Seeing Both the Forest and the Trees:

Distinct Resolution of Hierarchical Representations in Visual Working Memory

Tina T. Liu^{1,2*}, Zhongting Chen^{3*}, Matt Oxner^{1,4}, Wanying Wang³, Yixuan Ku^{5,6}, William G. Hayward¹

¹Department of Psychology, University of Hong Kong, Hong Kong, China

²Laboratory of Brain and Cognition, National Institutes of Health, Bethesda, United States

³School of Psychology and Cognitive Science, East China Normal University, Shanghai, China

⁴School of Psychology, Victoria University of Wellington, Wellington, New Zealand

⁵Guangdong Provincial Key Laboratory of Social Cognitive Neuroscience and Mental Health,
Department of Psychology, Sun Yat-Sen University, Guangzhou, China

⁶Peng Cheng Laboratory, Shenzhen, China

*Co-first authors with equal contribution

Correspondence:

William G. Hayward (whayward@hku.hk) or Yixuan Ku (kuyixuan@gmail.com)

Abstract

The capacity limit in visual working memory (VWM) has been well established, but the way in which it is limited remains under debate. For example, an item limit was proposed by the discrete slot models whereas a flexible allocation of information was suggested by the flexible resource models. Despite the discrepancies on the nature of VWM capacity limit, both models assume that items are processed independently. Recent evidence, however, revealed that VWM is hierarchically structured. Given that an item can be represented both at the ensemble and the item level, existing measures of VWM capacity cannot distinguish between them. To separately measure an item at both ensemble and item levels, we designed a novel paradigm by combining the continuous report paradigm with a categorical judgement task. In the present series of four experiments, we found that participants' performance at the ensemble and the item level was differentially affected by the layout and the number of objects in the display. Similarly, hierarchical Bayesian models fit to the empirical data suggest distinct effects of layout and set size on ensemble and item representations, respectively. Moreover, separating ensemble from item representations further allowed us to address their relationship and we uncovered a surprising possibility of independence between these two levels of representations. Taken together, the present dissociation and the demonstration of independence suggest distinct mechanisms of item and ensemble processing, which further sheds light on the nature and structure of representations in VWM.

Keywords: visual working memory, ensemble representation, memory precision, Bayesian hierarchical modeling

Introduction

Visual working memory (VWM) allows us to hold visual information actively in mind over a short-term delay and is central to many cognitive processes that involve vision and memory (Baddeley, 2007). One of the most distinctive characteristics of VWM is its limited storage capacity with the consequence that only a subset of information is selected for storage (Cowan, 2001, 2010; Pashler, 1988; Phillips, 1974). This capacity limit, however, opens a window to elucidate the properties of VWM and, as such, has been studied extensively over the past two decades (for a review, see Brady, Konkle, & Alvarez, 2011).

One of the critical debates on VWM capacity concerns the number and precision of objects that can be stored. According to the discrete slot models (Luck & Vogel, 1997; Zhang & Luck, 2008, 2011), only a limited number of items can be stored in a fixed number of slots in VWM with fixed precision in each slot. Once the item limit is reached, no further information can be stored from additional items. In contrast, the flexible resource models (Alvarez & Cavanagh, 2004; Alvarez & Cavanagh, 2010; Bays & Husain, 2008; Bays, Catalao, & Husain, 2009; Huang, 2010) propose no upper bound on the number of items that can be stored. According to this view, VWM capacity is limited in terms of the total amount of information that can be flexibly allocated among the items to be stored.

The discrete slot models and the flexible resource models are both well supported by empirical evidence (for reviews, see Luck & Vogel, 2013; Ma, Husain, & Bays, 2014). For example, typical results from the change detection task, in which participants are required to indicate whether two arrays viewed over a short delay are identical or not, have robustly demonstrated an item limit (about four) in VWM (Luck & Vogel, 1997). In contrast, object complexity was found to modulate the number of items stored in VWM. For instance, Alvarez

and Cavanagh (2004) showed that storage capacity was substantially lower for complex objects (e.g., shaded cubes) compared to simple objects (e.g., colored squares). This implies a tradeoff between number and complexity of objects.

Interestingly, evidence supporting both models can be found even within the same task. For example, Awh, Barton, & Vogel (2007) conducted a change detection task and showed that participants' performance in the within-category-change conditions (e.g., from a cube to another cube) varied as a function of object complexity, suggesting a limited but flexible resource rather than a fixed number of objects in VWM (Alvarez & Cavanagh, 2004). In contrast, capacity estimates for the cross-category-change conditions (e.g., from a cube to a Chinese character) were in line with Luck and Vogel (1997) that VWM holds a fixed number of objects regardless of object complexity. These apparently contradictory conclusions suggest that neither discrete slot models nor flexible resource models suffice to characterize the nature of VWM capacity limit.

Despite the extensive research on the number and precision of individual items stored, recent studies suggest that VWM representations have structures beyond the item level (e.g., Brady & Alvarez, 2011, 2015). Given that the structure of the visual world is inherently hierarchical (e.g., trees in a forest) and highly redundant (e.g., similar shape and color of leaves on a tree), Ariely (2001) proposes the visual system represents not only individual features but also overall statistical properties when sets of similar objects are present. In addition to the feature dimension, Jiang, Olson, and Chun (2000) showed that an item is not represented independently of its spatial context in working memory, uncovering the spatial ensemble dimension (for recent work, see Schurgin & Brady, 2019; Vidal, Gauchou, Tallon-Baudry, & O'Regan, 2005). Brady and Alvarez (2011) further argue that VWM encodes both properties of

the individual objects (e.g., color, orientation, size) and properties of the global scene, or “ensemble statistics” (e.g., mean color, mean orientation, mean size, and the relational information between items).

Based on the hierarchically structured memory framework (Brady & Alvarez, 2011, 2015) and the empirical work that extended the Awh et al. (2007) study to the spatial ensemble dimension (Schurgin & Brady, 2019), we reasoned that different capacity estimates of within-category and cross-category change in Awh et al. (2007) reflect different levels of representation in VWM. At the local level, each item needs to be stored at high resolution to allow for detailed comparisons between the test and the sample array for successful within-category change detection. At the global level, a low-resolution spatial ensemble representation of the whole display would suffice for detecting cross-category change(s). Given that participants could not anticipate which type of change occurred on each trial, storing both item and ensemble representations would possibly facilitate performance in the VWM change detection task.

Despite recent compelling evidence demonstrating that VWM stores both individual items and ensemble statistics (Nie, Müller, & Conci, 2017; Xie & Zhang, 2017), the exact relationship between the two levels of representation has not been addressed. One possibility is that both representations share a common pool of VWM resources, such that manipulating the number of individual items to be stored should affect the precision of the ensemble representation. Thus, remembering more individual objects would result in a worse performance in recalling the ensemble information. Alternatively, each level of representation can be independently processed or stored. As a result, measures of item representations would not be influenced by manipulations at the ensemble level (and vice versa).

In a series of four experiments, we systematically tested predictions from these two

competing hypotheses. To provide separate measures of the two types of VWM representations within each trial, we combined the categorical judgment task adapted from Awh et al. (2007) with the standard continuous report task (cf. Wilken & Ma, 2004). Take Experiment 1 as an example, in the categorical judgment task, participants reports the category of the probed item (e.g. cube or colored squares) which serves as a measure for low-resolution spatial ensemble information. Next, depending on the category chosen, the continuous report task measures the high-resolution item information (i.e., object identity) of this very same item-to-report. To foreshow the key results, we demonstrated an important dissociation between ensemble and item representations. Specifically, the precision of individual item representation, but not ensemble representation, decreased with the number of objects in the display (Experiments 1-3). In contrast, ensemble representation, but not individual item representation, was sensitive to different layouts of objects at the global level (Experiment 4). Converging evidence using Bayesian hierarchical modeling of ensemble and item representations further suggests distinct effects of layout and set size on the ensemble and the item representation, respectively.

Experiment 1

Experiment 1 was a conceptual replication and a validation of ensemble measure based on Awh et al. (2007), in which they outlined a two-factor model suggesting that number and resolution can be separate measures of working memory ability. Based on Schurgin and Brady (2019) which extended the Awh et al. (2007) study to the spatial ensemble dimension, we reasoned that different levels of representations might be simultaneously involved in this task, consisting of a low-resolution ensemble representation of all the items in the display which sufficed for detecting cross-category changes, and a high-resolution representation of the probed

item (if present in memory) to detect a within-category change. To test this hierarchical processing hypothesis, we asked participants to make two responses on every trial and made the nature of the second contingent on the first. First, participants had to make a categorical judgment, which allowed us to measure memory precision of ensemble representations at each set size. This was followed by a within-category discrimination task, to measure memory precision of an item in the display across different set sizes (cf. Wilken & Ma, 2004). Separating measures of general category membership and high-resolution identity thus allows us to examine how set size differentially affects ensemble and item representations in VWM. In addition, it serves as an important purpose to validate the design of experiment: If one could perform categorical judgment using individual item memory, then manipulating the set size in Experiments 1-3 (or layout in Experiment 4) should equally affect both levels of representations.

Method

Participants

Thirteen students from the University of Hong Kong participated in the study for course credit. All reported normal or corrected-to-normal visual acuity, and normal color vision. All provided informed consent to participate in the experiment approved by the Institutional Review Board of the University of Hong Kong. Two participants' performances were below 50% chance level in the categorical judgment task and were therefore discarded from further analysis.

Apparatus & Stimuli

Experiment was programmed using MATLAB and Psychophysics Toolbox extensions (Brainard, 1997). Stimuli were displayed on a 28-inch LCD monitor with a resolution of 1680×1050 pixels and a refresh rate of 60 Hz. Stimuli included both colored squares generated

online using MemToolbox (memtoolbox.org) and shaded cubes adapted from Alvarez and Cavanagh (2004). On each display, half were colored squares and the other half were cubes, and participants were asked to remember all of them. See Figure 1 for examples of the stimuli. Participants viewed the display from a distance of approximately 55 cm, therefore each item in the display subtended about $2.5^\circ \times 2.5^\circ$ of visual angle and was centered on an invisible circle with a radius of 6.5° . The positions were pseudo-randomly chosen from a set of eight locations equally spaced along the circle. Depending on participants' categorical judgment, either a 2° thick continuous color wheel containing 360 hues (radius: 9.7°), or an array of six cubes centered on an invisible circle (radius: 9.5°), was presented. In the color wheel display, the cued location changed its color according to the mouse position on the color wheel. In the cube array display, the background of the selected cube was highlighted in red.

Design

This study had one independent variable: set size (two, four, and six). The dependent variables were number accuracy, categorical accuracy, color error, and cube accuracy, at each set size. The operational definition of precision for ensemble representations was categorical accuracy, quantified by the percentage of trials that the categorical judgment was correct. The precision of item representations was quantified by color error and cube discrimination accuracy: a) color error was measured as the distance between the response and the target value in degrees on the color wheel. The color wheel contained 360 possible values, so the range of color error was between -180 and 180 ; and b) cube discrimination accuracy was measured as the percentage of trials in which the cube was correctly identified out of six cube choices. Color error and cube discrimination accuracy were calculated over only trials in which participants indicated the correct category.

Procedure

The sequence of displays in a single trial was illustrated in Figure 1. To prevent verbal encoding of the stimuli, two different digits (from 1 to 9) were randomly picked and presented aurally at the beginning of each trial. Participants were instructed to rehearse them (either silently or read-out-loud) throughout the trial. At the same time, a solid black dot was presented at the center of the white screen for 500ms. This was followed by a memory array, consisting of two, four, or six items staying on the screen for 0.5s, 1s, and 1.5s respectively. Half were shaded cubes and half colored squares. After a 1s retention interval, all the items were masked by gray squares, with one square darker than the others. This cue indicated the item to report and the location of the cue was chosen at random. The response consisted of two stages. In the first stage, the participant reported the category of the item-to-report by pressing the corresponding key (1 for “color” and 2 for “cube”). If a “color” response was made, a color wheel was shown, and the participant was instructed to match the color of the cued item by moving the mouse and clicking on the color wheel. If a cube response was made, a display of six cubes was shown and participants were asked to select the item that they had just seen. After that, two digits appeared centrally on the screen and participants were asked to judge whether they were the same or different as those they had been rehearsing throughout the trial. They were also aware that the auditory task was secondary to the primary visual task. The experiment comprised six blocks of 50 trials each, so there were 300 trials in total and the whole experiment took around 45 min to finish. Each participant finished a practice session of 24 trials prior to the experiment. Practice data were checked and in very rare cases when participants did not follow the instruction, they were asked to do one more practice session before proceeding to the experiment.

Figure 1. An illustration of trial sequence in Experiment 1.

Results and Discussion

In this experiment, we analyzed the response from each task using one-way repeated-measures Analysis of Variance (ANOVA) with set size (two, four, and six) as the within-subject factor. First, we performed a manipulation check of the articulatory suppression task. The purpose of this secondary task was to minimize verbal encoding and rehearsal of the visual display (Baddeley, 1986). Average number accuracy was very high ($97.47\% \pm 2.70\%$) and did not differ significantly across set sizes, $F(2, 20) = .875, p = .432$, partial $\eta^2 = .080$, suggesting that participants had effectively encoded and maintained the digits using verbal working

memory. Next, we examined responses from the visual task, which revealed a critical finding of differential set size effect on performance in the two responses. Figure 2 shows results from the categorical judgment task (measure of ensemble representation) and the within-category discrimination task (measure of item representation).

Figure 2. Results from Experiment 1. (a) Mean accuracy at each set size in the categorical judgment task (measure of ensemble representation). (b1) Color error at each set size in the within-category discrimination task (measure of item representation). (b2) Cube accuracy at each set size in the within-category discrimination task (measure of item representation). Error bars show ± 1 standard error of mean.

Ensemble representation

In the first categorical task, no significant main effect of set size was present, $F(2, 20) = 2.967$, $p = .074$, partial $\eta^2 = .229$. Categorical accuracy was 98.4%, 96.4%, and 97.8% at set size

2, 4, and 6, respectively. This absence of set size effect on ensemble representations revealed that memory precision at this level of representation was robust to changes in the number of items in the display (up to six items in this case).

Item representation

In the within-category discrimination, both color error and cube discrimination accuracy showed a main effect of set size, color error: $F(2, 20) = 35.979, p < .001$, partial $\eta^2 = .783$; cube discrimination accuracy: $F(2, 20) = 35.895, p < .001$, partial $\eta^2 = .782$. A follow-up paired-samples t-test revealed that there was a significant performance drop between any two adjacent set size conditions (color, set size 2 vs. 4: $t(10) = 4.758, p = .001$, *Cohen's d* = 1.504; color, 4 vs. 6, $t(10) = 4.929, p = .001$, *Cohen's d* = 1.559; cube, set size 2 vs. 4: $t(10) = 4.624, p = .001$, *Cohen's d* = 1.462 ; cube, set size 4 vs. 6: $t(10) = 4.891, p = .001$, *Cohen's d* = 1.547).

To summarize, Experiment 1 set out to examine how set size differentially affects ensemble and item representations. In line with the extensive literature on VWM precision, we found that precision of item representations decreased significantly with set size. Moreover, we showed that the precision of ensemble representations remained largely constant across all set sizes. Admittedly, the effect of set size in the categorical judgement task was approaching significance, which was contributed by set size 4 being lower than set sizes 2 and 6. This pattern, however, seems to be qualitatively different from the linear set size effect observed in the within-category discrimination task. More importantly, in our paradigm, both categorical judgment and within-category discrimination were based on the very same item, which enabled us to validate the design and tentatively arrive at the conclusion that VWM representation is inherently hierarchical: the representations at the ensemble level and at the item level are differently processed and stored for the same individual object.

Certainly, performance in the categorical judgment reached ceiling, such that any true drop in performance might be difficult to detect. This ceiling effect might be due to the ease of perceptual grouping in the current display, for example, all the color items may be perceptually grouped into a global representation of color set, so that simply remembering the locations of all the colors sufficed for making a categorical judgment (i.e., the rest were cubes). In Experiment 2, we attempt to minimize perceptual grouping by using compound Gabor patches (each had a color and an orientation dimension) to make the objects of two categories physically similar.

Experiment 2

In Experiment 1, color squares and shaded cubes were chosen as the two object categories to assess ensemble and item representations. These two types of stimuli are commonly used object categories in VWM research. They may not be ideal, however, for the present analysis for two reasons. First, physically dissimilar objects (such as color squares and shaded cubes) might be easy to group and categorize, which in turn, facilitates performance in the categorical task (Olsson & Poom, 2005). Second, the cube discrimination response (i.e., selecting the memorized one out of six cube choices) was not on a continuous dimension, which made it difficult to model the response data for further analysis. To address these issues, we used compound Gabor patches varying in two dimensions: color and orientation. In the categorical judgment task, we asked participants to respond which category (“color” or “orientation”) the probed Gabor belonged to. Moreover, the continuous nature of the color and orientation dimensions allowed us to model the continuous report task and fit to the standard mixture model of Zhang and Luck (2008).

The rationale of the Zhang and Luck (2008) model is that, if the probed item is present in VWM, responses should be near the correct color, with errors normally distributed around the correct color. In contrast, if the probed item is absent from memory, the responses should be drawn from a uniform distribution of errors, reflecting guessing. When the observer remembers the probed item on some trials and guesses randomly on other trials, the overall distribution of errors should be a mixture of a von Mises distribution¹ of errors (when probed item is present) and a uniform distribution of errors (when the probed item is absent). Maximum likelihood estimation is applied to decompose the data from each participant in each set size and two parameters are essential for the present analysis: for the errors from the uniform distribution of errors, P_m represents the probability that the probed item is present in memory (i.e., measure of item limit); in the von Mises distribution of errors, s.d. (standard deviation) reflects the width of the distribution of errors on trials when the probed item is present in memory, thus denoting the precision of the representation in VWM (measure of item precision).

Method

Participants

Sixteen students from the University of Hong Kong participated in the study for course credit. All reported normal or corrected-to-normal visual acuity, and normal color vision. Informed consent was obtained prior to the start of the experiment. Two participants' performance was below 50% chance level in the categorical judgment task and was therefore discarded from further analysis. One additional participant was excluded due to a strong bias for

¹ In the case of reporting color and orientation, the von Mises distribution is used instead of the normal distribution (or Gaussian distribution) as the both color space and orientation space is circular.

responding mostly to color items that resulted in a lack of sufficient orientation trials for analysis.

Apparatus & Stimuli

The apparatus and stimuli were identical to those in Experiment 1 except for the following changes: Stimuli were Gabor patches that each had a color and an orientation dimension (Figure 3), presented on a black background. The variable property denoted the category of the Gabor: half the items were color Gabors which had a standard orientation (vertical) but a variable color (Figure 3a), and the other half were orientation Gabors which had a standard color (red center with green surround) but variable orientation (Figure 3b).

Figure 3. An illustration of the compound “color” and “orientation” Gabor stimuli used in Experiment 2. The category of the Gabor refers to its variable property. Panel (a) shows a “color” Gabor as its color varies among all 360 possible values (from -180 to 180) on the color wheel, but its orientation is fixed (always vertical). Panel (b) shows an “orientation” Gabor, the orientation of which spans from -90 to 90, but its color is fixed (red center with green surrounds).

Design

This study had one independent variable: set size (two, four, and six). The dependent variables were categorical accuracy (measure of ensemble representation), P_m and s.d. (both were measures of item representation), and number accuracy, at each set size.

Procedure

The procedure of each trial in Experiment 2 is shown in Figure 4. Participants first reported the variable property (i.e. category) of the probed item, and then the specific value of that property. For example, a red-green Gabor oriented at the 30° would be first identified as “orientation”, then matched with the orientation wheel as close to 30° as possible. The experiment consisted of six blocks of 40 trials in each block, so there were 240 trials in total and the whole experiment took around 40 min to finish. Each participant finished a practice session of 24 trials prior to the experiment. Practice data were checked and in rare cases when participants did not follow the instruction, they were asked to repeat the practice session before proceeding to the experiment.

Model

The standard mixture model (“slot+averaging”) of Zhang and Luck (2008) was applied to model color and orientation trials using the MemToolbox (memtoolbox.org). By default, a Markov Chain Monte Carlo (MCMC) algorithm was used for model fitting and the maximum a posteriori (MAP) estimate was obtained for each parameters (g, s.d.), for each participant, at each set size (for more information, see Suchow, Brady, Fournie, & Alvarez, 2013). Note that the parameter g stands for guessing rate, and $P_m = 1 - g$.

Figure 4. Trial sequence in Experiment 2. The procedure² is identical to that of Experiment 1 except the substitution of color Gabor patches for color squares and shaded cubes.

Results & Discussion

Similar to Experiment 1, performance in the articulatory suppression task (i.e., number accuracy) was uniformly high ($97.37\% \pm 1.92\%$) across all set sizes in this experiment. Again, no effect of set size on number accuracy was found, $F(2, 24) = .333$, $p = .719$, partial $\eta^2 = .027$, suggesting no major contribution from verbal rehearsal of visual stimuli. Figure 5 shows results

² In the instruction, the "orientation" condition was shown as "angle" to minimize the character length on the display.

from the categorical judgment task in the left panel (a), and the continuous report task, in the right panels (b1 and b2).

Figure 5. Results from Experiment 2. (a) Average accuracy at each set size of the categorical judgment task (measure of ensemble representation). (b1) P_m for both color and orientation (measure of item representation) declined relatively slowly from set size 2 to 4, compared to the larger drop from 4 to 6. (b2) s.d. for both color and orientation (measure of item representation) increased as set sizes increased but remained flat after four items. Error bars show ± 1 standard error of mean.

Ensemble representation

Categorical accuracy was 98.17%, 97.98%, and 96.35% at set size 2, 4, and 6, respectively. Replicating findings from Experiment 1, no significant drop in categorical judgment performance was found, $F(1.293, 15.513) = 2.255$, $p = .127$, partial $\eta^2 = .158$, Greenhouse-Geisser corrections, suggesting that memory precision at the ensemble level was robust to changes in set sizes (two, four, or six items).

Item representation

Repeated-measures ANOVAs were run separately for P_m and s.d. (two key parameters in the standard mixture model), for color and orientation trials, respectively. The main effects of set size on P_m were significant in both color and orientation trials (*color* P_m : $F(1.223, 14.681) = 15.310$, $p = .001$, partial $\eta^2 = .561$, Greenhouse-Geisser corrections; *orientation* P_m : $F(2, 24) = 16.698$, $p < .001$, partial $\eta^2 = .582$). As shown in panel b1 of Figure 5, P_m for both color and orientation declined relatively slowly from set size 2 to 4 [*color* P_m : $t(12) = 1.913$, $p = .080$, *Cohen's d* = .552; *orientation* P_m : $t(12) = 2.063$, $p = .061$, *Cohen's d* = .596], compared to a significant decline from 4 to 6 [*color* P_m : $t(12) = 3.861$, $p = .002$, *Cohen's d* = 1.115; *orientation* P_m : $t(12) = 4.124$, $p = .001$, *Cohen's d* = 1.191]. Together, color and orientation P_m converged to show that the probability of an item present in memory decreases suddenly between set size 4 and 6, suggesting an item limit (about 4 items) in VWM.

The other key parameter from the Zhang and Luck (2008) model, s.d., was plotted separately for color and orientation in the panel b2 of Figure 5. The set size effect on s.d. was significant for both color [$F(1.122, 13.462) = 5.657$, $p = .030$, partial $\eta^2 = .320$, Greenhouse-Geisser corrections] and orientation ($F(2, 24) = 18.541$, $p < .001$, partial $\eta^2 = .607$). Pairwise comparisons showed that s.d. for both color and orientation increased significantly from set size two to four [color s.d.: $t(12) = 7.220$, $p < .001$, *Cohen's d* = 2.084, orientation s.d.: $t(12) = 4.592$, $p = .001$, *Cohen's d* = 1.326] but not after set size reached four items [color s.d.: $t(12) = .591$, $p = .565$, *Cohen's d* = 0.171, orientation s.d.: $t(12) = 1.745$, $p = .107$, *Cohen's d* = 0.504]. Given that s.d. denotes item precision in VWM, the monotonic increase in color and orientation s.d. until set size reached four is consistent with the item limit of four inferred from the color and orientation P_m .

The present experiment replicated and extended the findings from Experiment 1. In particular, the current statistical modeling approach allowed us to systematically examine how precision of ensemble and item representations change with varying set size. Same as in Experiment 1, we found that the precision of ensemble representations was not affected by the present set size manipulation, while precision of item representations plateaued after set size reached four, indicating an item limit in VWM (Zhang & Luck, 2008, 2011). This differential set size effect on the precision of representations implied that participants stored both low-resolution information about the entire display and high-resolution representations of a subset of items in the display. Therefore, our results suggest distinct ensemble and item representations, which are in line with the hierarchically structured memory account (Brady & Alvarez, 2015).

Experiment 3

In Experiments 1 and 2, the differential set size effects on ensemble vs. item representations allowed us to arrive at the tentative conclusion that an item in VWM is simultaneously represented at both the ensemble and the item levels. A potential concern, however, was that the task demand (e.g., category-then-item report procedure) may have encouraged observers to prioritize category encoding. In addition, the serial nature of the task may have introduced memory decay as another potential confound, given that the categorical judgment task always preceded the continuous report task. In this study, we address these concerns by comparing the within-subject VWM measures with and without the categorical judgment task. Our reasoning is that, if the presence of the first categorical judgment task led to 1) a prioritization of the ensemble task and/or 2) a decay of item representations in VWM, then we would expect to see better item measures without performing the categorical judgment task.

In contrast, if ensemble and individual representations were processed simultaneously in VWM, performing the prior categorical task should not affect item measures in the second continuous report task.

The no-categorical-judgment condition of the present experiment further allows us to examine whether the mere presence of object category can affect measures of VWM. Although distinct color squares and shaded cubes (Experiment 1) were changed to physically similar color and orientation Gabors (Experiment 2) to minimize the category boundary, we still do not know if item measures would be sensitive to the presence of object category in the display. To this end, we split the no-categorical-judgment condition into two sub-conditions, one-category-only (color-only or orientation-only trials) and both-category-present (half-color and half-orientation trials, same as Experiment 2). Again, note that the category boundary here refers strictly to the feature of stimuli since no categorical judgment task (i.e., task demand) is involved. If categorical representation is automatically perceived and shares the common pool of VWM resources with item representations, we would expect worse performance at the item level in the both-category-present condition (due to the mere presence of two categories) than in the one-category-only condition.

In many aspects, Experiment 3 serves as an important control experiment in which the critical comparisons are threefold. First, we compare measures of item representations with versus without performing the categorical judgment task to rule out the possibility that either task demand or memory decay explained the differential set size effect on ensemble vs. item representation found in Experiments 1-2. Second, within the no-categorical-judgment condition, we compare item performance when only one object category is present versus both categories are present to examine whether categorical representations share the common pool of VWM

resources with item representations. Lastly, we increase the number of items in the display in this experiment to examine if the previous differential set size effects at ensemble and item representations can be generalized to a larger set size manipulation.

Method

Participants

14 students from East China Normal University participated in the study for monetary compensation. All reported normal or corrected-to-normal visual acuity, and normal color vision. All provided informed consent to participate in the experiment approved by the Institutional Review Board of East China Normal University.

Apparatus & Stimuli

The stimuli were identical to those in Experiment 2, except they were displayed on a 27 inch monitor with the same resolution of 1920×1080 pixels and a refresh rate of 60 Hz.

Design

At the item level, there were three within-subject variables in this study: set size (four, six, and eight), task condition ('color-only', 'orientation-only', 'color and orientation', 'color and orientation with categorical judgment') and trial type (color, orientation). The dependent variables (P_m and s.d.) were the same as those tested in Experiment 2. At the ensemble level, the independent variable was set size (four, six, and eight) and the dependent variable was categorical judgement accuracy. In addition, given the null set size effect of number accuracy in both Experiments 1 and 2 and the three additional task conditions included here, the secondary articulatory suppression task was removed in the present experiment to increase the number of trials per condition.

Procedure

The general procedure in this experiment was similar to that in Experiments 1 and 2 except the exclusion of the secondary articulatory suppression task. In addition, the presence of the categorical judgment task only applied to the ‘color and orientation with categorical judgment’ sub-condition. The one-category-only conditions consisted of either all color or all orientation trials and participants were required to report the property of the probed stimulus using the corresponding color wheel (‘color-only’) or orientation wheel (‘orientation-only’). In the ‘both-category-present’ conditions, a display of half-color and half-orientation items was either followed by a categorical judgment task (‘color and orientation with categorical judgment’) or directly proceeded to the corresponding continuous report task (‘color and orientation’). Moreover, to make the timing of the two ‘both-category-present’ conditions comparable, an additional randomly varied interval between 500ms and 1000ms was added to the 1s retention interval in the ‘color and orientation’ condition.

Each participant performed all four task conditions within the same session. For all participants, two one-category-only tasks were run first, which were followed by two both-category-present tasks. The orders of the two one-category-only tasks and the two both-category-present tasks were counterbalanced across participants, respectively. Each one-category-only block included 75 trials and each both-category-present block included 150 trials. Each participant completed 450 trials, blocked by four task conditions. Each participant finished a practice session of 12 trials prior to the start of each task block. The whole experiment took around one hour to complete.

Model

Same as Experiment 2, the standard mixture model (“slot+averaging”) of Zhang and Luck (2008) was applied to model color and orientation trials using the MemToolbox (memtoolbox.org). By default, a Markov Chain Monte Carlo (MCMC) algorithm was used for model fitting and the maximum a posteriori (MAP) estimate was obtained for each parameters (g, s.d.), for each participant, at each set size (for more information, see Suchow, Brady, Fougner, & Alvarez, 2013). Note that $P_m = 1 - g$.

Results and Discussion

Same as in Experiment 2, statistical modeling of color/orientation trials was performed using the MemToolbox (memtoolbox.org). P_m and s.d. were fitted using the slot+averaging model (Zhang & Luck, 2008) for each condition at each set size. This ensured that item representations can be directly compared across the ensemble-then-item report condition (‘color and orientation with categorical judgment’) and three individual-item-only conditions (‘color-only’, ‘orientation-only’, ‘color and orientation’).

Ensemble representation

Categorical accuracy was 96.43%, 95.43%, and 93.86% at set size 4, 6, and 8, respectively. As observed in the previous two experiments, no significant drop in categorical judgment performance was found, $F(2, 26) = 2.656$, $p = .089$, partial $\eta^2 = .170$. This finding replicated the previous observation that memory precision of at the ensemble level was insensitive to set sizes change, which can now be generalized to larger set sizes.

Figure 6. Results from Experiment 3. (a) Mean accuracy in the categorical judgment task (measure of ensemble representation). (b1) Color s.d. and color P_m plotted separated for color-only, color and orientation, and color and orientation with categorical judgment (CJ) conditions (measures of item representation). (b2) Orientation s.d. and orientation P_m plotted separately for orientation-only, color and orientation, and color and orientation with CJ. Error bars show ± 1 standard error of mean.

Item representation

For each trial type (color or orientation), we first carried out an omnibus repeated-measures ANOVA to test the overall hypothesis in this experiment: whether individual-item representations at each set size (four, six, and eight) are affected by task demand (e.g., the serial

nature of the category-then-item report procedure). Specifically, we examined 1) color s.d. and P_m under task conditions of ‘color-only’, ‘color and orientation’, ‘color and orientation with categorical judgment’ with varying set size for color trials, and 2) orientation s.d. and P_m under task conditions of ‘orientation-only’, ‘color and orientation’, ‘color and orientation with categorical judgment’ at each set size for orientation trials.

Relevant to this overall hypothesis, we found no significant main effect of task demand on either s.d. or P_m for color-item representations (s.d.: $F(2, 26) = .301, p = .743$, partial $\eta^2 = .023$; P_m : $F(2, 26) = 3.009, p = .067$, partial $\eta^2 = .188$) or orientation-item representations (s.d.: $F(2, 26) = .139, p = .871$, partial $\eta^2 = .007$; P_m : $F(2, 26) = .764, p = .476$, partial $\eta^2 = .056$). As seen in Figure 5, comparable color s.d. and P_m can be found under conditions without categorical judgment (‘color-only’, ‘color and orientation’) and with categorical judgment (‘color and orientation with categorical judgment’) (Figure 5, b1). This is true for individual orientation representation as well (Figure 5, b2).

Despite a robust set size effect at the item level (Figure 5, b1-b2, also mentioned in the next paragraph], the interaction between set size and task demand was not significant either for color-item representations [color s.d.: $F(4, 52) = 1.953, p = .116$, partial $\eta^2 = .132$; color P_m : $F(4, 52) = .544, p = .704$, partial $\eta^2 = .040$), or orientation-item representations [orientation s.d.: $F(4, 52) = .411, p = .800$, partial $\eta^2 = .031$; orientation P_m : $F(4, 52) = .634, p = .640$, partial $\eta^2 = .047$]. Thus, in Experiment 3, we did not find that the inclusion of the prior categorical judgment task has an effect on the item measures in the second continuous report task. Moreover, this null finding did not change as a function of the number of items presented in the display. In other words, the general finding of set size effect at the individual-item level (i.e., item limit)

throughout Experiments 1-3 was not due to the serial nature of the category-then-item task or its associated memory decay.

As briefly mentioned above, the set size effect on item representations was again replicated in this experiment. For color-item representations, the significant main effects of set size were found for both s.d. ($F(2, 26) = 8.663, p = .001$, partial $\eta^2 = .400$) and P_m ($F(2, 26) = 36.328, p < .001$, partial $\eta^2 = .736$). Pairwise comparisons indicated that the average color s.d. (collapsed across all relevant conditions: ‘color-only’, ‘color and orientation’, ‘color and orientation with categorical judgment’) at set size eight was significantly larger than set size six [$F(1, 13) = 6.672, p = .022$, partial $\eta^2 = .342$] and set size four [$F(1, 13) = 8.538, p = .012$, partial $\eta^2 = .396$], whereas no significant difference was found between set size four and six [$F(1, 13) = 2.493, p = .138$, partial $\eta^2 = .161$]. In terms of color P_m , there was a significant drop between any two adjacent set sizes [4 vs. 6: $F(1, 13) = 37.719, p < .001$, partial $\eta^2 = .744$; 6 vs. 8: $F(1, 13) = 21.668, p < .001$, partial $\eta^2 = .625$].

For orientation-item representations, we also observed a significant main effect of set size in orientation s.d. ($F(2, 26) = 3.852, p = .034$, partial $\eta^2 = .229$) and orientation P_m ($F(2, 26) = 11.878, p < .001$, partial $\eta^2 = .477$). Paired sample t-test showed that the mean orientation s.d. (collapsed across all relevant conditions: ‘orientation-only’, ‘color and orientation’, ‘color and orientation with categorical judgment’) was significantly smaller at set size four than set size six ($F(1, 13) = 10.531, p = .006$, partial $\eta^2 = .448$), whereas the difference between set size six and eight was not significant ($F(1, 13) = .003, p = .960$, partial $\eta^2 < .001$). Similar to color P_m , here orientation P_m also showed a monotonic decline across all set sizes (four vs. six: $F(1, 13) = 5.440, p = .014$, partial $\eta^2 = .381$; six vs. eight: $F(1, 13) = 8.002, p = .036$, partial $\eta^2 = .295$). Taken together, we replicated the set size effect at the individual-item level found in

Experiments 1-2 and further extended it by shifting the set size from two, four, and six to four, six and eight.

Last but not least, even though the main effect of task demand was not significant, we still explored the potential effect of the presence of object category (sub-conditions of task demand) on item measures using a repeated-measures ANOVA with category presence (present or absent) and set size (four, six, and eight) as within-subject factors. We found no significant effect of the presence of object category on individual color-item representations [color s.d.: $F(1, 13) = .050, p = .826$, partial $\eta^2 = .004$; color P_m: $F(1, 13) = .1.193, p = .294$, partial $\eta^2 = .084$], or individual orientation-item representations [orientation s.d.: $F(1, 13) = .190, p = .670$, partial $\eta^2 = .014$; orientation P_m: $F(1, 13) = 1.763, p = .207$, partial $\eta^2 = .119$]. These null results demonstrate that item representation is not sensitive to the presence of multiple object categories (color and orientation) in the display.

If the ensemble and item levels of representations were stored in a common pool of VWM resource, the absence of the categorical judgment task should free up some VWM resource to store more information about individual items. This was not what we found in Experiment 3. The estimates of item representations were not affected by the inclusion of a category judgment task, implying an independent relationship between the two levels of VWM representations. In addition, given that object categories can be perceived in a rapid and automatic fashion (for a review, see Feldman, Griffiths, & Morgan, 2009), it is striking that item measures were not affected by the presence of object category in the display. Perhaps the perception of object category is, instead, directly linked to the global level of the display which cannot be measured in the present design (i.e., no categorical judgment task in the single category conditions (color- or orientation-only)).

Taken together, in this experiment we demonstrated that estimates of item representations were not affected by manipulation at the ensemble level, such as the inclusion of the category judgment task or the presence of object category in the display. It is also important to note another general finding replicated in all experiments conducted so far (Experiments 1-3): set size effect was found in item measures but not in the ensemble measure. Thus, a critical implication from these observations is a potentially independent relationship between the ensemble and item representations, which we further address in Experiment 4.

Experiment 4

In Experiment 4, we tested the idea of independence between ensemble and individual representations that we had already alluded to in Experiment 3. Specifically, we asked if ensemble and item representations are differentially sensitive to the layout of items in the display. Previous studies have converged to suggest that estimates of VWM capacity (usually from change detection tasks) can be sensitive to different grouping principles (Brady & Tenenbaum, 2013). It is not clear, however, which level of representation contributes to the inflated VWM capacity. Our hypothesis was that, if ensemble and item representations are integrated and stored in a shared pool of VWM resources, then varying layout would affect both types of representation similarly. In contrast, if these two levels of representation are independent from each other, the extent to which they are subject to layout manipulations should be different.

To avoid circular reasoning, we elect to use an a-priori definition of perceptual grouping for layout manipulation rather than quantify the amount of perceptual grouping in the display using post-hoc analyses of participants' responses. In this regard, Brady and Tenenbaum (2013)

observed that displays where the items are heterogeneous and intermixed were much harder to remember than displays where physically similar items were spatially clustered, e.g., all the left-most items are warm colors and all the right-most items are cold colors. Based on their findings we systematically vary the amount of heterogeneity in the display by intermixing color and orientation Gabors (Figure 7). In the ‘easy’ condition, the layout of items can be easily separated into two salient spatio-categorical groups: color Gabors on one side and orientation Gabors on the other side along the circular array. In contrast, to maximize the amount of heterogeneity in the ‘hard’ condition, each color Gabor is placed next to an orientation Gabor such that the stimulus category is fully intermixed in the circular array. We piloted the layout manipulation prior to this experiment which included another ‘medium’ condition to allow us to parametrically vary the level of perceptual grouping³. See Figure 7 for a schematic illustration of different layout conditions.

³ Note that although the layout conditions were referred to as ‘easy’, ‘medium’, ‘hard’, we do not imply a strictly linear scale in the level of perceptual grouping.

Figure 7. A schematic illustration of different layout and set size conditions in Experiment 4.

Gabors are placed in a circular array (same as in previous experiments), in which color Gabors are marked as red ‘C’ and orientation Gabors as blue ‘O’ for illustration purposes. In the easy layout condition, color Gabors are placed on one half and orientation Gabors on the other half of any imaginary diameter of the circular array. In the hard layout condition, color and orientation Gabors are fully intermixed in a way that no Gabors of the same category are presented next to each other. In the medium conditions, we ensure that adjacent neighbors of each Gabor come from different categories (one color and one orientation).

Method

Participants

13 students from the University of Hong Kong participated in the study for course credit. All reported normal or corrected-to-normal visual acuity, and normal color vision. Informed

consent was obtained prior to the start of the experiment and the experiment protocol was approved by the Institutional Review Board of the University of Hong Kong.

Apparatus & Stimuli

The apparatus and stimuli were identical to those used in Experiments 2, except that the item layout in the display was systematically manipulated. At each set size, there were three different layout conditions, corresponding to how easy it was to perceptually group items in the display into the color or orientation category. For illustration purposes, the location of color items was marked as “C” and the location of orientation items was marked as “O” in Figure 7. In the easy condition, color Gabors were placed on one half and orientation Gabors on the other half of any imaginary diameter of the circular array. In the hard condition, color and orientation Gabors were fully intermixed in a way that no Gabors of the same category are presented next to each other. In the medium condition, we ensured that the two neighbors of each Gabor came from different categories (one color and one orientation).

Design

This study had two independent variables: layout (easy, medium, hard), and set size (six, eight). The dependent variables were categorical accuracy (measure of ensemble representation), P_m and s.d. (measures of item representation) of each layout condition, at each set size.

Procedure

The procedure was identical to the ‘color and orientation with categorical judgment’ condition in Experiment 3. There were 70 trials per grouping condition at each set size, so there were 420 trials in total and the whole experiment took around 55 min to finish.

Model

Same as Experiments 2-3, the standard mixture model (“slot+averaging”) of Zhang and Luck (2008) was applied to model color and orientation trials using the MemToolbox (memtoolbox.org). By default, a Markov Chain Monte Carlo (MCMC) algorithm was used for model fitting and the maximum a posteriori (MAP) estimate was obtained for each parameters (g, s.d.), for each participant, at each set size (for more information, see Suchow, Brady, Fougner, & Alvarez, 2013). Note that $P_m = 1 - g$.

Results and Discussion

Ensemble representation

As a preliminary analysis, a 3×2 within-subject ANOVA with layout (easy, medium, and difficult) and set size (6, 8) as independent variables was conducted on the categorical accuracy. There was a significant main effect of layout, $F(2, 24) = 6.673, p = .005$, partial $\eta^2 = .357$, but neither the main effect of set size [$F(2, 24) = 3.660, p = .080$, partial $\eta^2 = .234$] nor the interaction between layout and set size was significant [$F(2, 24) = .976, p = .391$, partial $\eta^2 = .075$]. Therefore, categorical judgments at set sizes 6 and 8 were collapsed for further analysis and the average categorical accuracy was 93.79%, 91.65%, and 90.05% in the ‘easy’, ‘medium’, and ‘hard’ conditions, respectively. A follow-up paired-samples t-test showed a significant difference between performance in the easy and hard conditions, $t(12) = 3.208, p = .008$, *Cohen’s d* = .926.

To summarize key results on ensemble representation in this experiment, we replicated the null effect of set size on ensemble representation for the fourth time in the current series of studies. More importantly, Experiment 4 critically showed that ensemble measures could be

sensitive to different layouts of items in the display, such that reporting the category of an item was more challenging in the hard condition than in the easy condition.

Figure 8. Results from Experiment 4. (A) Mean accuracy of the categorical judgment in the easy, medium, hard layout conditions (measure of ensemble representation). (B1) The s.d. and P_m for color (measures of item representation). (B2) The s.d. and P_m for orientation (measures of item representation). Error bars show ± 1 standard error of mean.

Item representation

A repeated-measures ANOVA with factors layout (easy, medium, hard) and set size (6, 8) was run separately for s.d. and P_m . The right panel of Figure 8 (b1 & b2) shows that the main effect of layout was not significant in either color s.d., $F(2, 24) = .456$, $p = .639$, partial $\eta^2 = .037$,

or in orientation s.d., $F(2, 24) = 1.322, p = .285$, partial $\eta^2 = .099$. Similarly, the effect of layout was not found in color P_m , $F(2, 24) = 1.149, p = .334$, partial $\eta^2 = .087$, or orientation P_m , $F(2, 24) = .099, p = .906$, partial $\eta^2 = .008$. Here the absence of layout effect in both item measures was in sharp contrast with the significant main effect of layout in the ensemble measure mentioned earlier, but is consistent with our assumption that ensemble and item representations function independently. That is, the layout manipulation only affected representation at the ensemble level, but not at the item level. This conclusion aligns with the results of a recent study by Brady and Alvarez (2015) that found clustering of items on the display allows people to detect across-category changes better but has no effect on within-category changes.

Although the set size effect was not significant in either color s.d. [$F(2, 24) = 1.738, p = .212$, partial $\eta^2 = .127$], or orientation s.d. [$F(2, 24) < .001, p = .994$, partial $\eta^2 < .001$], the set size effect was significant in orientation P_m [$F(2, 24) = 10.952, p = .006$, partial $\eta^2 = .477$], and was approaching significance in color P_m [$F(2, 24) = 4.584, p = .053$, partial $\eta^2 = .276$], indicating that P_m decreased when the set size increased from 6 to 8. The observed drop in P_m with increasing set size is consistent with the prediction from the discrete slot model (Zhang & Luck, 2008) and replicates the previous observations in Experiment 3. There was no interaction between layout and set size in any s.d. or P_m parameter for color or orientation (s.d. for color: $F(2, 24) = .839, p = .444$, partial $\eta^2 = .065$; P_m for color: $F(2, 24) = 1.207, p = .317$, partial $\eta^2 = .091$; s.d. for orientation: $F(2, 24) = .158, p = .854$, partial $\eta^2 = .013$; P_m for orientation: $F(2, 24) = .855, p = .438$, partial $\eta^2 = .067$).

Modeling ensemble and item representations

To further examine effects of set size and layout on ensemble and item representations respectively, we adopted Bayesian hierarchical modeling approach to fit empirical observations

of ensemble and item representations, and then compared different models using standard model comparison techniques. Three nested models were constructed to fit performance in the ensemble task. One included layout but not set size as a predictive factor (*Model_{ensemble-layout}*), one included set size but not layout (*Model_{ensemble-set size}*), and another without any predictive factors as the baseline model (*Model_{ensemble-simple}*). Likewise, we constructed three models to fit performance in the item task (*Model_{individual-layout}*, *Model_{individual-set size}*, and *Model_{individual-simple}*). Our prediction is that, if ensemble representations are sensitive to different layouts of items in the display, *Model_{ensemble-layout}* should outperform the other two models. In contrast, if item representations are sensitive to different set sizes in the display, *Model_{individual-set size}* should outperform the other two models.

All the models were fitted via Gibbs sampling method, using the package of JAGS (Plummer, 2003) in the environment of R project. For each model, we generated one chain with 100,000 samples and thinned the chain by a factor of 10 to de-correlate the samples, resulting in 10,000 final samples of the posterior distribution in total. More details of the model structure can be found in Appendix A.

Based on the posterior samples, we performed model comparison using calculated WAICs (Watanabe, 2010, 2013). Table 1 lists the WAICs for all models. Amongst all models of ensemble representation, *Model_{ensemble-layout}* was favored by its lowest WAIC value. The difference in WAIC between *Model_{ensemble-simple}* and *Model_{ensemble-set size}* was 0.79, smaller than one standard error of their WAIC difference (2.90), indicating no particular advantage over the other model.

The lowest WAIC amongst all models that fit the item representations was found in *Model_{individual-set size}*, which was lower than the other two models by 10 times of the +/-1 standard

error of their WAIC difference. By comparison, $Model_{individual-layout}$ (13984.89) showed numerically higher WAIC than $Model_{individual-simple}$ (13976.82), although such difference (8.07) is inconclusive given the ± 1 standard error of their WAIC difference (5.24). Taken together, the conclusions from model comparisons are in line with the dissociation observed in the empirical results that layout was the single best predictor of performance in the ensemble task, whereas set size best predicted performance in the item task.

Table 1

Model fits for models of ensemble and item representations applied to Experiment 4.

Ensemble representation		Item representation	
Model	WAIC	Model	WAIC
$Model_{ensemble-simple}$	2620.21	$Model_{individual-simple}$	13976.82
<i>*$Model_{ensemble-layout}$</i>	2604.76	$Model_{individual-layout}$	13984.89
$Model_{ensemble-set\ size}$	2619.42	<i>*$Model_{individual-set\ size}$</i>	13936.37

* indicates the best fitting models according to WAIC.

In sum, the results of Experiment 4 indicate that ensemble representations are affected by layout manipulation but not set size, whereas item representations are not sensitive to layout manipulation but are subject to set size changes (replicating results from Experiments 1-3). Such findings suggest distinct mechanism and parallel processing of hierarchical VWM representations, which further corroborate findings from Experiment 3 that the ensemble and item representation levels are stored independently in VWM.

Although the layout manipulation was the only factor so far that seems to substantially affect ensemble representation in VWM (in contrast to the absence of any set size effect at the ensemble level in Experiments 1-4), it is worth revisiting the operational definition of the layout to better understand the nature of this effect. In this experiment, we followed the heterogeneity principle from Brady and Alvarez (2015) to intermix item category in our display. According to this principle, the hardest layout condition would require all adjacent items come from different categories. A potential concern we had for the layout manipulation was that participants may figure out the adjacent nature of item category and develop strategies to successfully report the category of the probed item simply based on the distance between this item and the remembered item(s) in the circular array. If this were the case, the layout effect at the ensemble level should be diminished. This was not what we found, however, categorical judgment accuracy was significantly higher in the ‘easy’ than in the ‘hard’ condition. Indeed, from the pilot testing prior to this experiment, participants noted that the physical similarity of color/orientation categories, random mixture of conditions, fast presentation of arrays, and serial nature of the tasks discouraged them to develop this strategy. In this regard, it is useful to pilot the design in an independent subject population to allow us to validate the layout manipulation.

Ensemble representation at different set sizes: Bayesian hierarchical models

Despite a highly consistent absence of set size effect on ensemble representation across experiments, it is likely that this null effect results from a lack of statistical power due to a relatively small sample size in each experiment. It was, however, not our intention for a sufficient demonstration that the precision of ensemble representations are robust to any number of items on the display, and it certainly would not be, in the most extreme case. Admittedly,

except in Experiment 1, there is a global trend of decreasing accuracy with increasing set size in the ensemble measure, leaving the possibility of a smaller set size effect on ensemble representation accuracy.

We use the Bayesian hierarchical modeling approach to go beyond the descriptive summary statistics of the empirical data. Here, rather than dwelling on the difference between significant and not-significant (Gelman & Stern, 2006; Nieuwenhuis, Forstmann, & Wagenmakers, 2011) set size effect on the precision of ensemble and item representations, we aim to inform theories of VWM using theoretically interpretable model parameters.

Modeling ensemble representation under different set sizes

To increase statistical power, performance in the categorical judgement task from Experiments 2, 3 and 4, which shared the same stimuli and similar procedure, was combined in the two nested Bayesian hierarchical models described below. In the first hierarchical model (*Model_{ensemble-combined-set size}*), we made a key assumption that a decrease in performance with increasing set size is expected with the worst performance at set size 8. Thus, in each of the other set size conditions, a free parameter called “performance difference” (i.e. $\Delta A_{8 \rightarrow 2}$, $\Delta A_{8 \rightarrow 4}$, $\Delta A_{8 \rightarrow 6}$) was used to describe difference in accuracy between this particular set size and set size 8 (see Appendix B for model details). If the set size effect on ensemble representation exists, the fitted “performance difference” parameter should be positive (value > 0), indicating better performance with fewer items on the display.

Model construction. This model (*Model_{ensemble-combined-set size}*) was pit against the other model (*Model_{ensemble-combined-simple}*) which had the exact same structure except no inclusion of the parameter “performance difference” ($\Delta A_{8 \rightarrow 2}$, $\Delta A_{8 \rightarrow 4}$, $\Delta A_{8 \rightarrow 6}$). That is, performance in the categorical judgement task was modeled assuming no effect of set size on ensemble

representations. As a result, *Model_{ensemble-combined-simple}* predicts that performance in the categorical judgement task should not differ across set size conditions. These two models (*Model_{ensemble-combined-set size}*, *Model_{ensemble-combined-simple}*) were then compared using standard model comparison techniques.

Model fitting. Performance in the categorical judgement task was fitted by each model using JAGS (Plummer, 2003). We generated one chain for each model, which had 100,000 samples and was thinned by a factor of 10 to de-correlate the samples, resulting in 10,000 final samples of the posterior distribution in total. As mentioned above, we focus on the posterior samples of $\Delta A_{8 \rightarrow 2}$, $\Delta A_{8 \rightarrow 4}$, and $\Delta A_{8 \rightarrow 6}$ (Figure 9) in order to compare these two models.

Model comparison. The WAIC difference between *Model_{ensemble-combined-set size}* (4054.17) and *Model_{ensemble-combined-simple}* (4056.93) was merely 3.77, smaller than one standard error of their WAIC difference (4.78). Thus, neither model was favored. Next, to elucidate the effect of set size on ensemble representation, we further examined the range of 95% highest density intervals (*HDI*) from the only model that contained the predictor of set size (*Model_{ensemble-combined-set size}*). In Figure 9, the 95% HDI was 0 – 0.0250 in $\Delta A_{8 \rightarrow 2}$ and 0 – 0.0233 in $\Delta A_{8 \rightarrow 4}$, suggesting numerically higher precision of ensemble representations when set size decreased from 8 to 2 or 4. Note that these two posterior means were only around 0.012 and the lower bounds of their 95% HDIs were higher than 0 by a negligible amount, suggesting that the effect of set size is highly limited, if there is any. Moreover, the effect of set size change from 8 to 6 was not conclusive according to the model because the 95% *HDI* of $\Delta A_{8 \rightarrow 6}$ ranged from -0.001 to 0.0170, which included values smaller than 0.

Although the mixed results from the model do not allow us to conclude an effect of set size on ensemble representation, this effect, if there is any, seems to be highly limited according

to a non-conclusive WAIC difference between $Model_{ensemble-combined-set\ size}$ and $Model_{ensemble-combined-simple}$. Moreover, the set size effect on ensemble measures (categorical accuracy) and item measures (e.g. P_m , probability that the probed item is present in memory) were far from comparison. For example, when set size decreased from 8 to 4, the mean ensemble accuracy only increased by 1.17% (Figure 9, middle panel); by contrast, the mean P_m decreased more than 10% (see Figure 8).

Figure 9. Histograms of the posterior samples of $\Delta A_{8 \rightarrow 2}$ (left), $\Delta A_{8 \rightarrow 4}$ (center), $\Delta A_{8 \rightarrow 6}$ (right) from the fitted $Model_{ensemble-combined-set\ size}$.

General Discussion

In the present series of experiments and models, we try to address two important questions regarding the structure of VWM: 1) whether precision of ensemble/item representations changes with varying set size/layout, and 2) whether ensemble and individual representations are stored and/or processed in an integrated or independent fashion. First, we found qualitatively different effects of set size and layout on the precision of visual memory

representations of the same item at these two levels of VWM. Throughout Experiments 1-4, the measure of ensemble representation (i.e., accuracy of categorical judgments) was not affected by set size, whereas the item representations (i.e., color/orientation report or cube discrimination) showed robust set size effects. By contrast, ensemble representation, but not item representation, was sensitive to layout manipulations in the display (Experiment 4). In a complementary fashion, hierarchical Bayesian models fit to the empirical data suggest distinct effects of layout and set size on ensemble and item representations. Second, the removal of the ensemble measure did not free up VWM resources to enable more storage or higher precision of individual items (Experiment 3), implying that the storage of the representations at these two levels may rely on different memory resources. Converging results suggest an independent relationship between the ensemble and the item representations in VWM. Below we discuss how the two broad conclusions we draw may have implications for the existing theories of VWM. Most, if not all, hold an implicit assumption that items in the display are stored independently and do not consider the effect of contextual information in the display on VWM measures (Bays et al., 2009; Zhang & Luck, 2008, 2011).

Measuring Visual Working Memory

The change detection paradigm has remained one of the most prominent approaches in the study of VWM capacity (for a review, see Luck & Hollingworth, 2008), as it is primarily sensitive to VWM (i.e., non-memory processes involved in this task are minimal) and it closely mimics the way in which VWM is used in the natural environment, such as integrating views of the world across saccades (Hollingworth, Richard, & Luck, 2008), and allowing the comparison of objects that cannot be attended at once (Pomplun, Reingold, & Shen, 2001). Recent evidence, however, challenges the notion of capacity estimates derived from the change detection

paradigm, as it does not take into account the precision of memory required in order to detect a change. Awh et al. (2007) first showed that detecting a change between the sample and the test display could reflect either the number of objects stored (cross-category change) or the precision of an object in VWM (within-category change). Using similar displays and the same paradigm as in the Awh et al. (2007) study, Schurgin and Brady (2019) further demonstrated that cross-category change detection is primarily sensitive to low-resolution spatial ensemble information, which in turn, leads to inflated item capacity when detecting large changes in the display.

In light of this limitation of the traditional change detection paradigm, we developed a novel dual report task to provide measures of ensemble and item representations of the same object. Although previous research has attempted to measure the influence of categorical and precise continuous representations in delayed estimation tasks (Bae, Olkkonen, Allred, & Flombaum, 2015; Hardman, Vergauwe, & Ricker, 2017; Pratte, Park, Rademaker, & Tong, 2017), the present study is the first attempt to measure both ensemble and item representations within a single trial. Because the second within-category discrimination task is contingent on the response from the first categorical judgment task, an item in VWM can be simultaneously construed at hierarchical levels.

Given that this is a new task, a number of caveats and limitations are warranted. First, categorical arrays encourage color and orientation to be sampled independently (Bays, Wu, & Husain, 2011; Nelson Cowan, Blume, & Scott Saults, 2013; Fougine & Alvarez, 2011). Thus, as set size increases, it is increasingly likely that participants will have either color or shape/orientation in memory for a particular item, leading to a shallower effect of set size for the category judgement. If this were the case that participants only remember color or shape/orientation at larger set sizes, manipulating the layout of the display should not affect

ensemble measures and/or equally affect both levels of representations. This is not what we found in Experiment 4. Instead, the ensemble measure (categorical judgement) was sensitive to different layouts of items in the display, but individual item measure (color or orientation task) was not, suggesting qualitatively different roles of layout and set size in the ensemble representation.

The second caveat relates to the sensitivity of each measure to detect set size/layout effect, such that the two tasks actually measure the same item representation but have different power to detect memory precision with set size. Although it is a likely alternative explanation, two aspects in our data showed compelling evidence against the possibility that the item task simply has more power to detect a significant effect. First, we included a critical manipulation of layout in Experiment 4 and our reasoning is that, if the two tasks measure the same underlying construct or if the item task is simply more sensitive than the ensemble/category task, then it is not possible to find a significant layout effect in the ensemble measure, but not in the individual item measure. Second, we reasoned that if the ensemble and item levels of representations were stored in a common pool of VWM resource, removing the categorical judgment task should free up VWM resource to store more information about individual items. This was not what we found in Experiment 3. The estimates of item representations were not affected by the inclusion of a category judgment task, implying an independence between the two levels of VWM representations. Both observations indicate that mechanism underlying these two tasks is not the same and that the two tasks seem to measure fundamentally different constructs.

Despite these potential limitations or alternative explanations, the finding of distinct ensemble and item representations speaks directly to the divergence on item limit between discrete slot and flexible resource models. According to the discrete slot models, an item is either

stored or not stored, and no further information can be stored from additional items once the item limit is reached (Luck & Vogel, 1997, 2013). Partially consistent with this claim, measures at item level in our design showed that the probability of a given item stored in VWM reduced significantly and its precision falls to its nadir once the “item limit” is reached (Zhang & Luck, 2008). However, the close-to-perfect performance of this very same item at the ensemble level beyond the “item limit” clearly indicates that the item representation is not the only representation present in VWM. In other words, an item in the display can be represented at the ensemble level even though it is not precisely “stored” at the individual level. This is consistent with the default assumption in flexible resource theories that all items are stored in VWM (Alvarez & Cavanagh, 2004, 2010; Bays & Husain, 2008; Bays et al., 2009; Wilken & Ma, 2004). Although our design does not allow us to speak to flexible resource allocation in VWM, the qualitatively different set size effect on item and ensemble representations suggests that participants obligatorily encode both low-resolution ensemble representations of the whole display and high-resolution representations of a subset of individual items in the display.

Parallel Processing of Individual Objects and Ensembles

Although it has been increasingly acknowledged that working memory representations are hierarchically structured (for a review, see Alvarez, 2011), the exact relationship between item and ensemble level representations is not clear. Previous studies related to ensemble representations showed that ensemble information can be quickly extracted from the visual display, such as mean size (Ariely, 2001; Chong & Treisman, 2003, 2005), mean orientation (Dakin & Watt, 1997; Lund, Parkes, Solomon, Angelucci, & Morgan, 2002), mean emotion or gender (Haberman & Whitney, 2007), and spatial context of the items (Jiang et al., 2000; Schurgin & Brady, 2019; Vidal et al., 2005). Furthermore, Brady and Alvarez (2011)

demonstrated that memory for individual items in a display can be biased by ensemble statistics of the whole display, implying that information across these two levels is integrated in VWM.

However, the present dissociation between ensemble and item information (which is not available in previous paradigms) allows us to suggest independent and parallel processing of objects and ensembles at different levels. For example, the item representation was sensitive to the number, but not to the spatial layout of items in the display. In contrast, ensemble representation was affected by spatial layout of items in the display but was robust to the number of items (up to eight in our experiments). In addition, removing the ensemble measure did not improve the precision of item representations, suggesting that these two levels of representation are unlikely to share a common pool of VWM resources.

How do we reconcile these two incompatible views on the relationship between ensemble and item representations? One possibility is the task context which can encourage integrated encoding. For example, Experiment 1 of Brady and Alvarez (2011) instructed participants to attend to the red and the blue circles (e.g., color-based encoding) which biased observers' size report of a given circle by the size of the other circles of the same color in VWM. This bias occurred, because an item's reported feature (i.e., item representation) can be constructed by combining information with other items along the same feature. Such bias was not found, however, when color was no longer a relevant feature to encode in Experiment 2 of Brady and Alvarez (2011). Taken together, a non-automatic integration of information between individual and ensemble levels can be demonstrated as a result of bias from task demand.

On the other end of the continuum is the demonstration of independence between ensemble and item representations. For example, in our present task, knowing the ensemble information (i.e., object category) did not provide detailed information about individual items for

successful within-category discrimination. Given that the task context did not encourage integration across levels, this allowed us to observe an independence between item and ensemble representation in VWM. Consistent with our view of independent processing, Brady and Alvarez (2015) also found that clustering of items on the display allowed people to detect cross-category changes better, but not within-category changes.

Implications of dissociable VWM representations

The present dissociation between ensemble level and individual level representations in VWM is in line with the neural object-file theory (Xu & Chun, 2009) and the two-factor model (Awh et al., 2007). According to the neural object-file theory (Xu & Chun, 2009), coarse-resolution object representations (e.g., object category) and fine-resolution object representations (e.g., object identity) might be stored independently of each other. Thus, parallel processing ensemble and item representations demonstrated in the current study could be related to distinct neural correlates of VWM representations at different levels of processing (Xu & Chun, 2006, 2009).

Indeed, while the lateral occipital complex and superior intra-parietal sulcus are known to be involved in individuating and selecting multiple objects for detailed processing (Xu & Chun, 2006, 2009), the anterior-medial ventral visual cortex was found to be sensitive to ensemble representations (Cant & Xu, 2012), such as object density (Cant & Xu, 2015) and object shape and surface properties (Cant & Xu, 2017). Moving forward, our current dual-report task could be adapted for potential use in fMRI to provide further support for the independent neural mechanisms of ensemble and item processing.

Our results are consistent with idea of separate working memory dimensions reflecting distinct aspects of working memory ability. In regards to the two-factor model (Awh et al.,

2007), our ensemble measure resembles the number of representations that can be maintained whereas our item measure reflects the resolution of those representations. Relevant to the understanding of dissociable VWM representations and cognitive function, Awh and colleagues show that the number of items that can be maintained in working memory is strongly correlated with measures of fluid intelligence (Fukuda, Vogel, Mayr, & Awh, 2010), whereas working memory resolution is sensitive to the level of perceptual expertise for specific object categories (Scolari, Vogel, & Awh, 2008). Thus, individual differences in number and resolution have implications for cognitive functions more broadly and need to be explored in depth in future research.

The dissociation between ensemble and individual level representations can also be relevant to the understanding of the core systems of number in the field of numerical cognition. We posit the low-resolution ensemble representation underlies processing of large, approximate numerosities, while the high-resolution item representations subserve processing of small, precise individual objects (Feigenson, Dehaene, & Spelke, 2004). Such dissociation has further implications on math aptitude as a number of studies converged to suggest that individual differences in representing approximate numerosities of items correlate with math aptitude and achievement (Halberda, Mazzocco, & Feigenson, 2008; Libertus, Feigenson, & Halberda, 2011).

Conclusions

This series of studies converged to demonstrate the nature of the hierarchical structure in VWM. Consistent with recent observations that VWM is inherently hierarchical (Brady & Alvarez, 2011, 2015; Schurgin & Brady, 2019), we found that VWM holds both representations of low-resolution ensemble information and a subset of high-resolution individual items in the

display. Furthermore, we suggested these two levels of representation are independent and likely to be processed in parallel. The observed dissociation between the ensemble and the item representation suggests distinct mechanisms underlying human ability to remember both “forest” and “trees” in the scene.

References

- Alvarez, G. A., & Cavanagh, P. (2004). The Capacity of Visual Short-Term Memory Is Set Both by Visual Information Load and by Number of Objects. *Psychological Science*, 15(2), 106–111. <https://doi.org/10.1111/j.0963-7214.2004.01502006.x>
- Alvarez, G. A., & Cavanagh, P. (2010). The capacity of visual short-term memory is set by total information load, not number of objects. *Journal of Vision*, 2(7), 273–273. <https://doi.org/10.1167/2.7.273>
- Alvarez, George A. (2011). Representing multiple objects as an ensemble enhances visual cognition. *Trends in Cognitive Sciences*, Vol. 15, pp. 122–131. <https://doi.org/10.1016/j.tics.2011.01.003>
- Ariely, D. (2001). Seeing Sets: Representation by Statistical Properties. *Psychological Science*, 12(2), 157–162. <https://doi.org/10.1111/1467-9280.00327>
- Awh, E., Barton, B., & Vogel, E. K. (2007). Visual working memory represents a fixed number of items regardless of complexity. *Psychological Science*, 18(7), 622–628. <https://doi.org/10.1111/j.1467-9280.2007.01949.x>
- Baddeley, A. (2007). Working Memory, Thought, and Action. In *Oxford University Press*. <https://doi.org/10.1093/acprof:oso/9780198528012.001.0001>
- Bae, G. Y., Olkkonen, M., Allred, S. R., & Flombaum, J. I. (2015). Why some colors appear more memorable than others: A model combining categories and particulars in color working memory. *Journal of Experimental Psychology: General*, 144(4), 744–763. <https://doi.org/10.1037/xge0000076>
- Bays, P., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science*, 321(5890), 851–854. <https://doi.org/10.1126/science.1158023>

- Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, 9(10), 7–7.
<https://doi.org/10.1167/9.10.7>
- Bays, Paul M., Wu, E. Y., & Husain, M. (2011). Storage and binding of object features in visual working memory. *Neuropsychologia*.
<https://doi.org/10.1016/j.neuropsychologia.2010.12.023>
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of Vision*, 11(5), 4–4.
<https://doi.org/10.1167/11.5.4>
- Brady, Timothy. F., & Alvarez, G. A. (2011). Hierarchical encoding in visual working memory: Ensemble statistics bias memory for individual items. *Psychological Science*, 22(3), 384–392. <https://doi.org/10.1177/0956797610397956>
- Brady, Timothy. F., & Alvarez, G. A. (2015). Contextual effects in visual working memory reveal hierarchically structured memory representations. *Journal of Vision*, 15(15), 6.
<https://doi.org/10.1167/15.15.6>
- Cant, J. S., & Xu, Y. (2012). Object Ensemble Processing in Human Anterior-Medial Ventral Visual Cortex. *Journal of Neuroscience*, 32(22), 7685–7700.
<https://doi.org/10.1523/JNEUROSCI.3325-11.2012>
- Cant, J. S., & Xu, Y. (2015). The impact of density and ratio on object-ensemble representation in human anterior-medial ventral visual cortex. *Cerebral Cortex*, 25(11), 4226–4239.
<https://doi.org/10.1093/cercor/bhu145>
- Cant, J. S., & Xu, Y. (2017). The Contribution of Object Shape and Surface Properties to Object Ensemble Representation in Anterior-medial Ventral Visual Cortex. *Journal of Cognitive*

- Neuroscience*, 29(2), 398–412. https://doi.org/10.1162/jocn_a_01050
- Chong, S. C., & Treisman, A. (2003). Representation of statistical properties. *Vision Research*, 43(4), 393–404. [https://doi.org/10.1016/S0042-6989\(02\)00596-5](https://doi.org/10.1016/S0042-6989(02)00596-5)
- Chong, S. C., & Treisman, A. (2005). Attentional spread in the statistical processing of visual displays. *Perception & Psychophysics*, 67(1), 1–13. <https://doi.org/10.3758/BF03195009>
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(1), 87–114. <https://doi.org/10.1017/S0140525X01003922>
- Cowan, Nelson. (2010). The magical mystery four: How is working memory capacity limited, and why? *Current Directions in Psychological Science*, 19(1), 51–57. <https://doi.org/10.1177/0963721409359277>
- Cowan, Nelson, Blume, C. L., & Scott Saults, J. (2013). Attention to attributes and objects in working memory. *Journal of Experimental Psychology: Learning Memory and Cognition*. <https://doi.org/10.1037/a0029687>
- Dakin, S. C., & Watt, R. J. (1997). The computation of orientation statistics from visual texture. *Vision Research*, 37(22), 3181–3192. [https://doi.org/10.1016/S0042-6989\(97\)00133-8](https://doi.org/10.1016/S0042-6989(97)00133-8)
- Feigenson, L., Dehaene, S., & Spelke, E. (2004). Core systems of number. *Trends in Cognitive Sciences*, 8(7), 307–314. <https://doi.org/10.1016/j.tics.2004.05.002>
- Feldman, N. H., Griffiths, T. L., & Morgan, J. L. (2009). The Influence of Categories on Perception: Explaining the Perceptual Magnet Effect as Optimal Statistical Inference. *Psychological Review*, 116(4), 752–782. <https://doi.org/10.1037/a0017196>
- Fougnie, D., & Alvarez, G. A. (2011). Object features fail independently in visual working memory: Evidence for a probabilistic feature-store model. *Journal of Vision*.

<https://doi.org/10.1167/11.12.3>

Fukuda, K., Vogel, E., Mayr, U., & Awh, E. (2010). Quantity, not quality: The relationship between fluid intelligence and working memory capacity. *Psychonomic Bulletin and Review*, 17(5), 673–679. <https://doi.org/10.3758/17.5.673>

Gelman, A., & Stern, H. (2006). The difference between “significant” and “not significant” is not itself statistically significant. *American Statistician*.
<https://doi.org/10.1198/000313006X152649>

Haberman, J., & Whitney, D. (2007). Rapid extraction of mean emotion and gender from sets of faces. *Current Biology*, 17(17), R751–R753. <https://doi.org/10.1016/j.cub.2007.06.039>

Halberda, J., Mazocco, M. M. M., & Feigenson, L. (2008). Individual differences in non-verbal number acuity correlate with maths achievement. *Nature*, 455(7213), 665–668.
<https://doi.org/10.1038/nature07246>

Hardman, K. O., Vergauwe, E., & Ricker, T. J. (2017). Categorical working memory representations are used in delayed estimation of continuous colors. *Journal of Experimental Psychology: Human Perception and Performance*, 43(1), 30–54.
<https://doi.org/10.1037/xhp0000290>

Hollingworth, A., Richard, A. M., & Luck, S. J. (2008). Understanding the Function of Visual Short-Term Memory: Transsaccadic Memory, Object Correspondence, and Gaze Correction. *Journal of Experimental Psychology: General*, 137(1), 163–181.
<https://doi.org/10.1037/0096-3445.137.1.163>

Huang, L. (2010). Visual working memory is better characterized as a distributed resource rather than discrete slots. *Journal of Vision*, 10(14), 8–16. <https://doi.org/10.1167/10.14.8>

Jiang, Y., Olson, I. R., & Chun, M. M. (2000). Organization of Visual Short-Term Memory.

- Journal of Experimental Psychology: Learning Memory and Cognition*, 26(3), 683–702.
<https://doi.org/10.1037/0278-7393.26.3.683>
- Libertus, M. E., Feigenson, L., & Halberda, J. (2011). Preschool acuity of the approximate number system correlates with school math ability. *Developmental Science*, 14(6), 1292–1300. <https://doi.org/10.1111/j.1467-7687.2011.01080.x>
- Luck, S. J., & Hollingworth, A. (2008). Visual memory. In *Visual memory*. Oxford University Press.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390(6657), 279–284. <https://doi.org/10.1038/36846>
- Luck, S. J., & Vogel, E. K. (2013). Visual working memory capacity: from psychophysics and neurobiology to individual differences. *Trends in Cognitive Sciences*, 17(8), 391–400. <https://doi.org/10.1016/j.tics.2013.06.006>
- Lund, J., Parkes, L., Solomon, J. A., Angelucci, A., & Morgan, M. (2002). Compulsory averaging of crowded orientation signals in human vision. *Nature Neuroscience*, 4(7), 739–744. <https://doi.org/10.1038/89532>
- Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, 17(3), 347–356. <https://doi.org/10.1038/nn.3655>
- Nie, Q.-Y., Müller, H. J., & Conci, M. (2017). Hierarchical organization in visual working memory: From global ensemble to individual object structure. *Cognition*, 159, 85–96. <https://doi.org/10.1016/j.cognition.2016.11.009>
- Nieuwenhuis, S., Forstmann, B. U., & Wagenmakers, E. J. (2011). Erroneous analyses of interactions in neuroscience: A problem of significance. *Nature Neuroscience*. <https://doi.org/10.1038/nn.2886>

- Olsson, H., & Poom, L. (2005). Visual memory needs categories. *Proceedings of the National Academy of Sciences*, 102(24), 8776–8780. <https://doi.org/10.1073/PNAS.0500810102>
- Pashler, H. (1988). Familiarity and visual change detection. *Perception & Psychophysics*, 44(4), 369–378. <https://doi.org/10.3758/BF03210419>
- Phillips, W. A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception & Psychophysics*, 16(2), 283–290. <https://doi.org/10.3758/BF03203943>
- Pomplun, M., Reingold, E. M., & Shen, J. (2001). Peripheral and parafoveal cueing and masking effects on saccadic selectivity in a gaze-contingent window paradigm. *Vision Research*, 41(21), 2757–2769. [https://doi.org/10.1016/S0042-6989\(01\)00145-6](https://doi.org/10.1016/S0042-6989(01)00145-6)
- Pratte, M. S., Park, Y. E., Rademaker, R. L., & Tong, F. (2017). Accounting for stimulus-specific variation in precision reveals a discrete capacity limit in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 43(1), 6–17. <https://doi.org/10.1037/xhp0000302>
- Schurgin, M. W., & Brady, T. F. (2019). When “capacity” changes with set size: Ensemble representations support the detection of across-category changes in visual working memory. *Journal of Vision*, 19(5), 3. <https://doi.org/10.1167/19.5.3>
- Scolari, M., Vogel, E. K., & Awh, E. (2008). Perceptual expertise enhances the resolution but not the number of representations in working memory. *Psychonomic Bulletin and Review*, 15(1), 215–222. <https://doi.org/10.3758/PBR.15.1.215>
- Suchow, J. W., Brady, T. F., Fougny, D., & Alvarez, G. A. (2013). Modeling visual working memory with the MemToolbox. *Journal of Vision*, 13(10), 9–9. <https://doi.org/10.1167/13.10.9>

- Vidal, J. R., Gauchou, H. L., Tallon-Baudry, C., & O'Regan, J. K. (2005). Relational information in visual short-term memory: The structural gist. *Journal of Vision*, 5(3), 244–256. <https://doi.org/10.1167/5.3.8>
- Watanabe, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, 11, 3571–3594.
- Watanabe, S. (2013). A widely applicable bayesian information criterion. *Journal of Machine Learning Research*, 14(1), 867–897.
- Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, 4(12), 11. <https://doi.org/10.1167/4.12.11>
- Xie, W., & Zhang, W. (2017). Discrete item-based and continuous configural representations in visual short-term memory. *Visual Cognition*, 25(1–3), 21–33. <https://doi.org/10.1080/13506285.2017.1339157>
- Xu, Y., & Chun, M. M. (2006). Dissociable neural mechanisms supporting visual short-term memory for objects. *Nature*, 440(7080), 91–95. <https://doi.org/10.1038/nature04262>
- Xu, Y., & Chun, M. M. (2009). Selecting and perceiving multiple visual objects. *Trends in Cognitive Sciences*, 13(4), 167–174. <https://doi.org/10.1016/j.tics.2009.01.008>
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, 453(7192), 233–235. <https://doi.org/10.1038/nature06860>
- Zhang, W., & Luck, S. J. (2011). The number and quality of representations in working memory. *Psychological Science*, 22(11), 1434–1441. <https://doi.org/10.1177/0956797611417006>

Appendix A

In this appendix, we describe two sets of Bayesian hierarchical models that fit the ensemble and the item representation performance in Experiment 4. The ensemble accuracy was fit with three models ($Model_{ensemble-simple}$, $Model_{ensemble-layout}$ and $Model_{ensemble-set\ size}$), separately. For $Model_{ensemble-simple}$, we presumed two unknown parameters, α and β , each of which was given an uninformative hyperprior information:

$$\alpha, \beta \sim \text{Gamma}(0.001, 0.001)$$

Then, we presumed an expected ensemble accuracy, A_i , for each individual participant. $A_1, A_2, \dots, A_n$ ($n = \text{participant number}$) all followed a beta distribution with the parameters α and β :

$$A_1, A_2, \dots, A_n \sim \text{Beta}(\alpha, \beta)$$

We also calculated an expected grand ensemble accuracy, A_G , using the expected value of the fitted beta distribution:

$$A_G = \frac{\alpha}{\alpha + \beta}$$

Then, the responses of each participant in the ensemble task followed a Bernoulli distribution with the parameter A_i :

$$\sum_{j=1}^{n_i} \text{resp}_{ij} \sim B(n_i, A_i)$$

where resp_{ij} is a binary variable which represents the response of *Participant i*, *Trial j*, and n_i represents the trial number of *Participant i*.

The model construction and fitting of $Model_{ensemble-layout}$ and $Model_{ensemble-set\ size}$ were the same as ones of $Model_{ensemble-simple}$ except the following differences:

For $Model_{ensemble-layout}$, we presumed an expected ensemble accuracy of the hard layout condition ($A_{1-hard}, A_{2-hard}, \dots, A_{n-hard}$), rather than collapsed across the layout conditions, for each

participant that follows the beta distribution with the free parameters α and β . We also presumed two additional unknown parameters, $\Delta A_{hard \rightarrow easy}$ and $\Delta A_{hard \rightarrow medium}$, that represented the expected ensemble accuracy in the categorical judgement task to accommodate the difference between the hard layout condition and the easy/medium layout conditions. Responses followed a Bernoulli distribution with the parameters of A_{i-hard} and $\Delta A_{hard \rightarrow easy}$ (or A_{i-hard} and $\Delta A_{hard \rightarrow medium}$):

$$\sum_{j=1}^{n_{i-easy}} resp_{i-easy\ j} \sim B(n_{i-easy}, A_{i-hard} + \Delta A_{hard \rightarrow easy})$$

$$\sum_{j=1}^{n_{i-medium}} resp_{i-medium\ j} \sim B(n_{i-medium}, A_{i-hard} + \Delta A_{hard \rightarrow medium})$$

$$\sum_{j=1}^{n_{i-hard}} resp_{i-hard\ j} \sim B(n_{i-hard}, A_{i-hard})$$

where $resp_{i-easy\ j} / resp_{i-medium\ j} / resp_{i-hard\ j}$ represents the response of *Participant i*, *Trial j* in the easy/medium/hard layout condition and $n_{i-easy} / n_{i-medium} / n_{i-hard}$ represents the trial number of *Participant i* in the easy/medium/hard layout condition.

Likewise, we assumed expected ensemble accuracy of the set size 8 condition ($A_{1-8}, A_{2-8}, \dots, A_{n-8}$, n = participant number) of each participant for *Model_{ensemble-set size}*. We also used a parameter $\Delta A_{8 \rightarrow 6}$ to represent the expected ensemble representation accuracy change from set size 8 to 6.

Responses followed a Bernoulli distribution with the parameters of A_{i-8} and $\Delta A_{8 \rightarrow 6}$):

$$\sum_{j=1}^{n_{i-6}} resp_{i-6\ j} \sim B(n_{i-6}, A_{i-8} + \Delta A_{8 \rightarrow 6})$$

$$\sum_{j=1}^{n_{i-8}} resp_{i-8\ j} \sim B(n_{i-8}, A_{i-8})$$

where $resp_{i-6j} / resp_{i-8j}$ represents the response of *Participant i*, *Trial j* in the condition of set size 6/8 and n_{i-6}/n_{i-8} represents the trial number of *Participant i* in the condition of set size 6/8.

The item representation performance was fit with another set of three models (*Model_{individual-simple}*, *Model_{individual-layout}* and *Model_{individual-set size}*), separately. *Model_{individual-simple}*, was a hierarchical mixture model and predicted the response errors of the item representation task according to a von Mises distribution with the parameters of κ_{ij} and z_{ijk} :

$$e_{ijk} \sim VM(0, \kappa_{ij} \bullet z_{ijk})$$

where e_{ijk} represents the response error of *Participant i*, *Trial k* in the *j*th stimulus condition (1: color; 2: orientation). z_{ijk} was a latent binary represents whether *Participant i* held the representation of the probed item of *Trial k* or not (held: $z_{ijk} = 1$; not held: $z_{ijk} = 0$). κ_{ij} was always larger than 0 represents the precision of the response error if the representation of the probed item was held. When $z_{ijk} = 0$, the distribution was a uniform distribution on the circle and represented a totally random performance; when $z_{ijk} \neq 0$, the response errors followed a von Mises distribution center on 0 and with precision of κ_{ij} .

κ_{ij} was a variable at the individual participant level and followed a Gamma distribution with the unknown parameters of $\alpha_{gamma-i}$ and $\beta_{gamma-i}$ in the *i*th stimulus condition:

$$\kappa_{ij} \sim Gamma(\alpha_{gamma-j}, \beta_{gamma-j})$$

where $\alpha_{gamma-j}$ and $\beta_{gamma-j}$ both followed an uninformative hyperprior Gamma distribution:

$$\alpha_{gamma-j}, \beta_{gamma-j} \sim Gamma(0.001, 0.001)$$

z_{ijk} followed a Bernoulli distribution with a parameter of P_{m-i} :

$$z_{ijk} \sim B(P_{m-ij})$$

where P_{m-ij} represents the probability that *Participant i* held the representation of the probed stimuli in *j*th stimulus condition. P_{m-ij} followed a Beta distribution with the unknown parameters of α_{beta-i} and β_{beta-i} in the *i*th stimulus condition:

$$P_{m-ij} \sim \text{Beta}(\alpha_{beta-j}, \beta_{beta-j})$$

where α_{beta-j} and β_{beta-j} both followed an uninformative hyperprior Gamma distribution:

$$\alpha_{beta-j}, \beta_{beta-j} \sim \text{Gamma}(0.001, 0.001)$$

$Model_{individual-layout}$ and $Model_{individual-set\ size}$ had the similar structures with $Model_{individual-simple}$ except that additional parameters were introduced to represent the grand effects of layout and set size. For $Model_{individual-layout}$, we presumed two unknown parameters, $\kappa_{ij-hard}$ and $P_{m-ij-hard}$, for each participant for the hard layout condition, rather than collapsed across the layout conditions. Then, we presumed $\Delta\kappa_{j-hard \rightarrow easy} / \Delta\kappa_{j-hard \rightarrow medium}$ and $\Delta P_{m-j-hard \rightarrow easy} / \Delta P_{m-j-hard \rightarrow medium}$ to represent the layout effect change from the hard layout condition to the easy / medium layout condition. So the response errors from different layout conditions followed different hierarchical models:

$$e_{ijk-easy} \sim VM(0, (\kappa_{ij-hard} - \Delta\kappa_{j-hard \rightarrow easy}) \bullet z_{ijk-easy})$$

$$z_{ijk-easy} \sim B(P_{m-ij-hard} + P_{m-j-hard \rightarrow easy})$$

$$e_{ijk-medium} \sim VM(0, (\kappa_{ij-hard} - \Delta\kappa_{j-hard \rightarrow medium}) \bullet z_{ijk-medium})$$

$$z_{ijk-medium} \sim B(P_{m-ij-hard} + P_{m-j-hard \rightarrow medium})$$

$$e_{ijk-hard} \sim VM(0, \kappa_{ij} \bullet z_{ijk-hard})$$

$$z_{ijk-hard} \sim B(P_{m-ij-hard})$$

Likewise, we presumed two parameters, κ_{ij-8} and P_{m-ij-8} , for each participant at set size 8 for $Model_{individual-set\ size}$. We then presumed $\Delta\kappa_{j-8 \rightarrow 6}$ and $\Delta P_{m-j-8 \rightarrow 6}$ to represent the set size effect change

from the condition of set size 8 to one of set size 6. Response errors from different set size conditions followed different hierarchical models:

$$e_{ijk-6} \sim VM(0, (\kappa_{ij-6} - \Delta\kappa_{j-8 \rightarrow 6}) \bullet z_{ijk-6})$$

$$z_{ijk-6} \sim B(P_{m-ij-6} + P_{m-j-8 \rightarrow 6})$$

$$e_{ijk-8} \sim VM(0, \kappa_{ij} \bullet z_{ijk-8})$$

$$z_{ijk-8} \sim B(P_{m-ij-8})$$

The posterior samples of all the unknown parameters were generated using Gibb's sampling method and the package of JAGS (Plummer, 2003). For each model, one chain was generated, which had 100,000 samples and thinned by a factor of 10 to de-correlate the samples, resulting in 10,000 final samples of the posterior distribution in total.

Appendix B

This appendix describes a Bayesian hierarchical model that fits the ensemble representation accuracy across Experiments 2, 3 and 4. This model had a similar structure with *Model_{ensemble-set size}* described in Appendix A. For each participant in each experiment, we presumed that the ensemble accuracy at set size 8 ($A_{1-8}, A_{2-8}, \dots, A_{n-8}$, n = participant number across Experiments 2, 3 and 4) follows a Beta distribution with unknown parameters, α, β , which were given uninformative hyperprior Gamma distribution for model fitting:

$$A_{1-8}, A_{2-8}, \dots, A_{n-8} \sim \text{Beta}(\alpha, \beta)$$

$$\alpha, \beta \sim \text{Gamma}(0.001, 0.001)$$

At each other set size (i.e. set size = 2, 4 or 8), one additional parameter (i.e. $\Delta A_{8 \rightarrow 2}$ / $\Delta A_{8 \rightarrow 4}$ / $\Delta A_{8 \rightarrow 6}$) describes the expected difference in ensemble accuracy between this particular set size and set size 8. Thus, for set size s ($s < 8$), responses followed a Bernoulli distribution with the parameters of A_{i-8} and $\Delta A_{8 \rightarrow s}$:

$$\sum_{j=1}^{n_{i-s}} \text{resp}_{i-s j} \sim B(n_{i-s}, A_{i-8} + \Delta A_{8 \rightarrow s})$$

where $\text{resp}_{i-s j}$ represents the response of *Participant i* in *Trial j* at set size s , and n_{i-s} represents the trial number of *Participant i* at set size s . The responses at set size 8 follows another Bernoulli distribution with the parameter of A_{i-8} only:

$$\sum_{j=1}^{n_{i-8}} \text{resp}_{i-8 j} \sim B(n_{i-8}, A_{i-8})$$

Note that there was no set size 8 condition in Experiment 2, so the fitted A_{i-8} of the participants from Experiment 2 can only denote hypothetical meanings. In fact, the lack of set size 8 did not influence the model fitting and the evaluation of $\Delta A_{8 \rightarrow 2}$ / $\Delta A_{8 \rightarrow 4}$ / $\Delta A_{8 \rightarrow 6}$, for

Experiments 2-4 sharing the same experiment procedures so we could safely assume that all A_{i-8} (including the hypothetical ones in Experiment 2) share same posterior distribution

As a reference model for comparison, we also constructed the $Model_{ensemble-simple}$ which does not include set size as a predictive factor for ensemble accuracy and the structure of the model was the same as $Model_{ensemble-simple}$ in Experiment 4, which is described in Appendix A.

For model fitting, the posterior samples of all the unknown parameters were generated using Gibb's sampling method and the package of JAGS (Plummer, 2003). For each model, one chain was generated, which had 100,000 samples and was thinned by a factor of 10 to de-correlate the samples, resulting in 10,000 final samples of the posterior distribution in total.