Vision Research 79 (2013) 8–16

Contents lists available at SciVerse ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

Size averaging of irrelevant stimuli cannot be prevented

Chris Oriet[a,*], John Brand[b]

[a] Department of Psychology, University of Regina, Canada
[b] Department of Psychology, Concordia University, Canada

ARTICLE INFO

Article history:
Received 29 May 2012
Received in revised form 12 November 2012
Available online 27 December 2012

Keywords:
Perceptual averaging
Divided attention
Set representation
Attention
Gist

ABSTRACT

Research suggests that subjects can compute the mean size of two sets of interspersed objects concurrently, but that doing so incurs a cost of dividing attention across the two sets. Alternatively, costs may arise from failing to exclude irrelevant items from the calculation of mean size. Here, we examined whether attention can be selectively deployed to prevent the inclusion of items from an irrelevant, concurrently displayed set in the computation of the relevant set's mean size. The results suggest that mean size is computed prior to the deployment of attention, failing to exclude processing of items that are irrelevant to the task. The influence of the irrelevant items is evident both with brief exposures of the set (200 ms) and in a simultaneous judgment task with unlimited viewing time, suggesting that attention cannot be effectively deployed to facilitate selective averaging of the size of the relevant set. Size averaging appears to precede the deployment of selective attention, suggesting that it may be carried out automatically, without intention.

1. Size averaging of irrelevant stimuli cannot be prevented

The tendency of the visual system to extract statistical properties from a set of items rather than to encode its individual elements is well documented. Researchers have demonstrated that observers can compute the average orientation (Dakin & Watt, 1997; Parkes et al., 2001), brightness (Bauer, 2009), direction of motion (Watamaniuk, Sekuler, & Williams, 1989), or size (Ariely, 2001; Chong & Treisman, 2003, 2005a) of sets of similar items, and can even summarize higher order properties of sets of more complex objects, such as the mean emotion (Haberman & Whitney, 2007, 2009), sex (Haberman & Whitney, 2007), or identity (de Fockert & Wolfenstein, 2009) of sets of faces. Such statistical summaries are argued to be formed outside of the focus of attention (Alvarez & Oliva, 2008; Chong & Treisman, 2003) and to rely on representations established relatively early in perceptual processing before subjects become aware of the identity of individual objects (Choo & Franconeri, 2010; Corbett & Oriet, 2011; Haberman & Whitney, 2011; but see Jacoby, Kamke, & Mattingley, in press).

Much recent research has focused on observers' abilities to compute the average size of objects, and the extent to which this ability can be characterized as an automatic process. Although debate persists about whether the number of items that can be summarized lies within known limits of focal attention and working memory (de Fockert & Marchant, 2008; Marchant, de Fockert, & Simons, 2013; Myczek & Simons, 2008; Simons & Myczek, 2008) or exceeds those limits (Chong et al., 2008; Chong & Treisman, 2003, 2005a, 2005b; Robitaille & Harris, 2011), researchers are generally in agreement that some sort of ensemble representation forms the basis of judgments in size averaging tasks (Alvarez, 2011). The extent to which size averaging is automatic, however, is a subject of current debate, with both the speed of the process (Whiting & Oriet, 2011) and the extent to which attention can be divided to compute the mean sizes of two sets of objects without cost (Brand, Oriet, & Sykes Tottenham, 2012) called into question. In this paper, we address a different aspect of the automaticity of perceptual averaging, focusing on the extent to which size averaging is carried out without conscious intent for items that are to be ignored.

Visual sets often contain distinct sub-groups, and occasionally it is useful to compare the statistical properties of one group to the statistical properties of another. For example, consider the situation in which the driver of a large vehicle is looking for a parking spot in a crowded parking lot. One strategy for locating a large spot is to quickly estimate the average size of vehicles in the different parts of the lot; the part of the lot with the larger vehicles is more likely to have the larger spots. In this example, the ability to rapidly compare the average size of two spatially-segregated sub-groups is beneficial in circumventing the need to undertake time-consuming serial search. As such, it seems reasonable that the visual system should be able to compute the mean size of two different sets of items concurrently after segregation into sub-groups, with a fair

* Corresponding author. Address: Department of Psychology, University of Regina, 3737 Wascana Parkway, Regina, Saskatchewan, Canada S4S 0A2.
E-mail address: chris.oriet@uregina.ca (C. Oriet).

0042-6989/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.visres.2012.12.004

degree of precision. Such a process would be especially useful if this segregation could be done preattentively, minimizing the potential for interference between the groups to be compared.

Chong and Treisman (2005b) addressed these possibilities by examining whether mean size can be computed for two different sets of objects (circle sets distinguished by color). Subjects viewed a display containing two sets of circles, one in red and one in green, randomly distributed throughout the display. Following the offset of this display, two probe circles were shown. The size of one circle corresponded to the mean size of one of the sets (the relevant set) and the size of the other circle was titrated to yield 75% correct performance across trials. Subjects were instructed to determine which of the two probes corresponded to the average size of the circles shown in the display in the relevant color. On some trials, the relevant set was pre-cued with two lines appearing in the relevant color prior to the onset of the display. On other trials, the relevant set was post-cued; subjects learned which color set's mean was to be reported when the probes appeared, with the probes sharing the relevant set's color. In a control condition, only one of the two color sets was displayed.

Chong and Treisman (2005b) reasoned that if mean size is computed following the segregation of items into sets, subjects should be able to reliably identify which of two probes corresponds to the mean size of one of the sets. Further, if the mean size of both sets can be computed concurrently, subjects should be able to perform this task at rates better than expected by chance even when they are not told in advance which set will be cued. Finally, if the two means can be computed concurrently without cost, they reasoned that there will be no advantage to telling subjects in advance which mean will have to be reported at the end of the trial because the means of both sets will be computed in parallel. The authors found that pre-cueing the relevant color set yielded no better performance for estimating the mean size of the relevant set than displaying the sets first and then post-cueing the relevant color. Chong and Treisman interpreted this finding as evidence that subjects were able to first preattentively segregate the two sets of items on the basis of color and then calculate the mean size of the two sets concurrently, without cost. The authors concluded that size averaging is an automatic process, able to be carried out for two sets of items as easily as for one.

Recently, Brand et al. (2012) confirmed that subjects could indeed compute two means concurrently, but challenged Chong and Treisman's (2005b) assertion that attention could be divided across the two sets without cost. Brand et al. argued that the method used to probe performance in their task allowed subjects to circumvent the need to compute two means, which in turn yielded similar performance across the pre-cued, post-cued, and single set conditions. By eliminating this confound, Brand et al. demonstrated reliable costs of post-cueing the relevant set in Chong and Treisman's task, suggesting that although two means could be computed concurrently, there was a clear cost in performance when the relevant set was post-cued.

The fact that performance improves when the relevant set is pre-cued suggests that selective attention functions in some way to protect the calculation of the relevant set's mean from interference from the irrelevant set. One possibility is that selectively attending to the pre-cued set allowed subjects to filter out the irrelevant set altogether, restricting attention to processing only the items in the relevant set. Consistent with this interpretation, Brand et al. (2012), like Chong and Treisman (2005b) observed similar performance in the pre-cued and single set conditions, as would be expected if pre-cueing the relevant set effectively reduced the pre-cue condition to the single set condition.

Alternatively, if mean size is computed for all visible items in a display prior to selection, subjects may not be able to prevent inclusion of items from the irrelevant set in the calculation of the mean size of the relevant set. Consistent with this interpretation, Brady and Alvarez (2011) found that subjects appeared to encode separate mean sizes of two colored sets of circles when instructed to do so, but encoded the mean of the entire set when not instructed to encode separate means. This suggests that encoding two means is indeed more effortful than encoding one, but it does not speak to the present question of whether subjects can selectively deploy attention to a relevant set of items to prevent including items from an irrelevant set when computing the mean size of the relevant set.

A limitation of the cueing paradigm discussed above is that it is not adequate for addressing the question of whether calculation of mean size can be prevented. In the post-cued condition, either set can be cued, so subjects must try to compute both sets' means, but in the pre-cued and single set conditions, attention could either be restricted to one set, or the mean sizes of both sets could be computed concurrently. Both accounts predict similar performance across the single set and pre-cued conditions in which only one set's mean needs to be remembered, and both predict costs relative to the post-cued condition, in which both sets' means need to be remembered. Thus, the post-cueing paradigm cannot address the question of whether attention can be selectively deployed to the relevant set, excluding processing of the irrelevant set altogether.

In the present work, we introduce a new paradigm that allows us to determine whether selective attention can be deployed to the relevant set to prevent processing of the irrelevant set. We asked subjects to compute the mean length of one set of objects while the mean length of a second set was experimentally manipulated. If variation in the irrelevant set influences judgments of the mean size of the relevant set despite the instruction to focus attention on one set only, then selective attention is unable to protect the calculation of the relevant set's mean from interference that arises from the presence of an irrelevant set. In this scenario, finding that the mean size is biased by the presence of irrelevant items, even when processing these is unnecessary and potentially detrimental to performance, would be a clear indication that calculation of the relevant set's mean size occurs prior to deployment of attention to the relevant items. If judgments of mean size are accurate in the presence of this potentially interfering information, then selective attention serves to prevent processing of the irrelevant set altogether, suggesting that calculation of mean size is not strongly automatic and can be postponed until attention has selected the relevant set when it is advantageous to do so.

In Experiment 1, we demonstrate that items from an irrelevant set are included in computation of a relevant set's mean even when subjects attempt to selectively attend to one set and ignore the other. In Experiment 2, we confirmed that unintentional processing of items from an irrelevant set interferes with judgments of the relevant set's mean length by showing that interference persists even when subjects are given unlimited time to view the two sets.

2. Experiment 1

In Experiment 1, subjects viewed displays of horizontal lines interspersed among vertical lines, and were instructed to compute the average length of one set while ignoring the other. Subjects compared two sets of six items directly, one on the left half of the display and one on the right (Fig. 1); this method was previously used by Chong et al. (2008) among others. In two conditions, each (relevant) set to be compared was accompanied by an irrelevant set. In these conditions, we varied whether the mean length of the irrelevant set matched or mismatched (i.e., was 20% or 40% shorter or longer than) the mean length of the relevant set. If subjects can filter out the irrelevant set altogether, no effect of varying