The infidelities of summary statistics

Sujala Maharjan^{*}, Chad Dubé[†], and Robert Sekuler

Brandeis University and University of South Florida[†]

Abstract

The briefest glimpse of some scene can suffice to allow the extraction of a scene's overall character. That brief glimpse produces a compressed, prototype representation of the scene, that is, its gist. Many studies have provided evidence for the remarkable accuracy of such compressed/prototype representations, but few have explored factors that control the emergence of such representations, or possibly distort the products. We addressed these questions by using small arrays of luminances. We focused on how clustering or repeating items impact perception of an array. The results showed that repetitions lowered estimation errors, though the mean of the array continued to strongly predict responses. The data suggest that subjects may capitalize on informational redundancy and emergent structure in computing summary representations. The implications for theoretical accounts of ensemble coding are discussed.

Keywords: Summary statistics, Ensemble statistics, Redundancy

From just a fleeting glimpse, a viewer can extract the general sense (gist) of a complex display or scene (Biederman, Rabinowitz, Glass, & Stacy, 1974). This perceptual ability, which

^{*} Now at Lincoln Financial Group, Boston. Address correspondence to Chad Dubé, chaddube@usf.edu

entails a substantial compression of information, is successful with a wide range of stimuli, including crowds of people whose facial expressions vary (Haberman & Whitney, 2007, 2009),
5 collections of simple geometric forms (Chong & Treisman, 2003), displays of moving random dots (Williams & Sekuler, 1984), arrays of items that vary in color (Maule & Franklin, 2015), and collections of Gabor patterns that vary in spatial frequency (Dubé & Sekuler, 2015). Recently, clear evidence has shown that such effects are not limited to perception, but also can impact visual short term memory (Dubé, 2019; Tong, Dubé, & Sekuler, 2019).

10 To quantify the accuracy with which a scene's information has been compressed and its gist extracted, laboratory studies have tended to hone in on one metric: the perceived or remembered average of the stimulus array. A great deal of the work dealing with this metric has focused on effects of the heterogeneity of stimuli in an ensemble.

For example, Chong and Treisman (2003) asked subjects to compare the mean sizes
15 of two stimulus arrays. Specifically, each display in their study was divided into two halves vertically, with each half containing 12 circles. These circles in each stimulus set were either identical (homogeneous) or were a mixture of four sizes (heterogeneous). Subjects could report the side of the display with the larger mean size for diameter differences as small as 6-8% when the display was homogeneous. This suggests that the redundant presentation of
20 multiple identical circles facilitates the averaging process. Similar results have been reported with color averaging where displays consisting of perceptually similar colors were accurately averaged compared to displays with circles of perceptually different colors (Maule & Franklin, 2015).

Extracting the mean from a briefly presented array (as short as 50 ms in Chong and
25 Treisman's work) is a striking perceptual achievement, but relatively little attention has been given to what makes it possible, other than the array's heterogeneity, exposure duration, and set size (Haberman & Whitney, 2010; Chong & Treisman, 2003; Maule & Franklin, 2015). However, these are clearly not the only variables that can bias the perceived average: an item's (or items') position within a sequence of items also has a substantial effect (Hubert-
30 Wallander & Boynton, 2015; Tong et al., 2019). Our overall aim was to investigate how the

arrangement of items within an array would influence the accuracy of *ensemble statistics*. For this, we chose variables that, in other, very different conditions, had been shown to affect salience (Itti & Koch, 2001) and perceptual grouping (Wagemans et al., 2012). In addition, in order to maximize ability to detect any effects of our experimental manipulations, we designed arrays comprising a relatively small number of items.

Experiment One

Although Chong and Treisman (2003)'s results do show that inter-item similarity facilitates extracting the mean size of array items, their results do not explain how ensemble perception is influenced by the presence of both similar and dissimilar items. Furthermore, results on the extraction of mean size do not necessarily generalize to the computation of other averages, such as mean luminance, mean direction, etc. Therefore, we asked if ensemble processing differs between (i) arrays with random luminances, and (ii) arrays with some repeating luminances among the random ones. Does the repetition of items facilitate or hinder our ability to compute ensemble statistics? We also asked whether ensemble coding is more sensitive to similar items when they are grouped rather than distributed. In other words, do details of how repeated items are organized influence ensemble coding and, if so, by how much?

Incidental and implicit learning via repetitions has long been known to strengthen subsequent memory retrieval, likely via transfer to secondary storage (Hebb, 1961; Kimble, 1967). Additionally, informational redundancy, in the form of repeated items, is often confounded with perceptual similarity. In fact, some classic work in verbal learning of meaningless stimuli (e.g., random sequences of digits) suggests that redundancy may actually be a more powerful determinant of short-term memory performance than similarity alone (Kintsch, 1970). Finally, recent work on sequential averaging of visual stimuli has suggested weights are assigned to items in proportion to the fidelity of their memory representations (Tong et al., 2019). Our study attempts to determine whether repetition and clustering of luminances facilitate computation of *ensemble statistics*.

The work discussed above suggests a reasonable hypothesis, namely that, overall,

repetitions within a spatial array will disproportionately influence ensemble coding, and that clustering these repetitions together will further bias judgments toward the repeated items. We also hypothesize that error in *ensemble statistics* will decrease as stimulus duration increases because longer exposures could amplify the influence of any repeated items.

To test these hypotheses, we exploited one of psychophysics' most venerable tools, direct matches made by the subject. Specifically, subjects were asked to adjust a random initial luminance to match from memory the mean luminance of a just-seen array. This method generates an analogue assay of a stimulus' appearance on each trial. Analogue measures afford sensitivity sufficient to measure important effects that might be too subtle to detect in the less sensitive, binary (*e.g.*, "Yes" and "No") responses generated by other methods.

Methods

Subjects

Twenty young adults, 18 to 30 years of age, were recruited for this study. Subjects had normal or corrected to normal vision (measured with Snellen targets), and were compensated \$10 (U.S.) each for participating. Subjects gave written consent to a protocol approved by Brandeis University's Committee for the Protection of Human Subjects.

Apparatus

Stimuli were generated in Matlab (version 8.5) using the Psychophysics Toolbox extensions (Brainard, 1997). An Apple iMac computer presented stimuli on a cathode ray tube display (Dell M770) set at 1024 X 768 pixels resolution, and 75 Hz frame rate. A constant, gray background on the display was fixed at 31.68 cd/m². A chin rest enforced a viewing distance of 57 cm. The room was darkened during the experiment.

Stimuli

Stimulus luminances were drawn from a library of 16 different luminances whose items varied from 1.05 to 99.83 cd/m², in 15 equally-spaced logarithmic steps. The sequence

(a)

(b)

Figure 1. a: The sequence of events within a trial. After fixation, a spatial array appeared for 40, 80, or 120 ms, following which subjects reproduced the average luminance of the spatial array using a response dial. *b:* Schematic of *Random*, *Clustered* and *Distributed* spatial arrays. In this (and all figures), colors are intended to approximate those used in the experiment and should not be taken to represent those colors precisely.

of events on a trial is shown schematically in Fig.1a. Each trial began with a brief (~400 ms) presentation of a central fixation cross, after which a spatial array consisting of eight luminances was presented simultaneously. The array, which was centered on the display, spanned ~10° in width. The array's total extent, 10°, was designed to ensure the perceptibility of its elements even under the brief exposure durations we use (as short as 40 ms; see Experimental task) (Woutersen, van den Berg, Boonstra, Theelen, & Goossens, 2018). Adjacent luminances were separated by 4 minutes of arc, making neighboring items perceptually distinct from one another. After the luminance array disappeared, a single random probe luminance was presented as a 1.25 × 1.25° square. The probe's luminance was a random sample from a uniform distribution of 216 luminances, on a linear scale ranging from 1.05 to 99.83 cd/m². Subjects adjusted the random luminance probe to reproduce from memory the average luminance of the spatial array that had just been seen.

Fig.1b illustrates the three stimulus conditions of the experiment. On a given trial, subjects saw either a random luminance sequence ("Random") with no repeated luminances, a sequence with repeated luminances that were clustered in sequence ("Clustered"), or a sequence with repeated luminances that were not clustered ("Distributed").

Experimental task

The experimental design was mainly defined by two independent variables: Stimulus duration and type, with three levels each. Stimulus types were *Clustered*, *Distributed* and *Random*; durations were ~40 ms, ~80 ms and ~120 ms, which we will designate "short," "medium" and "long", respectively. Previous work showed that judgments of mean size were quite robust with stimulus duration, remaining accurate down to quite brief presentations (~50 ms) (Chong & Treisman, 2005b) and over a wide range of sizes (Chong & Treisman, 2003). The three types of trials, *Clustered*, *Distributed*, and *Random*, were interleaved and presented for either 40 ms, 80 ms, or 120 ms. In *Random* trials, all the luminances within an array were different from each other. In a *Distributed* trial, one luminance was chosen at random and repeated three times within a spatial array. The algorithm that generated stimulus arrays

110 prevented repeated items from occupying adjacent locations. In a *Clustered* trial, one of the array's luminances was chosen at random and repeated three times, with all three repeated item occupying adjacent locations in the array. The three repeated items could be presented anywhere within a spatial array.

Immediately following each spatial array, a fixation point appeared and stayed on
 115 the screen for ~ 400 ms. Once the fixation point disappeared, subjects reported the mean luminance of the spatial array via an adjustable dial (PowerMate USB Multimedia controller, Griffin Technology, USA). The dial had 100 steps, with the lower limit set at 1.05 cd/m^2 and the upper limit set at 99.83 cd/m^2 ; luminance on-screen varied linearly with dial position. Once subjects were satisfied with their response, they pressed the dial to register their choice. No
 120 feedback regarding accuracy was provided. The next trial began ~ 670 ms after subjects registered their response.

Design

There were 495 trials in total, of which the first 27 were treated as practice trials. A single test session lasted ~ 50 min. Of the 495 trials, 165 were short trials (arrays presented for
 125 ~ 40 ms each), which were further divided into *Clustered* (55), *Distributed* (55), and *Random* (55). The remaining 330 trials were equally divided into medium (presented for ~ 80 ms) and long (presented for ~ 120 ms) trials, each of which were further sub-divided among *Clustered*, *Distributed* and *Random* conditions. All conditions were equally represented in the 27 practice trials at the beginning of the experiment.

Results

The reproduction errors were calculated using data from each trial, as detailed in Equation 1. That is, for a particular subject's i^{th} array we divided the absolute difference between the actual arithmetic mean of the array ($\bar{M}_i$) and the subject's response ($\bar{m}_i$) by the actual mean of the array. The errors were then summed over the N arrays for each subject j
 135 within a given condition k . The result constituted the basic datum entered into the ANOVA

and other analyses reported below.

$$jk^{th} \text{ Reproduction Error} = \sum_{i=1}^N \frac{|\bar{M}_{ijk} - \bar{m}_{ijk}|}{\bar{M}_{ijk}} \quad (1)$$

One subject's data were discarded because they made the same response on nearly every trial, suggesting that, for whatever reason, their behavior was not under stimulus control. The initial analysis of results entailed a 3 (stimulus type: *Clustered*, *Distributed*, *Random*) × 3 (duration: 40, 80, 120 ms) repeated-measures ANOVA on reproduction error. There was no main effect of duration and no interaction between duration and stimulus type. But, there was a main effect of stimulus type ($F(2,36) = 8.97$, $p < .001$, partial $\eta^2 = .33$), as shown in Fig. 2. This indicated that the presence of repeated items within an array significantly influenced subjects' memory for the mean luminance of the items in the array. Planned comparisons showed a significant difference between *Random* and *Clustered* ($t(36) = -3.37$, $p < .01$) and between *Random* and *Distributed* ($t(36) = -3.90$, $p < .001$) stimuli. So, repetition of items, either as clusters or distributed over the array, does influence subjects's judgments of the array's mean. Inspection of Fig. 2 reveals the nature of that influence, which was to reduce error when items were repeated within an array.

150 **The influence of various stimulus attributes**

To determine how the arrangement of repeated items in *Clustered* and *Distributed* stimuli influenced subjects' reproduction errors, we did model selection with a hierarchical linear mixed effects regression analysis. The three nested models we evaluated are shown below.

$$\text{Response} = \beta_0 + \beta_1 \overline{\text{Array}} + (1|\text{SubjectID}) + \epsilon \quad (2)$$

$$\text{Response} = \beta_0 + \beta_1 \overline{\text{Array}} + \beta_2 \text{Probe} + (1|\text{SubjectID}) + \epsilon \quad (3)$$

$$\text{Response} = \beta_0 + \beta_1 \overline{\text{Array}} + \beta_2 \text{Probe} + \beta_3 \overline{\text{RepeatedItems}} + (1|\text{SubjectID}) + \epsilon \quad (4)$$

Figure 2. Mean reproduction error with *Random* (red), *Clustered* (green), and *Distributed* (blue) stimuli, over all stimulus durations. Error bars are $\pm$ within-subject standard errors of the mean (Cousineau, 2005).

Model 1 assumes that subjects' responses reflect only the actual arithmetic mean of the spatial array, plus a random error, ϵ . Model 2 allows for possibility that responses are also influenced by the luminance of the random initial luminance of the test probe that subjects adjusted. Model 3 assumes that in addition to the variables in Model 2, subjects' responses might be influenced also by the mean luminance of the repeated items.

Table 1 shows the results from comparing the three models. According to nested χ^2 tests, overall, Model 3 was the best-fitting model, producing a significant improvement over Model 2 in every condition (all $p < .01$). For most conditions (that is, excluding the Distributed 80 ms and Distributed 120 ms conditions), this improvement, due to the addition of the repeated-items coefficient in Model 3, was not solely the result of having one extra parameter, as Model 3 also had the lowest Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) values. These results suggest that subjects were influenced not only by the display's mean

Table 1
Evaluations of the nested regression models

Predictor	Stimuli					
	Clust40	Dist40	Clust80	Dist80	Clust120	Dist120
Model 1						
AIC	5894.5	5937.2	6031.6	5908.7	6060.1	5979.4
BIC	5913.2	5955.7	6050.3	5927.3	6078.8	5998.1
Model 2						
AIC	5840.7	5883.3	5998.7	5862.3	6003.5	5902.1
BIC	5864.0	5906.7	6022.0	5885.6	6026.8	5925.5
Model 3						
AIC	5811.0	5871.9	5972.1	5864.3	5995.8	5901.6
BIC	5839.0	5899.9	6000.0	5892.2	6023.8	5929.6
<i>Best model estimates</i>						
χ^2	31.691	13.475	28.627	48.348	9.696	79.273
<i>p-value</i>	1.80e-08	0.0002	8.77e-08	3.57e-12	0.001	<2e-16

Note: Highlighted values in the table show the lowest values of Akaike Information (AIC) and Bayesian Information (BIC) criteria. Model 3, which included three predictors, DisplayMean, InitialProbe, and RepeatedItems, gave the best fit for *Clustered* conditions, whereas Model 2, which excluded RepeatedItems as a predictor, was the best fit for the majority of the *Distributed* conditions. This suggests that when repeated items are clustered together in the array, they have greater influence on subjects' reproduction of the mean.

luminance but also by the initial luminance of the probe and, additionally, by the luminance of the repeated items. However, looking specifically at the *Distributed* conditions, it is clear that Model 2 was the better fitting model according to AIC and BIC in the 80 ms and 120 ms conditions.

Taken together, the modeling results indicate a significant influence of the repeated items on subjects' perception of the arithmetic mean of the array when the repeated items were arranged right next to each other, and a less than significant influence when those repeated items were distributed across the array. The use of nested models was especially helpful in ferreting out subtle differences between the impacts of the repeated items in *Clustered* and *Distributed* trials.

Figure 3. A. Relative importance of Model 3 predictors for 80 ms *Clustered* condition. DisplayMean and RepeatedItems are the two most influential predictors, carrying almost all of the variance explained by Model 3. **B.** Relative importance of Model 2 predictors for 80 ms *Distributed* condition. DisplayMean is the most influential predictor and carries almost all of the variance explained by the model. In both panels, error bars are 95% bootstrapped confidence intervals.

Relative importances of predictors

To characterize the relative importances of predictors in the best-fitting model, we calculated the relative importance metric using R's 'relaimpo' package (Grömping, 2006). In order to parse out each predictor's distinct influence, the 'lmg' metric (see Lindeman, Merenda, & Gold, 1980) averaged sequential sums of squares over all orderings of the predictors. This made it possible to rank the predictors in order of their importance in a best-fitting model, and to characterize their share of R^2 in the model. RepeatedItems influenced subjects' responses almost at the same level as DisplayMean in *Clustered* conditions (see Fig. 3a for 80 ms *Clustered* relative importance plot). DisplayMean was a major influencer in the best-fitting model corresponding to the *Distributed* conditions (see Fig. 3b for 80 ms *Distributed* relative importance plot). Again, these results further emphasize the important influence of the clustered repeated items on subjects' responses.

Discussion

190 When an array contained repeated items, those items had disproportionate influence on subjects' responses, causing the responses to more closely approximate the mean of the array (Figure 2). The influence of repeated items was amplified when those items clustered together in the array, rather than spread out over an array (Table 1).

One possible mechanism that could have given rise to our results is random sub-sampling
195 of the items within an array (de Fockert & Marchant, 2008; Myczek & Simons, 2008). On this view, subjects randomly sample some items from a display, and report the mean of the array based on the sampled items. This strategy of randomly sampling (or, non-selective readout; Averbach & Coriell, 1961) could explain why presentation duration had little or no effect: sampling just a few items from an array could have occurred quickly. Such a strategy
200 could also explain the reproduction error difference between *Random* vs. *Clustered* and *Distributed* arrays. Specifically, when a few items from either *Clustered* or *Distributed* arrays were randomly sampled, repeated items in the display would constitute the majority of the items within an array, whereas the same process for *Random* arrays would not lead to a similar result. This could be one of the potential reasons for higher reproduction error on *Random*
205 arrays compared to *Clustered* and *Distributed* arrays. This strategy could also account for the differential influence of repeated item arrangement on performance. The repeated items within a *Clustered* array could have been more prominent due to their arrangement resulting in their consistent selection during a *non-random* sub-sampling process (or, *non-selective* readout; Averbach & Coriell, 1961) whereas the spread out arrangement of the repeated
210 items in a *Distributed* array might have made them less prominent and prevented them from being consistently selected. As a result, performance might have been more influenced by the repeated items on the *Clustered* arrays compared to the *Distributed* arrays.

Nonetheless, the analysis of relative importance does suggest that subjects are unlikely to subsample only a single item, or we would not have obtained such a substantial contribution
215 from the display's mean luminance. In other words, subjects do appear to be integrating at least some of the information in the displays. Future work should further explore the influence

of item arrangement on *ensemble statistics* by comparing performance on *Clustered* and *Distributed* arrays that contain similar items.

In summary, using a method of adjustment to study the influence of repetitions on
220 *ensemble statistics*, we found that repetition of items plays an influential role in determining the
remembered average of an array. Further, this influence varies according to the arrangement of
the repetitions within an array. Although previous studies have shown an influence of similarity
in visual working memory (Kahana & Sekuler, 2002), this is the first study to demonstrate an
effect of repetition on *ensemble statistics* and more importantly, the difference in influence of
225 repetition depending on how such repetitions are arranged.

Experiment Two

Experiment One examined how stimulus duration, item repetition and item clustering
influenced *ensemble statistics* computation. To anticipate, the results showed that the re-
dundancy and arrangement of items that compose a visual display dictate how ensemble
230 representations are formed. More specifically, if the items within a display are exactly repeated,
those tend to weigh more on mean luminance judgment. Further, this influence of repetition
differs significantly between trials where repeated items are grouped together rather than
spread out. Repeated items presented adjacent to each other within an array exert a larger
influence on performance compared to repeated items that are distributed throughout an array.
235 Experiment Two's main aims were to assess how subjects' *ensemble statistics* computation
was affected by the number of repeated items within an array, and how that influence might
interact with the arrangement of such repeated items.

In determining how best to approach our questions related to arrangement and number
of repetitions, it is useful to consider prior results relating the density of spatial arrangements
240 to mean estimates. A study by Chong and Treisman (2005b) is particularly informative in
this respect. They tested the effects of number and density on mean size discriminations
between circles presented in two spatially segregated arrays. They varied the number of circles
and how densely or sparsely the circles were arranged in both halves of the display and had

their subjects identify the side with the larger mean size. Their results showed no significant
245 effect of either varying the number of circles or density. Accuracy at discriminating the half
with the larger mean size was comparable across displays and within a display. Particularly,
discriminating larger mean size on displays consisting of a dense half and a sparse half was
similar to discrimination made on displays with both halves either sparse or dense. This result
is also similar to one that Dakin (1997) obtained with the eponymous Glass patterns, Moire
250 fringes set up when two random dot patterns are overlaid and rotated relative to one another
(Glass, 1969). Dakin found that the pattern's sparseness significantly affected the accuracy
with which the mean orientation of the fringes could be discriminated. Taken together, these
results suggest that density has an inconsequential effect on the accuracy of averaging.

Although the absence of density effects in the studies reviewed above is convincing,
255 their investigation was incomplete because it did not directly examine the influence of density
on averaging performance. As the results from Experiment One showed, our subjects placed
more emphasis on the side-by-side presentation of replicated items over distributed replicated
items, suggesting a significant influence of density on mean luminance computation. To further
confirm the results we obtained from our previous study, we investigated the effects of density
260 under conditions in which the numerosity of the replicated items varied. This variation in
numerosity is especially critical as it reflects what might be encountered in nature. For example,
in natural scenes, similar elements may not only be present but also differ in quantitative
attributes and an observer might perceive a scene differently based on the frequency of such
elements. Importantly, it also helps establish the threshold at which the previously observed
265 density effect either disappears or is further enhanced. Moreover, it could also elucidate the
reasons behind the absence of numerosity effects in Chong and Treisman (2005b)'s study
in which mean size discrimination was unaffected by the difference in the number of circles
between the two halves of a display.

In the current study, we presented subjects with *Clustered* and *Distributed* trials. The
270 two types of trials consisted of varying numbers of repeated items. As in Experiment One,
subjects' task was to estimate and reproduce the average luminance of a given spatial array.

In addition to assessing the influence of the changing number of repeated items and their arrangement, this study also investigated the interaction between the two factors.

Figure 4. Schematic of 2rep, 3rep, and 4rep conditions within *Clustered* and *Distributed* spatial arrays. Note: The repeats could be presented anywhere within a spatial array. In this (and all figures), colors are intended to approximate those used in the experiment but should not be taken to represent those colors precisely due to the limited availability of colors in the software program chosen to design the schematic and display equipment.

Methods

Subjects

Eighteen young adults, between 18 and 30 years of age, with normal or corrected to normal vision, took part in this study. Subjects were compensated ten dollars (U.S.) for participating. Each gave written consent to a protocol approved by Brandeis University's Committee for the Protection of Human Subjects. None had served in Experiment One.

Apparatus and stimuli

Except where noted, the apparatus and stimuli were identical to those from Experiment One.

Experimental task

As in the previous experiment, subjects adjusted a probe stimulus to match their memory for the average luminance in each trial (Figure 1a). The experimental task was divided into 3×2 levels. The two main factors were: Repeat set size and stimulus type. Repeat set size had three levels, namely 2rep, 3rep, and 4rep and stimulus type had two levels, namely *Clustered* and *Distributed*. In repeat set size trials, a randomly chosen luminance within a spatial array was repeated twice (2rep), thrice (3rep), or four times (4rep). All of the repeat set size trials were presented in both a *Clustered* and *Distributed* manner. The conditions are shown schematically in Fig. 4.

As Experiment One showed no main effect of array duration, we decided to present all stimulus arrays for just a single duration, ~ 80 ms, the medium duration used in Experiment One. Prior to the main experimental task, in a short experimental session, subjects were asked to recreate a single luminance item presented to them for ~ 80 ms. This was done to get an estimate of error for a single item and could potentially be used to evaluate spatial array averaging error in the main experimental session. With the exception just mentioned, the timing of all stimulus events and the method by which subjects reported their responses were identical to those already described for the previous experiment.

Design

The short experimental task consisted of 30 trials of which the first 5 were treated as practice. The main task had 438 total trials, of which 24 were treated as practice trials. It took ~ 40 min to complete a session. Out of 438 trials, 50% were *Clustered* and the rest were *Distributed* trials. The 219 *Clustered* trials were divided into 2rep (73), 3rep (73), and 4rep (73) trials. The remaining 219 *Distributed* trials were apportioned in the same manner. In the 24

practice trials subjects received during the main experimental session, all types of trials were equally represented.

Figure 5. Mean reproduction error over all set sizes in Experiment One. Error bars represent within-subject standard errors (Cousineau, 2005).

Results

We calculated reproduction error for each trial using Equation 1. We then performed
 310 a 3 (repeat set size: 2rep, 3rep, 4rep) $\times$ 2 (stimulus type: *Clustered*, *Distributed*) repeated
 measures ANOVA to examine the influence of repeat set size and stimulus type on reproduction
 error. There was a main effect of repeat set size ($F(2,34) = 16.32$, $p < .001$, partial $\eta^2 = .49$)
 but there was neither a main effect of stimulus type ($F(1,17) = 0.25$, $p = .62$, partial $\eta^2 = .01$)
 nor a significant interaction between repeat set size and stimulus type ($F(2,34) = 0.01$, $p = .99$,
 315 partial $\eta^2 = .0003$). Planned comparisons between the repeat set size conditions revealed a
 significant difference between 2rep and 4rep ($t(34) = -4.742$, $p < .0001$) and 3rep and 4rep
 ($t(34) = -5.131$, $p < .0001$) conditions. This suggested that when half of the items within a

spatial array are repeated, subjects' perception of the mean grows significantly more accurate. Figure 5 shows how manipulating the number of repeated items within a spatial array affected subjects' reproduction of the array's mean. It appears, then, that reproduction error tracks the number of repeated items, with more repeated items producing smaller errors.

Influence of various features on performance

We did not observe any effect of stimulus type on reproduction error. There are two possibilities for this result: first, different arrangement of repeated items does not affect perception of the mean; second, the differing influence of repeated items arrangement is not expressed in reproduction error. Results from Experiment One make the latter possibility seem more likely. We, therefore, hypothesized that regression analysis, which was powerful for Experiment One, would be useful here for analyzing the effect of repeated items' arrangement. We carried out a hierarchical linear mixed effects regression analysis and evaluated the same three models that we used for Experiment One. To review, Model 1 assumed that responses reflected the array's arithmetic mean luminance, plus some random error. To this, Model 2 added the possibility that responses were influenced by the random initial luminance of the adjustable probe. Finally, Model 3 incorporated a separate term for the mean luminance of the repeated items.

Table 2 shows results from the comparison of the three regression models. For all the *Clustered* conditions, the best-fitting account was Model 3 that consisted of the three predictors DisplayMean, InitialProbe, and the RepeatedItems (lower AIC and BIC values and $p < .01$ in the nested χ^2 comparison with Model 2). This suggests that, in addition to the actual mean of the display and the random probe, the repeated items, when presented right next to each other, influence subjects' perception of the mean. For all the *Distributed* conditions, except 3rep, Model 2, which included just the DisplayMean and the InitialProbe, was the best-fitting model (lower AIC and BIC values and $p < .001$).

Although the comparison of the three regression models showed Model 3 to be the best-fitting model for the 3rep condition, AIC and BIC values for Models 2 and 3 were not

345 much different from one another. Additionally, Model 3 was significantly different from Model 2 at a value, $p=.02$, relatively high compared to the other significance levels shown in the table. Occam's razor tells us that when fewer predictors suffice, it is best to adopt the model with fewer predictors than a model with more. Therefore, it is safe to assume that for all our *Distributed* conditions, Model 2 was a better-fitting model. This suggests that repeated items
350 spread out across an array exert less influence on subjects' responses than they do when presented next to each other within an array.

Relative importances of the predictors

Although the best-fitting models informed us about the predictors that explain the total variance of the model, we do not fully understand the importance of each of the predictors in a
355 given model. To characterize the importances of each predictor, we calculated the relative importance metric using R's 'relaimpo' package (Grömping, 2006). In order to parse out the distinct influences of each predictor in the model, we used the 'lmg' metric (see Lindeman et al., 1980), which averages sequential sums of squares over all orderings of predictors. As in our analysis of Experiment One's results, we ranked the predictors in order of their importance
360 in a best-fitting model and characterized their share of R^2 in the model.

The results show that RepeatedItems influenced subjects' responses almost at the same level as DisplayMean in *Clustered* conditions. DisplayMean influenced performance significantly in the best-fitting model corresponding to the *Distributed* conditions (see Fig. 6 for a plot comparing the relative importance of predictors in the best-fitting models, in the 3rep *Clustered*
365 and *Distributed* conditions). These results were similar to the results from Experiment One and showed that the clustered repeated items exerted greater influence than the repeated items that were distributed throughout an array.

General Discussion

Experiment One investigated whether presenting the same items immediately adjacent
370 to one another would be disproportionately influential in determining the array's mean per-

Table 2

Summary of the comparison of the three regression models ($n=18$)

Models	Stimuli					
	Clump2rep	Dist2rep	Clump3rep	Dist3rep	Clump4rep	Dist4rep
Model 1						
AIC	9557.4	9324.9	9536.6	9442.7	9708.6	9464.5
BIC	9578.0	9345.4	9557.1	9463.2	9729.1	9485.0
Model 2						
AIC	9496.6	9250.8	9468.9	9332.4	9645.7	9424.2
BIC	9522.3	9276.4	9494.5	9358.0	9671.3	9449.8
Model 3						
AIC	9490.7	9252.5	9435.2	9329.0	9529.4	9415.1
BIC	9521.6	9283.2	9465.9	9359.7	9560.1	9445.9
<i>Best model estimates</i>						
χ^2	7.880	76.090	35.696	5.382	118.347	11.088
<i>p-value</i>	0.004	<2e-16	2.30e-09	0.02	<2e-16	0.0008

Note: Models 1, 2 and 3 were compared to determine the best-fitting model. All the highlighted values are the lowest AIC and BIC values which also represent the model with the best fit. Model 3 consisting of the three predictors, namely DisplayMean, InitialProbe, and RepeatedItems proved to be the best-fitting model for all the *Clustered* conditions whereas Model 2 consisting of two predictors, namely DisplayMean and InitialProbe proved to be the best fit model for the *Distributed* conditions. For *Distributed* 3rep condition, Model 3 was a slightly better model at a significance level of $p = 0.02$, which, compared to the significance levels of other best fit models, was substantially smaller. In addition, the AIC and BIC values for Model 2 and Model 3 were almost equal; therefore, we take this to mean that Model 2 with fewer predictors could be the better-fitting model for 3rep condition. Thus, this further suggests that repeated items presented as clumps exert greater influence on subjects' reproduction of the mean compared to when the repeated items are spread out.

Figure 6. Panel A. Relative importance of Model 3's predictors for 3rep *Clustered* stimuli. DisplayMean and RepeatedItems are the two most influential predictors that explain almost all of the variance explained by the model. **Panel B.** Relative importance of Model 2 predictors for 3rep *Distributed* stimuli. DisplayMean is the most influential predictor and explains almost all of the variance explained by Model 2. In both panels, error bars are 95% bootstrapped confidence intervals.

ceived luminance. We found that grouping repeated items near one another made them more influential than when they were distributed, and in fact gave them a perceptual influence that rivaled the influence of the entire array's actual mean. As the number of repeated items was constant, this result seemed to demand another study to understand fully whether varying the number of repeated items would have a comparable effect. Therefore, in Experiment Two, we delved deeper into the influence of presenting repeated items as clumps vs. spreading them throughout an array by varying the number of repeated items across arrays.

Subjects saw arrays of luminances with variable numbers —two, three or four— of repeated items. The repeated items within an array were either presented as clumps or separated from each other. As in Experiment One, at the end of each trial, a subject used a response dial to adjust an initial, random luminance so as to reproduce the average luminance of the array they had just seen. The analysis of these reproductions focused on reproduction error and on the relationship between performance and our independent variables.

Our main findings were, first, increasing the number of repeated items within an array produced more accurate responses (Fig. 5), and second, that repeated items had stronger

influence when they were near one another, potentially forming a perceptual group (Wagemans et al., 2012), than when they were spread out across the array (see Table 2). Moreover, repeated items presented near one another influenced subjects' judgments of the array's mean luminance about as much as the array's actual overall mean luminance did (Fig. 6).

390 These results further corroborate our previous findings from Experiment One.

Our first prediction was that subjects' accuracy in reproducing the array's mean luminance would increase as additional repeated items were inserted into the array. Within limits, the results confirmed this prediction: accuracy on arrays with four repeated items was significantly higher compared to the arrays with either two or three repeated items. However, 395 contrary to the prediction, we did not see a comparable increase in accuracy when the number of repeated items increased from two to three. One possible reason why accuracy increased for arrays with four repeated items was that with so many repeated items, there was an increased bias towards those repeated items. This suggests that some form of selective attention might be at play (e.g. Chong & Treisman, 2005a), which would cause the repeated items to have 400 particular influence in *ensemble statistics*. According to this view, when half the items are all the same, responses would tend to reflect those items. When those repeated items constituted such a substantial fraction of the entire array, subjects' biased perceptual judgments would reflect the array's actual mean luminance.

Our second prediction was that there would be a difference in accuracy at determining 405 the mean of an array between *Clustered* and *Distributed* arrays. The results of analyses of variance failed to support this prediction. Our results were similar to what we found in Experiment One, suggesting that subtle differences between *Clustered* and *Distributed* arrays might not be easily detectable with reproduction error, the measure we used to calculate accuracy. Indeed, taking nuisance variables and subject variability into account via hierarchical 410 regression, we were able to discern effects that were not detected by ANOVA. These regression results resemble the regression results of Experiment One: they showed small, but fairly consistent, differences, which suggests a greater role for repetitions when those repetitions were clustered rather than distributed within arrays.

Our third prediction was that with *Clustered* arrays but not *Distributed* arrays, accuracy
415 of reproduction would increase with the number of repeated items. However, we found no
such difference.

Our fourth prediction was that repeated items would have greater influence within
Clustered arrays than they would within *Distributed* arrays. Results from hierarchical multiple
regression supported that prediction: models that included repeated items gave the best fits
420 for *Clustered* arrays but not for *Distributed* arrays. This suggests that *ensemble statistics*
computations assign disproportionate weight to perceptually similar units that are grouped
together, perhaps because these arrangement causes attention to be drawn towards those
items (Chong & Treisman, 2005a; Brady & Alvarez, 2011; Brady & Tenenbaum, 2013). This
suggestion is consistent also with results from a very different task. Turk-Browne, Jungé, and
425 Scholl (2005) showed that signals that influence visual statistical learning, a form of learning
about sequential dependencies between objects, are gated by selective attention, even though
the task does not require such gating. Therefore in our study, it may be that when repeated
items are grouped together, the salience of that grouping give those items a disproportionate
influence on *ensemble statistics* computations. In contrast, when repeated items are spread
430 out, and are less salient, those repeated items would have had less impact on *ensemble*
statistics computation. So there are hints, from our own results and those of others, that
selective attention can play some role in ensemble perceptual judgments, but as yet we have
no direct evidence on this point.

Our study offers a new way of understanding *ensemble statistics* computation. Repetition
435 of items within an array increases the accuracy with which ensembles are reproduced, a fact
that might be understood as the brain's tendency to minimize the amount of information
that must be encoded. From an information-theoretic perspective, repetition of items within
an array, particularly when those items are grouped, would allow a sparse neural code to
suffice for the array (Barlow, 1961). This perspective is closely related to what is known as
440 "structural information theory" (Wagemans et al., 2012), and is broadly consistent with recent
evidence from a different task that individual items are weighted according to the likely fidelity

of those items' memory representations (Tong et al., 2019). If informational redundancy were an important variable, as our data suggest, then those processes that are sensitive to information per item, rather than number of items, would be implicated in the coding process.

445 For instance, it is well-known that visual short-term memory is information-limited rather than item-limited (Miller, 1956; Pollack, 1953), one feature that distinguishes it from the long-term store. The fact that reproduction error was reduced by redundancy (repetitions) suggests that the extraction or calculation of the mean luminance was sensitive to the amount of information in the display. Note that the brief duration of our stimuli (as short as 40 ms) and

450 the similarities of our results to those of Chong and Treisman (2003, 2005b), suggests that our subjects drew on an iconic store. Specifically, of the subcomponents of iconic memory (visual persistence, informational persistence, and afterimages), the demonstrated sensitivity to structural information suggests that informational persistence provides the raw material for the calculation of *ensemble statistics* in our task.

455 **Outstanding questions.** Although our study did succeed in uncovering previously unrecognized influences on ensemble perception, it left several questions unaddressed. Here, we consider two of those questions.

What stages of the visual system participate in *ensemble statistics*? The effect of the distribution of repeated items revealed here by hierarchical regression suggests that an emergent structure (e.g., light-dark boundary) guides *ensemble statistics*'s computation.

460 After all, stimulus structure does have a strong effect on later stages of processing, such as visual short-term and working memory, and stimulus structure is known to facilitate learning and memory for both verbal and non-verbal materials (Garner, 1974; Carterette & Friedman, 1978; Kintsch, 1970). Taken together with the foregoing arguments regarding informational

465 persistence and the information-limited character of VSTM, this suggests a process in which readout from the iconic store is held in VSTM and operated upon using control processes. This is important because subjects' memory-based matches to the mean luminance of a no longer visible stimulus array had to involve the contents of VSTM, and VSTM entails information-limited processing and sensitivity to structure and redundancy.

Brady, T. F., & Alvarez, G. A. (2011). Hierarchical encoding in visual working memory: ensemble statistics bias memory for individual items. *Psychological Science*, 22(3), 384-392. doi: 10.1177/0956797610397956

500 Brady, T. F., & Tenenbaum, J. B. (2013). A probabilistic model of visual working memory: Incorporating higher order regularities into working memory capacity estimates. *Psychological Review*, 120(1), 85-109. doi: 10.1037/a0030779

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10(4), 433-436.

Carterette, E., & Friedman, M. (1978). *Handbook of perception: Perceptual processing* (Vol. 9).
505 Academic Press.

Chong, S. C., & Treisman, A. (2003). Representation of statistical properties. *Vision Research*, 43, 393-404.

Chong, S. C., & Treisman, A. (2005a). Attentional spread in the statistical processing of visual displays. *Perception & Psychophysics*, 67(1), 1-13.

510 Chong, S. C., & Treisman, A. (2005b). Statistical processing: computing the average size in perceptual groups. *Vision Research*, 45(7), 891-900. doi: 10.1016/j.visres.2004.10.004

Cousineau, D. (2005). Confidence intervals in within-subject designs: a simpler solution to Loftus and Masson's method. *Tutorials in Quantitative Methods for Psychology*, 1(1), 42-45.

Dakin, S. C. (1997). The detection of structure in glass patterns: psychophysics and computational models. *Vision Research*, 37(16), 2227-2246.
515

de Fockert, J. W., & Marchant, A. P. (2008). Attention modulates set representation by statistical properties. *Perception & Psychophysics*, 70(5), 789-794.

Dubé, C. (2019). Central tendency representation and exemplar matching in visual short-term memory. *Memory & Cognition*, 47(4), 589-602.

520 Dubé, C., & Sekuler, R. (2015). Obligatory and adaptive averaging in visual short-term memory. *Journal of Vision*, 15(3), 1-13.

Garner, W. R. (1974). *The processing of information and structure*. Mahwah, New Jersey: Lawrence Erlbaum and Associates.

Glass, L. (1969). Moiré effect from random dots. *Nature*, 223(5206), 578-580. doi:

525 10.1038/223578a0

Grömping, U. (2006). Relative Importance for linear regression in R: The Package relaimpo. *Journal of Statistical Software*, 17(1), 1-27.

Haberman, J., & Whitney, D. (2007). Rapid extraction of mean emotion and gender from sets of faces. *Current Biology*, 17(17), 751-753. doi: 10.1016/j.cub.2007.06.039

530 Haberman, J., & Whitney, D. (2009). Seeing the mean: ensemble coding for sets of faces. *Journal of Experimental Psychology: Human Perception and Performance*, 35(3), 718-734. doi: 10.1037/a0013899

Haberman, J., & Whitney, D. (2010). The visual system discounts emotional deviants when extracting average expression. *Attention, Perception & Psychophysics*, 72(7), 1825-1838. doi: 10.3758/APP.72.7.1825

535

Hebb, D. O. (1961). Distinctive features of learning in the higher animal. In J. Delafresnaye (Ed.), *Brain mechanisms and learning* (p. 37-46). Oxford: Blackwell.

Hubert-Wallander, B., & Boynton, G. M. (2015). Not all summary statistics are made equal: Evidence from extracting summaries across time. *Journal of Vision*, 15(4), 1-12. doi: 10.1167/15.4.5

540

Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience*, 2(3), 194-203. doi: 10.1038/35058500

Kahana, M. J., & Sekuler, R. (2002, Aug). Recognizing spatial patterns: a noisy exemplar approach. *Vision Research*, 42(18), 2177-2192.

545 Kimble, D. P. (1967). *The anatomy of memory: Learning, remembering, and forgetting* (Vol. 2). Science and Behavior Books.

Kintsch, W. (1970). *Learning, memory, and conceptual processes*. John Wiley & Sons.

Knill, D. C., & Pouget, A. (2004). The bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neuroscience*, 27(12), 712-719. doi: 10.1016/j.tins.2004.10.007

550 Lindeman, R. H., Merenda, P. F., & Gold, R. Z. (1980). *Introduction to bivariate and multivariate analysis*. Glenview, IL: Scott, Foresman.

Maule, J., & Franklin, A. (2015). Effects of ensemble complexity and perceptual similarity on

rapid averaging of hue. *Journal of Vision*, 15(4), 1-18. doi: 10.1167/15.4.6

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity
555 for processing information. *Psychological review*, 63(2), 81.

Myczek, K., & Simons, D. J. (2008). Better than average: alternatives to statistical summary
representations for rapid judgments of average size. *Perception & Psychophysics*, 70(5),
772-788.

Pollack, I. (1953). Assimilation of sequentially encoded information. *The American journal of*
560 *psychology*, 66(3), 421--435.

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science.
Science, 237(4820), 1317--1323.

Tong, K., Dubé, C., & Sekuler, R. (2019). What makes a prototype a prototype? averaging visual
features in a sequence. *Attention, Perception, & Psychophysics*, 81(6), 1962--1978.

565 Turk-Browne, N. B., Jungé, J., & Scholl, B. J. (2005). The automaticity of visual statistical learning.
Journal of Experimental Psychology: General, 134(4), 552-564. doi: 10.1037/0096-
3445.134.4.552

Wagemans, J., Feldman, J., Gepshtein, S., Kimchi, R., Pomerantz, J. R., van der Helm, P. A.,
& van Leeuwen, C. (2012). A century of Gestalt psychology in visual perception: II.
570 Conceptual and theoretical foundations. *Psychological Bulletin*, 138(6), 1218-1252. doi:
10.1037/a0029334

Williams, D. W., & Sekuler, R. (1984). Coherent global motion percepts from stochastic local
motions. *Vision Research*, 24(1), 55-62.

Woutersen, K., van den Berg, A. V., Boonstra, F. N., Theelen, T., & Goossens, J. (2018). Useful
575 field of view test performance throughout adulthood in subjects without ocular disorders.
PLoS One, 13(5), e0196534. doi: 10.1371/journal.pone.0196534