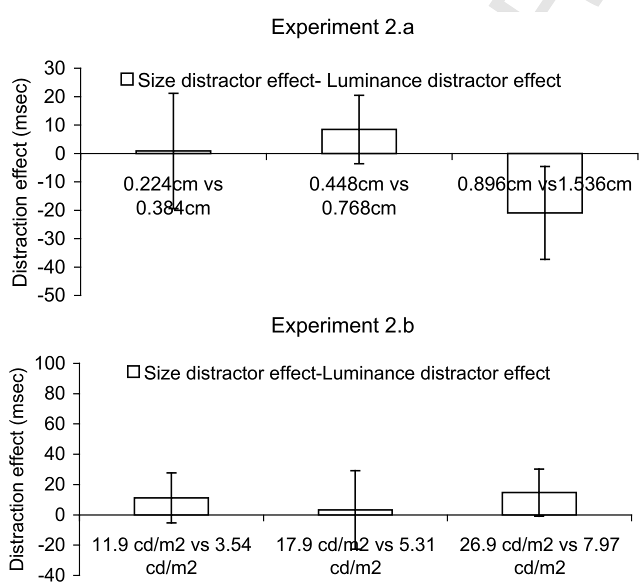

Fig. 4. Results of Experiment 2. The differences between distraction effects in the two dimensions of luminance and size: they are basically not significant in all conditions in Experiments 2a and 2b.

5. An approach to computing salience from feature values

Before stepping into further experiments, let us here propose an approach to computing salience from feature values.[4] Part of this approach has not been solidly supported, but we have tried to make it as natural as possible. The salience of a singleton against a homogeneous background of other items is determined by feature values in all dimensions. If the singleton is unique only in the dimension of size, its salience is:

$$Salience(size) = \ln(Size(Uniqueitem)/Size(backgrounditems))$$

If the singleton is unique only in the dimension of luminance, its salience is:

$$Salience(lum) \quad R_{lum} X \ln(lum(Uniqueitem)/lum(backgrounditems))$$

R_{lum} is the rate at which salience increases in the luminance dimension relative to the size dimension. (The rate of salience increase in the size dimension is defined as 1.)

It should be noted that Experiment 2 here supports this approach only up to the point of identifying the ratio between feature values as the key determinant of salience. We use the logarithmic function for several reasons: first, it makes the salience in one dimension linearly additive (e.g. the salience of a 1 cm object against 0.1 cm background objects is equal to the sum of the salience of a 1 cm object against 0.5 cm background objects and salience of a 0.5 cm object against 0.1 cm background objects); second, since this logarithm rule seems to be widely obeyed under the name of Weber's Law for near-threshold psychological measurement, it might also prove applicable here for suprathreshold measurement.

Now let us try to estimate the relative rates of increase of salience in the two feature dimensions (R_{lum}). In Experiment 1, we identified two pairs of feature values that have equal feature differences: 0.768 cm vs. 0.448 cm; 17.9 cd/m^2 vs. 5.31 cd/m^2. (These values were basically confirmed by Experiment 2.) They can be used to estimate R_{lum}.

The luminance ratio between the target and distractors is 3.37. The size ratio between the target and distractors is 1.75 ± 0.12.

$$\text{So}: \ln(1.75 \pm 0.12)/\ln(3.37) = 0.47 \pm 0.04$$

[4] Strictly speaking, the term feature difference is more appropriate than salience here, but since in this study we balanced the task relevances, feature difference and salience should be interchangeable.

If we express the “size” in terms of area rather than length, we get a relative rate of 0.94 ± 0.08. It seems that the rate at which salience increases with luminance and the rate at which it increases with area are almost the same.

6. Experiment 3

In the approach outlined above, we have suggested that the difference between the two functions governing the increase of salience can be simply compensated by a constant; that constant should be our R_{lum} value. However, if the functions are fundamentally different, R_{lum} might apply only to the single point that we derived it from. The purpose of Experiment 3 was to test whether it would also apply for other ratio values. The feature values of the target and background distractors remained constant (target: 17.9 cd/m^2 and 0.768 cm, distractor, 5.31 cd/m^2 and 0.448 cm). We tested the distraction effects of the size and luminance key distractors, each with 5 different salience levels. The feature values and predicted salience values (calculated using the equations and constant proposed above) of the 10 types of key distractors are given in Table 1.

The mean RT of Experiment 3 was 792 ms. The results are given in Fig. 5. The observed increasing distraction effect (y-axis) is plotted against key distractors' salience as calculated according to our proposed approach, as a function of key distractor–distractor ratios (compensated with the R_{lum} constant for luminance).The two curves thus produced, representing size and luminance, fit together pretty well, without any apparent systematic deviation. It seems it is indeed the case that the function relating salience increase to feature value ratios is basically the same for these two feature dimensions. This experiment confirms the value of R_{lum} and illustrates its general applicability.

Experiment 3 has a further implication. When the luminance key distractor was brighter than the target, its distraction effect still increased with luminance, even as it became less and less similar to the target. Similarly, the size key distractor's effect continued to increase with size even after its size had surpassed that of the target. This data represents a strong argument against any visual search theory holding that the underlying search mechanism only computes similarity. It supports a model holding that relative salience determines attentional distribution.

Table 1
The feature values and estimated salience of the key distractors in Experiment 3

Size (cm)	Luminance (cd/m^2)	Salience
0.576	5.31	0.251
0.672	5.31	0.405
0.768	5.31	0.539
0.864	5.31	0.657
0.96	5.31	0.762
0.448	10.7	0.329
0.448	14	0.456
0.448	17.9	0.571
0.448	22.6	0.681
0.448	28.3	0.786

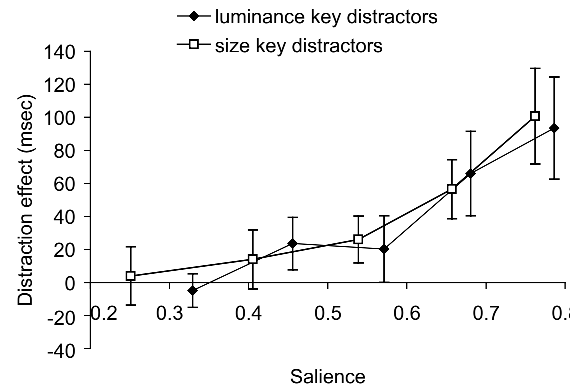

Fig. 5. Results of Experiment 3. The function how distraction effect vary with salience in two dimensions of luminance and size: the correspondence of the two curves confirms our equation for computing salience from feature values, as well as our constant R_{lum}: using that equation and that constant to compute key distractor salience, we find no systematic difference between the effects of distractors of the two feature dimensions.

7. Experiment 4

The purpose of Experiment 4 was to investigate the question of how salience from more than one dimension is summed. In this experiment, the feature values of the target and the background distractors remained constant (target: 17.9 cd/m^2 and 0.768 cm, distractor, 5.31 cd/m^2 and 0.448 cm). We measured the distraction effect of the size and luminance key distractors, each at 3 different salience levels, to provided a reference for other kinds of key distractors. There were four other kinds of key distractors (composite key distractors): two were defined mainly by size and secondarily by a very small luminance contrast (size composite key distractors); the other two were defined mainly by luminance and secondarily by a very small size contrast (luminance composite key distractors). The feature values and predicted salience values (in both dimensions) of the 10 types of key distractors are given in Table 2.

There are two usual answers to this question as it applies to quantities of physical magnitude: forces and momentums of more than one dimension are summed as vectors; mass and energy are summed as scalars. We fitted the results using two kinds of summation models: orthogonal vectors and scalars. The mean RT of

Table 2
The feature values and estimated salience of the key distractors in Experiment 4

Size (cm)	Luminance (cd/m^2)	Salience from size contrast	Salience from luminance contrast
0.768	5.31	0.539	0
0.864	5.31	0.657	0
0.96	5.31	0.762	0
0.768	6.69	0.539	0.109
0.864	6.69	0.657	0.109
0.448	17.9	0	0.571
0.448	22.6	0	0.681
0.448	28.3	0	0.786
0.512	17.9	0.133	0.571
0.512	22.6	0.133	0.681

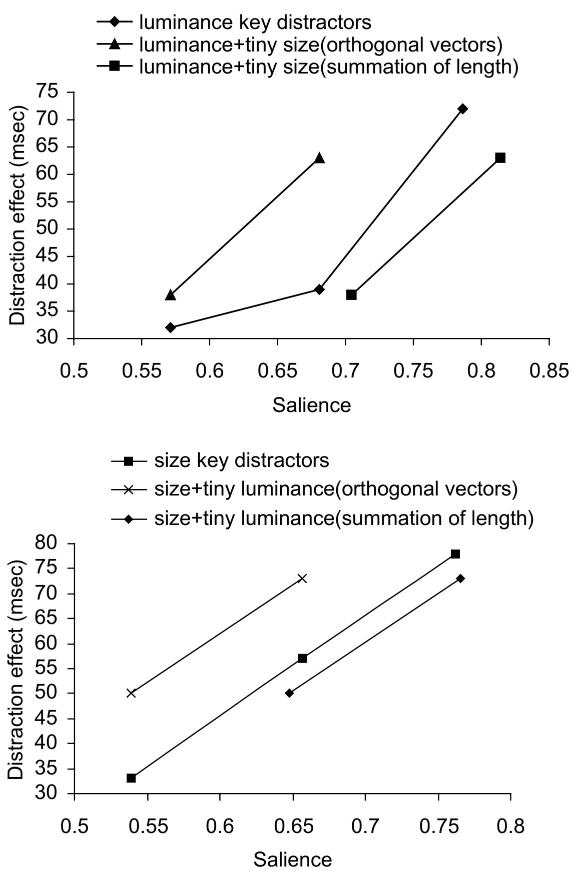

Fig. 6. Results of Experiment 4. The salience of composite key distractors is tentatively computed in two ways: the sum of the two orthogonal vectors, or the two orthogonal vectors. These two computations made distraction effect vs. salience curve of composite key distractors systematically higher (orthogonal vectors) or lower (scalars) than curves of single dimension key distractors. Therefore, the sum of salience arising from variation in two dimensions is larger than their vector sum (orthogonal vectors), but smaller than their scalar sum.

Experiment 4 was 756 ms. The results of the fitting are given in Fig. 6. These two computations made distraction effect vs. salience curve of composite key distractors systematically higher (orthogonal vectors) or lower (scalars) than curves of single dimension key distractors. Therefore, the sum of salience arising from variation in two dimensions is larger than their vector sum (orthogonal vectors), but smaller than their scalar sum.

For each composite key distractor, we used very different salience values for its two dimensions. The primary dimension's salience in each case was much greater than the secondary dimension's. How equal amounts of salience are combined is an interesting question, but unfortunately we could not ask it here, since a key distractor that was both much larger and much brighter than the background distractors would have looked very similar to the target and so would have taken considerably more time to be identified and rejected than the others in a later stage. This would have confounded the salience effect that we are investigating. Further study is needed along this line, perhaps with an improved version of our method.

In a study of how salience from different dimensions is combined, Nothdurft (2000) found that the combination is additive, but with some discount. Our results basically confirm his finding. In our experiment, the addition rate (defined as (*combined salience − main dimension salience*)/(*sub-dimension salience*)) can be estimated as 0.68 ± 0.11.

8. General discussion

In summary, our results indicate

(1) When the object/background feature difference increases in luminance or size, it becomes more salient. Its salience is decided by the ratio between its defining feature value and the corresponding feature value of background items; this conclusion is congenial to Weber's law. The function relating increase of salience to feature value ratios is similar for size and luminance dimensions; the difference between the two dimensions' rates of salience increase can be compensated with a constant. However, it should be mentioned that this conclusion probably does not apply to sizes beyond some maximum; if we test larger values, this conclusion (and also 2–3 below) will probably be invalid.

(2) Salience increases with increasing luminance at almost half the rate (0.47 ± 0.04) that it does with increasing size, so increases of object area and of luminance affect salience approximately to the same degree (0.86–1.02). It seems luminance and size are functionally related.

(3) The sum of salience arising from variation in two perceptual dimensions is larger than their vector sum (orthogonal vectors), but smaller than their scalar sum. When the salience from one dimension is much smaller than the other, 0.68 ± 0.11 of the salience in that secondary dimension is added to the overall salience.

Some of these findings are new to our knowledge, while others have been suggested in previous subjective measurements (Nothdurft, 1993a, 1993b, 2000). Even for those phenomena previously described, our results may be significant insofar as our method was very different from previous studies that used introspection or near-threshold measurement. There is no apparent conflict between previous findings and ours; the corroboration strengthens the validity of our results. The most important contribution this paper offers is probably the new approach of measuring salience. The equation we have proposed relating salience to feature value ratios (logarithm function) is apparently not applicable for some dimensions, like orientation. Further research is needed to investigate how salience increases in those dimensions.

The current study also offers some theoretical contribution to current issues in visual search. First, some investigators (Wolfe, 1994) have assumed that all feature dimensions can be weighted gradually by top-down control. Although that assumption is supported by some evidence of a priming effect (dimension-weighting account, Mueller et al., 1995), it had not, to our knowledge, been clearly demonstrated. Our Experiment 1 provides such a demonstration.

Second: if a search target has a certain size and brightness, and distractors are all smaller and dimmer, will an occasional distractor that is even larger or even brighter than target be more salient than target, or less? The results of our Experiments 3 and 4 suggest that a key distractor larger or brighter than the target is more salient than a distractor just as large or as bright as the target. This result argues against any model assuming that similarity is the only factor governing the search process. Our results indicate that at least in some early stage, salience is computed from feature values without relying on even a gross computation of similarity. It seems that similarity becomes important only in some later stage, like target identification, but has little role in the control of attention.

Finally, our technique for the objective measurement of suprathreshold salience may help researchers and engineers to test and improve a number of vision models and video encoding schemes, and to better predict the importance to human observers of different kinds of signal degradation across a wide variety of display technologies.

Acknowledgement

The research was supported by grant MH45584 from the National Institute of Mental Health.

Appendix A. Definitions, axioms and rationale in this study

We will first introduce definitions and mathematical equations we took for given.

Perceived overall difference (POD) is decided by two features in a certain dimension, reflecting how different they are psychologically. In this study we deal with POD_s and POD_l for size and luminance, respectively.

Task relevance of a dimension (TRD) is a 2-dimensional function (since only two dimensions are involved in this study): TRD (POD_s, POD_l). The output of TRD is also 2-dimensional (representing the TRDs that each dimension receives); to simplify, we call them TRD_s and TRD_l for size and luminance.

Salience is a 4-dimensional function: SAL (POD_s, TRD_s, POD_l, TRD_l).

The distraction effect of a certain key distractor is a 1-dimensional function of salience: DE (SAL)

Now we will introduce several axioms using these definitions:

1. The distraction effect of a certain key distractor is a monotonic function of its salience.
   If SAL1 > SAL2, then DE (SAL1) > DE (SAL2).
   That is to say, more salient key distractors will be more distractive.
2. SAL (POD1, TRD1, POD2, TRD2) can be simplified as SAL (POD1, TRD1) if POD2 is 0.
   That is to say, when there is no feature difference in one dimension, the salience is decided solely by the POD and TRD of the other dimension.
3. SAL (POD1, TRD1) is a monotonic function in both dimensions:
   If POD1 > POD2 and TRD1 > TRD2, then SAL (POD1, TRD1) > SAL (POD2, TRD2).
   That is to say, if one singleton has greater feature difference and greater task relevance of a dimension, then it is more salient.
4. if POD_l > POD_s, then
   TRD_l (POD_s, POD_l) > TRD_s (POD_s, POD_l).
   That is to say, a dimension with greater difference will be more important.
5. TRD (POD_s, POD_l) is a symmetric function in the sense:
   TRD_l (x, y) = TRD_s (y, x),
   TRD_s (x, y) = TRD_l (y, x).

That is to say, switching the relative importance of two dimensions will switch the TRDs.

Appendix B. Rationale of experiment 1

In experiment 1, as shown in Fig. 1, the size differences between the size key distractor and background distractors and between the target and background distractors were both POD_s. The luminance differences between the luminance key distractor and background distractors and between the target and background distractors were both POD_l.

According to Axiom 2, the salience of two kinds of key distractors are given as

SAL_s= SAL (POD_s, TRD_s (POD_s, POD_l)),
SAL_l= SAL (POD_l, TRD_l (POD_s, POD_l))

If POD_s > POD_l then
TRD_s (POD_s, POD_l) > TRD_l (POD_s, POD_l) (Axiom 4).
Then
SAL_s > SAL_l (Axiom 3).
Then
DE (SAL_s) > DE (SAL_l) (Axiom 1),
DITTO if POD_s < POD_l then DE (SAL_s) < DE (SAL_l).
Using reduction to absurdity, apparently if DE (SAL_s) = DE (SAL_l), then
POD_s = POD_l
So TRD_s (POD_s, POD_l) = TRD_l (POD_s, POD_l) (Axiom 5)
So SAL_s = SAL_l

Thus we prove that if the distraction effect is equal for these two kinds of key distractor, the PODs in the two dimensions are the same, the key distractors' salience is the same, and the two dimensions are weighted equally.

The rationale of Experiments 2–4 is simply based on Axiom 1.

Most axioms and rationales mentioned here have appeared in previous literature on visual search, though usually implicitly (e.g. Duncan & Humphreys, 1989; Yantis & Egeth, 1999; Wolfe, 1998; Wolfe, Cave, & Franzel, 1989).

References

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review, 96*, 433–458.

Lu, Z.-L., & Sperling, G. (1995). The functional architecture of human visual motion perception. *Vision Research, 35*, 2697–2722.

Mueller, H. J., Heller, D., & Ziegler, J. (1995). Visual search for singleton feature targets within and across feature dimensions. *Perception & Psychophysics, 57*, 1–17.

Nothdurft, H. C. (1993a). The conspicuousness of orientation and visual motion. *Spatial Vision, 7*, 341–366.

Nothdurft, H. C. (1993b). Saliency effects across dimensions in visual search. *Vision Research, 33*, 839–844.

Nothdurft, H. C. (2000). Salience from feature contrast: Additivity across dimensions. *Vision Research, 40*, 1183–1201.

Parkhurst, D., Law, K., & Niebur, E. (2002). Modeling the role of salience in the allocation of overt visual attention. *Vision Research, 42*, 107–123.

Titchener, E. B. (1908). Lectures on the elementary psychology of feeling and attention. New York: The MacMillan Company.

Yantis, S., & Egeth, H. (1999). On the distinction between visual salience and stimulus-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance, 25*, 661–676.

Weber, E. H. (1834). De pulsu, resorptione, auditu et tactu. *Annotationes anatomicae et physiologicae*. Leipzig: Koehler.

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

Wolfe, J. M. (1998). Visual search. In H. Pashler (Ed.), *Attention*. London UK: University College London Press.

Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance, 15*, 419–432.