

Structured Visuospatial Representations Revealed through Serial Reproduction

Yang Wang, Timothy F. Lew, Timothy F Brady, Edward Vul

Department of Psychology, University of California, San Diego

(Word count with exclusions: 5614)

Author Note:

Data, analysis code, and study materials are available at github repository (<https://github.com/yaw001/Serial-Reproduction>) and OSF (The study is not pre-registered). Correspondence concerning this article should be addressed to Yang Wang. Email: yaw001@ucsd.edu.

Abstract

Working memory is a reconstructive process that requires integrating multiple hierarchical representations of objects. This hierarchical reconstruction allows us to overcome perceptual uncertainty and limited cognitive capacity, but yields systematic biases in working memory as individual items are influenced by the ensemble statistics of the scene, or of their particular group. Given the importance of the hierarchical encoding of a display, we aim to characterize what structures people use to encode visual scenes using a nonparametric data-driven approach. In Experiment 1, we examine visuospatial memory for locations by asking participants to recall the locations of objects in a serial reproduction task. We show that people report items in a more compact structure than they initially were, and organize them into clustered spatial groups. In Experiment 2, we explicitly introduce discrete color groups, allowing us to test whether the color feature governs the spatial grouping. We find that the spatial structures were color-contingent. By analyzing color groups, we circumvent the grouping uncertainty in Experiment 1 and further reveal that people compress color groups into collinear structures with similar orientations and equidistant spacing.

Keywords: Visuospatial memory, serial reproduction, ensemble coding, memory biases

Public Significance Statement: What we perceive and remember is a result of both what we sense, and what we expect to be in the world. Here we use serial reproduction – a formalized version of the telephone game – to characterize what spatial arrangements people expect in visual scenes, and thus what kinds of errors they make when trying to remember a particular display.

Introduction

Although working memory possesses an extremely limited capacity for information (Cowan, 2001, 2005; Miller, 1956), it can make efficient use of its limited resources by exploiting statistical structure in the visual world to aid recall (Alvarez, 2011; Brady & Alvarez, 2011; Ariely, 2001; Orhan, et al., 2014; Brady, Konkle & Alvarez, 2009; Sims, Jacobs, & Knill, 2012). For example, an observer trying to remember the locations of people in a crowd might infer that individuals are organized into groups. Later on, the observer might have forgotten people's exact locations and compensate by remembering individuals' locations biased towards their group centers (Lew & Vul, 2015). Although relying on objects' ensemble statistical structures bias the memory of these objects, it can gain precision in the representation of the ensemble, compensating for noise in local feature representations (Alvarez & Oliva, 2009). Furthermore, encoding objects according to their statistical structure constrains the possible properties of those objects, allowing observers to remember the objects' exact features more precisely (Sims, Jacobs & Knill, 2012; Orhan, et al., 2014). For example, inferring that a set of objects generally fall on a horizontal line constrains their y-coordinates. This allows the observer to focus on encoding their x-coordinates with greater precision. However, the effectiveness of an encoding scheme depends on how well it matches the statistics of a stimulus (Orhan & Jacobs, 2014a). Consequently, when people's expectations about statistical structures fail to match what they observe, the fidelity of visual working memory will suffer. Orhan & Jacobs (2014b), for example, found that in a typical study of capacity, when the subject's visual systems' built-in expectations about regular line configuration (i.e., line segments with similar orientations or form continuous lines) mismatches with stimuli that have uniformly distributed features, the mismatch can detrimentally bias memory and potentially explain a significant portion of performance limitations. Therefore, how

people use the structure of displays to help encode visual information depends on what structures they have available in their visual system (e.g., object templates, like letters; oriented lines, etc) to encode objects in displays.

Typically, researchers examine what spatial organizations people encode by designing stimuli that test whether people use specific grouping strategies to facilitate task performance. In perception, this approach has allowed psychologists to identify a host of Gestalt grouping principles (Wertheimer, 1923). For example, visual perception processing is facilitated by grouping elements that are near each other (the principle of proximity) or are similar to each other (the principle of similarity) and perceiving continuous lines over segmented pieces (the principle of continuity). Furthermore, previous findings demonstrate that Gestalt grouping principles (i.e., Connectedness, common region, spatial proximity and similarity) facilitate visual working memory performance in change detection tasks (Woodman, Vecera & Luck, 2003; Xu, 2002, 2006; Xu & Chun, 2007; Jiang, Olson & Chun, 2000; Peterson & Berryhill, 2013).

Building upon these findings, we adopted a data-driven design to discover the grouping structures that people expect by virtue of the memory biases that arise in a serial reproduction paradigm.

Rather than testing whether people possess particular structures whose motivations are themselves poorly understood, this allows us to discover a wide variety of structures without specifying them in advance. In the current study, we had participants reveal their grouping expectations by performing a task similar to a game of Telephone: each participant studied and recalled the locations of objects and then the next participant studied and recalled the previous participant's responses, and so on. The serial reproduction process amplifies shared biases by aggregating systematic error and the serial reproduction chains will eventually converge toward responses which are more in line with participants "priors" – that is, which are more easily represented and reproduced by participants (Bartlett, 1932; Griffiths & Kalish, 2005; Sanborn &

Griffiths, 2007; Kirby et.al., 2008; Xu & Griffiths, 2010; Langlois et al., 2021) — in our case, the structural organizations that participants expected. One view of how this works is a Bayesian point of view on memory: peoples' reproduction reflects peoples' posterior, and that this posterior combines the sensory percept and their particular “priors” that might be relevant for a given memoranda. Thus, at every iteration, participants' posterior – upon which they base their response – will be influenced by the data and the relevant “prior”, leading to biases. Over iterations, priors that are in common across participants – but not those that are not – will be increasingly reflected in the responses (as the prior influences each iteration, and only does so in the “same direction” for each participant if it is in common between participants). Indeed, previous work has shown that people's reproductions are a combination of learned priors and their memory for the specific instance they are trying to recreate, leading their memory to be slightly pulled toward their prior when noise is present (e.g., Huttenlocher et al. 2001; Hemmer & Steyvers, 2009). When this process is repeated over and over, as in the case of serial reproduction, this results in the chains moving to situations that more and more closely resemble priors that are in common across participants. Thus, the serial reproduction paradigm emphasizes and exaggerates inductive biases through the process of repeated noisy reproduction, even if the serial reproduction chains do not perfectly converge to a single unique arrangement, they allow us to examine the structural characteristics that gradually surface over iterations.

In Experiment 1, we examine the structures people use to remember locations in working memory by asking participants to recall the locations of objects in a serial reproduction task. We show that people tend to misremember items toward a more globally compact structure, and organize them into clustered spatial groups. In Experiment 2, we explicitly introduce discrete color groups, allowing us to test whether the color feature governs the spatial grouping. We find that the spatial structures were color-contingent. By analyzing color groups, we circumvent the

grouping uncertainty in Experiment 1 and further reveal that people tend to compress color groups into dot line segments with similar orientations and equidistant spacing.

Experiments 1A&1B

Rather than designing predefined structural regularities that people possibly capitalize on, we took a data-driven approach to reveal the structures people use to encode locations using a serial reproduction paradigm.

In Experiment 1A, participants briefly saw a set of homogenous dots on a computer screen and clicked on the screen to recall where the circles had been after a short delay. Critically, the locations one participant reported were shown as the stimulus to the next participant, thus producing a serial reproduction chain. Experiment 1B was a perceptuomotor control experiment with the same serial reproduction structure except the stimulus was available throughout the reproduction of the dots. The perceptuomotor experiment allows us to validate whether any structural biases arose due to biases in visual memory or motor planning with perceptual noise.

Method

Stimuli

For the memory task, the subjects' goal was to remember the location of 15 grey dots within a circular display area. The display area had a radius of 275 pixels, while each of the dots had a radius of 10 px. In the first iteration of each chain, the locations of the dots were randomly drawn from a uniform distribution over the display area, subject to the constraint that the dots could not overlap. Perceptuomotor control task has a similar stimulus display. The size and resolution of

observers' computer monitors were not controlled. However, with an exclusion criterion, the goal was recording acceptable recalls within subjects rather than comparing the recall performances between subjects.

Procedure

In Experiment 1A, participants first performed a randomly generated practice trial to familiarize themselves with the task. The second trial was our main test in which participants saw locations of either a seed display (i.e., randomly distributed objects) of the chain or locations reported by the previous participant in the chain). Participants were not told that the stimuli they studied were another participant's responses and they were ignorant about being in a serial reproduction chain. In the third trial, participants studied the initial seed given the chain they were assigned (participants who received the seed display in the test trial would perform on the same seed display twice). The fourth trial was a randomly generated performance check: if a participant's score was below criterion on this test, their responses were not included in the serial reproduction chain to prevent a single inattentive subject from corrupting an entire chain.

On each trial (See Figure 1), participants observed the locations of the circles for 10 seconds, followed by a 1 second mask. Participants then recalled the locations of the circles by clicking the mouse. Participants had unlimited time to recall the locations of the circles and could move them by dragging the circles as much as they wanted. Once participants indicated that they were done reporting the locations (by pressing Enter), we gave them feedback by showing the correct and recalled locations along with lines indicating how far off they were. We determined the mapping between guesses and targets using a greedy search that minimized root mean square error (RMSE). Participants also received a score between 0 and 100 based on the average

distance between the recalls and targets normalized by the standard deviation of object locations.

To motivate participants, they were instructed that their final bonus would reflect their scores.

For the Experiment 1B (perceptuomotor control task) (See Figure 1), participants saw two environments side-by-side. The left environment contained the circles in the target locations and remained onscreen for the entire trial. The right environment was empty and participants were instructed to copy the locations from the left environment onto the right environment. Once the participant finished, they received feedback in the right environment using the same criteria as in the memory task.

Figure 1. Experiment trial. Memory trials (top row): Participants saw 15 grey dots for 10 seconds followed by a one-second mask. Participants then recalled the locations of all the dots and were told how many dots they had to recall. Participants could move around the dots until they were satisfied. Participants then saw the correct object locations (grey) and their recalled locations (red) and the mapping between the targets and their recalls (black lines). Perceptuomotor trials (bottom row): Participants saw a display with 15 grey dots and a blank display next to it and they were told to copy the dots' locations on the blank display. They received feedback after they were satisfied with their response.

Design

For the memory experiment (Experiment 1A), we showed subjects a display of 15 dots, and asked them to report the locations of all 15 dots after a brief mask interval. The positions one subject reported were presented as the stimulus to another subject. This process was repeated with 20 unique subjects, thus yielding a serial reproduction “chain” of 20 iterations.

The control experiment (Experiment 1B) was designed to elicit only perceptuomotor errors. The structure of the perceptuomotor task was similar to the memory task—participants studied and reported the locations of objects and their responses were passed on to the next participant.

However, instead of briefly studying and then recalling the objects, participants had access to the display they were instructed to reconstruct the entire time.

For both experiments, we generated 10 “seed” displays, each with 15 dots placed randomly in the display area. We set up 10 chains for each seed display and then ran each chain for 20 iterations. Thus, there were a total of 100 chains, each consisting of 20 iterations. Figure 2 shows some of the typical examples of serial reproduction chains.

Figure 2. Four example chains (rows) for the seed display, 1st, 5th, 10th, 15th, and 20th iterations (columns) from the memory experiment and the perceptuomotor experiment. For the memory experiment (Experiment 1A), despite objects being initially uniformly distributed in the displays, participants gradually organized them into complex organized structures. In contrast, for the perceptuomotor experiment (Experiment 1B), the reproduction preserves the overall pattern of the original seed.

Participants

For the memory task, we gathered participants from the Amazon Mechanical Turk and rewarded participants with a base payment and a performance-based bonus. We allowed participants to perform multiple trials of our experiment for different initial displays, resulting in 1614 unique participants performing a total of 2000 experiment runs. For the perceptuomotor control experiment, we gathered 1399 unique participants from the Amazon Mechanical Turk marketplace with a base payment.

Transparency and openness

All data, analysis code, and research are available at the github repository <https://github.com/yaw001/Serial-Reproduction>. The data from Experiment 1 were collected in 2015 and the data from Experiment 2 were collected in 2018. All data were analyzed using R. The study's design and its analysis were not pre-registered.

Results

Did participants' responses drift across iterations?

Figure 2 shows examples of serial reproduction chains evolving in the memory experiment and the perceptuomotor experiment. To quantitatively determine whether participants' responses drift, we measured the recall error distance using the mean absolute error (MAE) between the objects in a given iteration and the objects matched by the Hungarian algorithm (Kuhn, 1955) in the original seed. If the responses become increasingly dissimilar to the original seed, we would expect the MAE to increase across the chains. Figure 3 shows the MAE averaging across all chains for a given iteration. We found that the MAE averaging across all chains for a given seed display increased over iterations in memory ($b = 0.021$, 95% CI [0.020, 0.022]) and perceptuomotor ($b = 0.0155$, 95% CI [0.015, 0.016]) experiments indicating that patterns gradually diverged from the original seeds. However, the divergence in the memory experiment cannot be accounted for by perceptuomotor errors. The errors introduced in the memory experiment are significantly greater than those in the perceptuomotor experiment (Welch t-test: $t(3505.1) = 25$, $p < 0.001$, $d = 0.8$).

Figure 3. Mean MAE: The mean distance between the objects and the matched locations in the initial seed display averaging across 100 chains given each iteration. The red dots indicate the mean error distance in the memory experiment and the blue dots indicate the mean error distance

in the perceptuomotor experiment. Participants' responses initially resembled the seed displays, but became increasingly dissimilar over time. Error bars indicate SEM.

Altogether, participants appeared to introduce small errors, resulting in the locations of objects drifting over time and becoming increasingly dissimilar to the initial seed. Moreover, the errors are much smaller with direct reproduction than with memory recall, indicating that visual memory errors cannot be fully explained by perceptuomotor errors and the accumulating divergence from the random seed in the memory experiment must stem from systematic biases in visuospatial memory.

Did participants' patterns of responses converge over iterations?

We showed that the participants' responses drift across the serial reproduction chains and the final iterations significantly diverge from the random seeds. To examine what gives rise to this divergence, we sought to verify that the recalled patterns' gradual drift is towards relatively stable structures, as we would expect if people share a set of structures they use to encode visual displays (what in the Bayesian conception of memory would be considered their shared priors). To quantitatively determine the convergence of the recalled patterns, we measured one-back MAE which is the average recalled error distance of the current iteration to the immediately preceding iteration (e.g., 1st iteration vs. the random seed; 20th iteration vs. 19th iteration). A trend toward convergence in the memory experiment would result if earlier iterations tend to have larger drift from their immediately preceding iteration and these errors gradually decrease closer to the perceptuomotor errors towards the end of chains. This would be indicative of the structures becoming more stable over time.

Figure 4 shows the MAE of the iteration to its one-back iteration. We found that one-back MAE significantly decreases over iterations ($b = -0.0034$, 95% CI [-0.004, -0.0029]). Furthermore, the

one-back MAE for the last iteration is significantly closer to the average perceptuomotor one-back MAE than the first iteration in the memory experiment (Welch t-test: $t(185.18) = 6.2$, $p < 0.001$, $d = 0.88$) indicating large recall errors were driven by the divergence of the dots from the structures people use to encode visual displays in the early iterations, and the errors approach to perceptuomotor errors towards the end of the serial reproduction chains as the chains approach more stable configurations that are more congruent with the structures people use to represent dots. Altogether, the results suggest that the serial reproduction chains converged to relatively stable configurations. Thus, after participants' biases influenced the displays, the displays were easier for subsequent participants to accurately remember. This is consistent with the fact that structures emerged that subsequent participants could use to help encode the display more accurately. In the Bayesian view of memory, this suggests the displays have become more in line with participants' visuospatial priors.

Figure 4. One-back MAE: the MAE between a given iteration and its immediately preceding iteration. Red dots represent the memory experiment and blue dots represent the Perceptuomotor experiment. The one-back MAE significantly decreases approaching the average perceptuomotor one-back MAE indicating the serial reproduction chains converge towards stable configurations. Error bars indicate SEM.

What characteristics of patterns emerge across serial reproduction chains?

While each chain seems to converge to a unique configuration, some general characteristics of configurations manifest themselves across chains as they diverge from the initial seeds. In this section, we quantitatively measure pattern dispersion and clusterability.

Pattern Dispersion

To quantitatively measure the dispersion of a display, we calculated the mean determinant of the covariance matrix for each iteration aggregating across all chains. A smaller mean covariance matrix determinant indicates a more compact (and lower entropy) structure. The serial reproduction chains in the visual memory experiment tend to converge towards significantly more compact structures ($b = -0.19$, 95% CI [-0.20,-0.18]), while those in the perceptuomotor control experiment tend to drift towards slightly more dispersed patterns ($b = 0.028$, 95% CI [0.024, 0.032]) with much smaller absolute slope (See Figure 5) (Welch t-test for the difference of average deviation between memory and perceptuomotor experiments: $t(2287.2) = -49.7$, $p < 0.001$, $d = -1.57$).

Figure 5. Dispersion of patterns in memory and perceptuomotor experiments. The dispersion of a display is measured by the determinant of the covariance matrix of the object locations. The mean determinant of the covariance matrix is computed for each iteration aggregating across all chains given an iteration. For perceptuomotor experiments (blue dots), the dispersion of the patterns across chains increases slightly over iterations. In contrast, the dispersion significantly decreases in the visual memory experiment (red dots) suggesting that configural representations underlie a convergence toward more compact patterns.

Clusterability

In the memory experiment, patterns not only compacted, but also tended to be clustered. To examine the clusterability of patterns, we applied a measure using the Dip test on pairwise distances (dip-dist test) (Adolfsson et al., 2018). Multiple modes in the distribution of pairwise distances implies the presence of clusters (See Figure 6A), because within-cluster pairs will be close together, while between-cluster pairs will be far apart. Thus, to test for spatial clustering in our displays, we calculate the set of pairwise Euclidean distances in a given display, and use them as inputs into the distribution Dip test (procedure detailed in Hartigan, 1985). We found that the proportion of displays that had significant ($p < 0.05$) clustering based on the Dip test on pairwise distance distributions increased over iterations in serial reproduction chains ($b = 0.011$,

95% CI [0.009,0.014]) in the memory experiment (See Figure 6B). Although there was a slight increase of clustering in the perceptuomotor chains ($b = 0.0009$, 95% [0.0004, 0.0014]), the significant clustering tendency in the memory experiment could not be fully attributed to the perceptuomotor biases (paired-t test: $t(19) = 8.2$, $p < 0.001$, $d = 2.47$).

Figure 6. Analysis of clusterability. (A) Distance dip test. Top row: An unclustered arrangement of dots and its pairwise distance distribution is unimodal. Bottom row: A clustered arrangement of dots and its pairwise distance distribution is bimodal. (B) Proportion of significantly clustered reports. The proportion of clustered iterations increases across serial reproduction chains in the memory experiment while very few arise in the perceptuomotor experiment.

Discussion

We set out to characterize the structures people use to encode visuospatial displays using a serial reproduction paradigm in which participants had to reproduce a spatial arrangement of 15 homogenous dots on a screen after a delay. As intended, the serial reproduction chains reveal the structured memory biases: the spatial arrangements diverge from the original seed at a greater rate when each reproduction trial requires a memory delay. Not only is the rate of divergence greater for reproduction from memory, but also memory reproduction introduces systematic biases into the structure of the displays. Reproductions in the Memory experiment, but not in the Perceptuomotor experiment, yield spatial arrangements with more compact structures over iterations. Such a tendency is consistent with the view of an adaptive bias that compensates for the spatial memory uncertainty by reducing the magnitude of encoded relative distances between objects (Lew & Vul, 2015) and potentially increases the average accuracy of the stimulus reproduction (Huttenlocher and Hedges, 2000). Moreover, we find that memory, but not perceptuomotor reproductions, tend toward clustered groups. In contrast to the traditional assumption that objects in visual memory are encoded independently, the result indicates that

people appear to encode higher-order structures of objects and utilize such structured memory to aid recall (Brady & Alvarez, 2011).

Although the overall trend toward clustering of spatial positions in memory reproduction is robust, it fails to capture richer grouping organizations to which the memory serial reproduction chains converge. In particular, some displays converge to rich structures, and these are sometimes part of separate clusters (Figure 7). While some aspects of the within-group structure might be estimated via data-driven clustering, it is challenging to analyze this data in a cluster-by-cluster manner when there is no direct way to assign different items to different clusters *a priori*. To address this, we attempted to better characterize these higher order group structures without introducing our own within-cluster distribution assumptions by explicitly adding a grouping cue. In particular, in Experiment 2, we add colors to the dots, to induce color-specific grouping, thus allowing us to analyze within-group and between-group structural biases.

Figure 7. Three examples of structures. Response “Seed 8, Chain 3, Iteration 20”: The mean nearest neighbor clustering analysis is explicitly aimed at capturing such unambiguous clustered grouping structure. Response “Seed 2, Chain 10, Iteration, 20” and Response “Seed 5, Chain 10, Iteration 20”: Two examples of grouping structures cannot be captured by mean nearest neighbor clustering analysis despite there exists possible grouping structures.

Experiment 2

In Experiment 2, we explicitly introduced a grouping cue: color. We aim to test two questions: (1) Do the color cues serve a strong grouping cue that induce color-specific spatial biases in serial reproduction chains? (2) If (1) were true, would different structures be used to represent the within-group and between-group spatial arrangement?

Method

Stimuli

The stimuli were similar to the ones in Experiment 1 except that the 15 dots were randomly assigned with red, green and blue colors (i.e., 5 red dots, 5 green dots and 5 blue dots).

Procedure

All observers were presented with 6 trials: 3 practice trials and 3 test trials. The three practice trials consisted of 3 dots, 9 dots and 15 dots. The 3 test trials consisted of 15 dots. The goal was to memorize, and subsequently recall, the configuration of colored dots. In each trial (See Figure 8), participants observed the locations of specified numbers of dots for 15 seconds, followed by a 1-second mask. Participants then recalled the locations of the dots by clicking the mouse and moving the dots to the recalled position. Participants had unlimited attempts to recall the locations of the colored dots and drag them at their discretion. Once participants indicated that they were done reporting the locations (by pressing enter), we gave them feedback by showing the correct and recalled locations along with connected lines indicating the error. We determined the mapping between color-matched guesses and targets using a greedy search that minimized mean absolute error (MAE). We assumed the empirical measure of chance performance was $MAE = 0.4$ and iterations with MAE greater than 0.4 were excluded.

Figure 8. Experimental Trial: Participants saw 15 dots with three different colors of equal number for 15 seconds followed by a 1-second mask. Participants then recalled the locations of all the dots. Participants could move around the dots until they were satisfied. Participants received feedback: the correct objects location and their recall (connected with lines).

Design

The serial reproduction design is the same as in Experiment 1. We generated 10 initial seed displays, each containing 15 colored dots with uniformly distributed locations. For each seed, we collected 10 serial reproduction chains with 25 iterations for each. Figure 9 shows three typical examples of serial reproduction chains.

Figure 9. Three typical chains starting with random seeds. The random seed, 5th iteration, 10th iteration, 15th iteration and 20th iteration are displayed. .

Participants

The experiment was conducted on the Amazon Mechanical Turk Marketplace (who performed our study for payment and a performance-based bonus). We allowed participants to perform multiple trials of our experiment for chains with different initial displays, resulting in a total of accepted 2732 experiment runs. Participants were not told that the stimuli they studied were another participant's responses.

Results

Do people use color-contingent grouping structures?

We first replicated the result from Experiment 1 that the overall dispersion of the patterns decreases across serial reproduction chains ($b = -0.06$, 95% CI [-0.08,-0.05]). Given the decreasing dispersion of overall patterns, the serial reproduction chains can converge either to random color-mixed structures or color-contingent grouping structures. To examine this account, we randomized the colors within each iteration. If color did not play a role in grouping, then the within-color dispersion will be unaffected by shuffling color labels across dots; however, if people grouped dots by color, and then systematically reported same-color dots as closer together, then the within-color dispersion will be greater after shuffling color labels (See Figure

10A). We found that all color groups became increasingly dispersed across serial reproduction chains (red group: $b = -0.12$, 95% CI [-0.14, -0.10]; blue group: $b = -0.13$, 95% CI [-0.16, -0.11]; green group: $b = -0.14$, 95% CI [-0.16, -0.11]) and color randomization significantly increased the dispersion of all the color groups (paired t-test, red group: $t(24) = 15.12$, $p < 0.001$; blue: $t(24) = 11.08$, $p < 0.001$; green group: $t(24) = -10.72$, $p < 0.001$), indicating that subjects' recalled patterns are color-contingent (See Figure 10B). Therefore, we conclude that serial reproduction chains gradually converge toward globally compact structures with color-specific groups.

Figure 10. Color grouping analysis. (A) Random color-mixed structure (top row) vs. color-contingent grouping structure (bottom row). The distinction between the two structures is whether the dispersion of the observed dots (solid line) changes significantly after color randomization (dotted line). For the color-mixed pattern, color randomization does not influence the dispersion of the color group. However, color-contingent groups are disrupted by color randomization and their dispersion increases significantly as a result. (B) Dispersion of original color groups (solid lines) vs. dispersion of color groups after color randomization (dotted lines). Error bars indicate ± 1 SEM. Iteration 0 represents the random seed. Pattern dispersion for color groups after randomization is significantly above that of the observed color groups indicating serial reproduction chains gradually converge towards color-contingent grouping structures.

What characteristics of structural biases emerge within color groups?

Since color features induce color-contingent grouping biases in visual memory, we can use the color-specific grouping structures to investigate the characteristics of within-group and between-group structural biases. In particular, we examine tendencies toward collinearity, orientation alignment across groups, and regularity of spacing.

Collinear grouping tendency

With color-specific grouping structures, we analyzed the organized shapes within color groups.

In general, the organized shapes can be categorized into isotropic structures, anisotropic structures and collinear structures (See Figure 11A). These organized shapes can be differentiated by the ratio of their principal components. The first principal component (PC1) represents the maximum variance direction in the data and the second principal component (PC2) is orthogonal to the first principal component. Mathematically, the two eigenvalues of the covariance matrix give the amount of variance carried along the two principal component directions. The ratio of these two eigenvalues captures the overall shape of the group (See Figure 11A). A greater ratio (i.e., large eigenvalue for PC1 and small eigenvalue for PC2), indicates that the overall configuration is more univariate, anisotropic, and line-like.

We classified color groups with PC ratios between 1 and 2 as isotropic (15% of the data), groups with PC ratios greater than 14 (25% of the data) as collinear and the rest as anisotropic (60% of the data). The partition assures adequate sample sizes and reasonable shape classification (i.e., groups whose ratios between 1 and 2 are approximately isotropic and groups whose ratios greater than 14 highly resemble linear structure). We found that the proportion of isotropic groups did not vary substantially over the 20 serial reproduction iterations ($b = -0.0002$, CI [-0.002, 0.001]). However, the proportion of linear groups ($b = 0.01$, CI [0.009, 0.011]) increased significantly, trading off with the proportion of anisotropic groups ($b = -0.009$, CI [-0.010, -0.007]) (See Figure 11B). This pattern suggests that a tendency for misreporting objects closer to their group centers is greater along axes on which the group has less variance, thus causing slightly anisotropic clusters to be reported as even more anisotropic, gradually converging to line-like structures. Altogether, this yields a systematic bias toward collinear structures in visuospatial memory.

Figure 11. Analysis of cluster shapes. (A) Categories of shapes: Isotropic group, anisotropic group, collinear group. The greater the ratio, the more anisotropic and collinear the configuration is. (B) The change in proportion of cluster shapes represented by area. Top area: collinear group (ratio > 14). Middle area: anisotropic group ($2 < \text{ratio} < 14$); Bottom area: isotropic group ($1 < \text{ratio} < 2$). The proportion of isotropic groups stays roughly the same across chains and the proportion of collinear groups increases significantly while the proportion of anisotropic groups significantly decreases. It implies visual memory of the locations tends to be biased towards collinear structures.

Similarity in orientation between color groups

With three clearly defined color-specific groups, we can evaluate the feature similarity between groups. Given that the majority of the color groups are anisotropic, one of the features that characterize the group similarity is its orientation. Because isotropic color groups do not have defined orientations, we first identified the iterations (66% of the data) in which all three color groups are anisotropic (i.e., PC ratios > 2). We used the direction of PC1 as an indicator of the shape orientation for the color groups and the orientation similarity was measured by the magnitude of the normalized vector sum of the first eigenvectors (unit vectors) (See Figure 12A), which yielded a number between 0 and 1. The larger the magnitude of the mean vector, the more similar their orientations are. Since the orientation of a group can be equivalently represented by the obtained first eigenvector direction or its opposite direction, we searched for a combination of three vectors that had minimum angle difference (i.e., maximum magnitude of vector mean) for each iteration. This process yields a higher average orientation similarity even on null data, so we obtained the null sampling distribution of this similarity measure for three randomly sampled unit vectors for each iteration. Because we only considered trials where all three color groups were anisotropic, the dispersion of the null distribution or orientation similarities varied across iterations due to different numbers of included trials.

We found that orientation similarity increased over the 25 serial reproduction iterations, indicating that successive reproductions made the color groups more aligned with one another ($b = 0.0014$, 95% CI [0.0009, 0.002]; See Figure 12B). In most of the later iterations, orientation similarity was significantly greater than expected from three randomly oriented vectors (e.g., $M(\text{iteration}=20) = 0.911$, while the 95% confidence interval on the simulated null was [0.858, 0.894]) indicating that the orientations were more congruent than expected by chance at $p < 0.05$. Furthermore, we would expect such orientation similarity among groups to be more pronounced when the groups themselves are more clearly oriented. To evaluate this, we considered all 1717 anisotropic trials in which all 3 groups had anisotropy ratios greater than 2, and asked whether trials in which the average anisotropy of the 3 groups was greater also had greater orientation similarity among groups. Specifically, we divided these anisotropic trials into five equal-sized bins ($n=343$ or 344) according to the mean anisotropy of the three color groups and examined the average orientation similarity within each bin. We found that trials in which the color groups were highly anisotropic (i.e. individual objects within color groups were arranged into something resembling a line), the three color groups tended to have much more similar orientations than expected by chance ($M(\text{level 4}) = 0.909$, $M(\text{level 5}) = 0.933$, compared to the 95% null hypothesis interval of [0.868, 0.886]). However, the orientation similarity on iterations when groups are only moderately anisotropic was no different from chance ($M(\text{level 1}) = 0.878$, $M(\text{level 2}) = 0.887$, $M(\text{level 3}) = 0.885$, 95% null hypothesis interval [0.868, 0.886]; See Figure 12C). Together, these results indicate that orientation similarity between groups emerges when the three groups are clearly collinear. This result makes some sense, because it is the most linear patterns that have the most identifiable orientations, and this result is also consistent with orientation similarity emerging in later iterations, as the most linear groups also arise only in later iterations.

Figure 12. Analysis of Orientation similarity. (A) Normalized vector sum. The direction of the first eigenvector represents the orientation of each color group. Because the vector direction and its opposite direction can both represent the same orientation, we choose the vector combination that maximizes the vector sum. The vector sum is normalized by the group number so that the max vector sum equals one indicating perfect alignment. A larger normalized vector sum implies more consistent orientations among color groups. (B) Mean normalized vector sum. Each dot represents the mean vector sum given the iteration where all 3 groups are oriented with an anisotropy ratio greater than 2. Error bars indicate SEM. The purple ribbon represents the estimated 95% confidence band of the means by simulations of 3 randomly oriented vectors conditioned on the sample sizes given the iterations (sample size varies across iterations). The green dots indicate the orientations of the color groups are significantly more similar than chance. The significant positive slope suggests that the color groups tend to be increasingly arranged in similar orientations across serial reproduction chains. (C) Orientation similarity in iterations where all 3 groups are oriented with an anisotropy ratio greater than 2, binned by the average anisotropy ratio of the three groups (x-axis). The higher the mean anisotropy on a given trial, the more line-like the groups in that trial (by definition) are, and the more the orientations are among the 3 groups (y-axis). Each dot represents the average normalized vector sum of all trials in a given bin, and error bars indicate ± 1 SEM. The purple ribbon represents the estimated 95% null hypothesis interval, obtained by simulating 3 randomly oriented vectors with the same number of simulated trials as represented in the bins ($n=343$ trials). Red dashed line is the mean of the null sampling distribution. When the dots are reported in 3 highly anisotropic (oriented) groups, those groups tend to be reported as parallel.

Spacing regularity in collinear arrangements within color groups.

Finally, we evaluated whether people systematically report objects along a line to be regularly spaced. Such a regularity would manifest as equidistant spacing among objects, as opposed to random dispersion. Measuring the extent to which objects are equidistantly spaced is equivalent to measuring the variability of the pairwise distance between the objects. We projected the dots in the collinear color group onto its first eigenvector and normalized the projected locations on a scale of 0 to 1. The normalized projected locations are used to calculate the standard deviation of the pairwise distances (See Figure 13A). Smaller standard deviation implies that the arrangement of dots has more equidistant spacing.

We examined the spacing regularity on the most collinear groups (25% of the data) whose PC ratios are greater than 14. We found that variability of the pairwise distances decreased across the serial reproduction chains ($b = -0.0065$, CI 95% [-0.0077, -0.0054]) and the spacing of the dots significantly differed from the randomly-spaced baseline arrangement (See Figure 13B). It implies that visual memory for the locations of dots in these collinear structures are biased towards regular spacings.

Figure 13. Analysis of spacing regularity. (A) Measure of spacing variability. Dots from linear color groups are projected onto the first Eigenvector and the standard deviation of pairwise distances between the projected locations measure how equidistantly spaced the dots are. (B) The change of spacing variability across chains. Purple ribbon represents the 95% confidence interval of the mean SD of pairwise distances estimated by simulating uniformly distributed dots on a unit vector conditioned on sample size for a given iteration. The green dots indicate significant deviance from the random baseline. Overall, the spacing variability decreases and objects tend to be arranged in equidistant spacing in the collinear structures.

Discussion

Experiment 1 suggests that memory reproductions tend toward clustered spatial arrangements. However, when cluster membership is only indicated by the spatial arrangement, we cannot independently estimate which cluster a given object belongs to, and what the within-cluster spatial memory biases are. Consequently, the goal of Experiment 2 was to induce grouping independently, by introducing a salient color feature, and investigate the within-group and between-group spatial configuration biases. First, we confirmed that serial reproduction chains converged toward increasingly compact color-specific groups, indicating that object color was an overwhelming grouping cue (Wertheimer, 1923; Quinlan & Wilton, 1998), and that color-grouped objects were recalled with increasing spatial compactness. The color-specific grouping structure is consistent with Boolean map theory which explains visual attention is constrained to

access only a single feature (e.g., green) associated with multiple spatial locations at one moment (Huang & Pashler, 2007). Color-specific configuration allows “chunking” of locations that form compressed representations (Cowan, 2001; Brady, Konkle & Alvarez, 2009). Because working memory is severely capacity-limited, the configuration-chunking strategy may greatly enhance processing efficiency (Jiang, Olson & Chun, 2000) and bias the spatial locations of the objects. Within color groups, visual memory reconstructed groups into more regular collinear structures. Previous research suggests that people are very sensitive to collinearity as cue to the presence of ecologically realistic contours (Schwartzkopf and Kourtzi, 2008). The anisotropic groups that manifest weak collinearity are biased toward strong collinear structures in favor of more ecological realistic arrangements. Thus, such canonical representations of collinear ensemble structures may reflect the sensitivity to collinearity. Between collinear color groups, color groups in the same iteration had increasingly similar orientations across the serial reproduction chains. Within collinear color groups, locations of dots that formed the collinear groups were biased towards collinear arrangements with equidistant spacings. These biases suggest that the visuospatial working memory system capitalizes on the regularities that create redundancies in order to encode more items (Cover & Thomas, 1991; Brady, Konkle & Alvarez, 2009).

General Discussion

Using a serial reproduction task, we revealed structures people use to represent the spatial arrangement of objects. In Experiment 1, we examined memory for locations by asking participants to recall the locations of objects in a serial reproduction task. We showed that people misremembered specific items toward a globally compact structure, and organized them into clustered spatial groups. In Experiment 2, we first confirmed that discrete color features

explicitly introduce color-contingent spatial configuration. By analyzing color groups, we circumvented the grouping uncertainty in Experiment 1 and further revealed that people compress color groups into dot line segments with similar orientations and equidistant spacing.

Serial reproduction paradigm

Bartlett's "serial reproduction" experiments (Bartlett, 1932) were the first psychological investigation to examine how memory biases influence information transmission. The method has been extended to study language evolution (Kirby, 2001; Kirby et al., 2008; Griffiths & Kalish, 2005), cultural transmission (Atran, 2001, 2002; Mesoudi, 2007), inductive biases in function learning (Kalish, Griffiths, & Lewandowsky, 2007), category structures (Sanborn & Griffiths, 2007) and spatial memory biases (Langlois et al., 2021). However, there is an important distinction between a serial reproduction paradigm, as employed here, and the MCMC with people paradigm (Sanborn & Griffiths, 2007). MCMC with people sets up the proposal and acceptance distributions in such a way as to guarantee that the chains would converge to a stationary distribution that matches the prior. This is not the paradigm we employed, since our proposal and acceptance distributions are not controlled, and are entirely determined by subjects. Instead, we employ a serial reproduction paradigm, which does not converge to a stationary distribution that matches the prior, but by emphasizing and exaggerating inductive biases through the process of repeated noisy reproduction, wherein each noisy step offers another opportunity for biases to intervene (e.g., Huttenlocher et al. 2001; Hemmer & Steyvers, 2009), reveals structures people use to help encode the displays. Consequently, stationarity and convergence are not prerequisites for learning about the structures that lead to better encoding of the display from a serial reproduction paradigm. More formally, the distribution toward which the serial reproduction chains are converging are not unimodal in the space of positions, but are

lower entropy in the space of all possible positions: Consider starting with a uniform distribution over all pixel images and converging to images of English letters – while the English letters are all qualitatively different, the distribution of images that contain English letters is much lower entropy than the distribution of all pixel images. So convergence toward easier to encode displays (for new participants), as we found here, is entirely possible while maintaining many modes and qualitative differences across chains. Despite the diversity and idiosyncrasy of the patterns in the serial reproduction chains, our analyses reveal some salient and consistent structural characteristics that people use to represent displays in visuospatial memory. Thus, our study further confirms the fruitful application of the serial reproduction paradigm in revealing structure in human cognition.

Implications for structured priors in visuospatial working memory

In contrast to the traditional assumption that objects in visual memory are encoded independently (Anderson, Vogel, & Awh, 2011; Bays & Husain, 2008; Zhang & Luck, 2008), our results indicate that people appear to encode higher-order hierarchical ensemble structures (Orhan & Jacobs, 2014) that biases spatial locations of objects. Working memory capacity is highly limited, the more items there to encode and store, the fewer bits there are available to represent each one (Neisser, 1967). People can utilize such structured representations that encode a set of objects to cope with limited capacity rather than simply reducing the resolution of local representation (Dan, 2001; Brady & Tenenbaum, 2013). From our results, it appears that such structures bias the locations of objects into globally compact clustering structures and individual color dots are attracted and compressed into color-specific clusters. By encoding the whole set and “chunks”, the configuration-chunking strategy greatly enhances processing efficiency by forming compressed representations (Cowan, 2001; Brady, Konkle & Alvarez, 2009; Jiang,

Olson & Chun, 2000). Moreover, biasing the recall of locations toward the center of sets can compensate for uncertainty and aid recall (Lew & Vul, 2015).

Furthermore, the emergence of color-specific grouping structures may reflect the inferences or predictions of partial data due to the limited working memory capacity and this ability depends on the existence of statistical dependencies in natural images (Kersten, 1987; Hansmann-Roth et al., 2021). For example, in a natural scene, colors of proximate locations are likely to be similar. The statistical priors about spatial configurations of similar colors bias the inferences of the locations when partial information is available.

Implications for regularities in visuospatial memory

According to information theory, more information can be stored if there are redundancies in the input in an optimal system (Cover & Thomas, 1991). In other words, creating redundancies makes it possible to encode more items within a limited capacity. The collinear structure bias, orientation similarity and spacing regularity all reflect the tendency of generating redundancy to cope with limited working memory. For examples, biasing the anisotropic structure toward a more collinear arrangement shrinks the variance on the PC2 and generates redundancy in the set of possible projected distances from the object locations to the PC2; spacing regularity in collinear structures produces spacing redundancy along the direction of PC1; orientation similarity increases redundancy in the set of possible orientations for groups. Removal of redundancy forms compressed and simplistic representations, allowing more items to be stored more accurately in memory (Brady, Konkle & Alvarez, 2009). These results are consistent with the view (Attneave, 1954; Barlow, 1959) that a principal mission of biological vision may be to encode the visual image into a less redundant form. For example, eigenvector transformation of

face images facilitates a large reduction of dimensionality which may be useful for economical representation and efficient retrieval (Sirovich & Kirby, 1987).

Visual system has evolved to utilize statistical regularities in the environment for extracting shape information from the noisy sensory input. Behavioral and computational work suggests that observers are better at detecting collinear edges (i.e., edges aligned along a path) (Dakin & Hess 1997; Field et al. 1993; Hess & Field 1999) that co-occur frequently and form contours in natural images (Geisler 2008; Geisler et al. 2001; Sigman et al. 2001). We found that the anisotropic groups are increasingly biased towards linear grouping structures. The sensitivity to collinearity may drive the anisotropic groups that manifest weak collinearity to be biased toward strong collinear structures in favor of more ecologically realistic arrangements.

Limitations and future directions

Our results, by no means, exhaustively capture the structures used in visuospatial memory. Specifically, we only focus on the most basic and general structures such as clustering and collinearity. However, observers frequently encode objects in complex shapes in the real-world. Likewise, a quick glance at responses in later iterations of our studies reveals structures like letters and shapes that suggest the use of long-term knowledge (Figure 14 displays several particularly notable structures). Although we were able to capture much of how people grouped and organized objects in visual memory, there are potentially many more complex and richer structures that are embedded in our data. Moreover, the patterns of convergence in our study demonstrate that observers' have at least some shared structures they use to store items Future work may further examine how observers represent the display using basic representational units like individual elements and ensembles to more sophisticated structures like parts of objects, whole objects and scenes (Palmer, 1977; Biederman, 1987; Orhan & Jacobs, 2014). We have

made all of our data and analysis code available for others to explore to identify further systematic structures that emerge through serial spatial memory reproduction.

Figure 14. Examples of sophisticated structures that we were unable to account for through our analyses.

References

Alvarez, G. A. (2011). Representing multiple objects as an ensemble enhances visual cognition. *Trends in Cognitive Sciences*, 15, 122–131.

Alvarez, G.A., & Oliva, A. (2009). Spatial ensemble statistics are efficient codes that can be represented with reduced attention. *Proceedings of the National Academy of Sciences*, 106(18), 7345—7350.

Attneave, F., "Informational aspects of visual perception," *Psych. Rev.* 61, 183-193 (1954).

Austerweil, J. (2014). Testing the psychological validity of cluster construction biases. In *Proceedings of the 36th annual meeting of the cognitive science society*. (pp. 122—127).

Barlow, H. B., "Sensory mechanisms, the reduction of redundancy, and intelligence," in *Proceedings of the National Physical Laboratory Symposium on the Mechanization of Thought Processes* (H. M. Stationery Office, London, 1959), No. 10, pp. 535-539.

Bartlett, F.C. (1932). *Remembering: A Study in Experimental and Social Psychology*. Cambridge University Press.

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

Brady T. F., Konkle T., Alvarez G. A. (2009). Compression in visual working memory: Using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General*, 138 (4), 487–502.

Brady, T.F., & Tenenbaum, J.B. (2013). A probabilistic model of visual working memory: Incorporating higher order regularities into working memory capacity estimates. *Psychological review*, 120(1), 85—109.

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.

Cowan N. (2005). Working memory capacity. Hove, East Sussex, UK: Psychology Press.

Dakin 1997. Dakin SC, Hess RF. The spatial mechanisms mediating symmetry perception. *Vision Res* 37: 2915–2930, 1997

Ariely, D. (2001) Seeing sets: representation by statistical properties. *Psychol. Sci.* 12, 157–162
Lew TF, Vul E. Ensemble clustering in visual working memory biases location memories and reduces the Weber noise of relative positions. *J Vis.* 2015;15(4):10.

Thomas A. Langlois, Nori Jacoby, Jordan W. Suchow, Thomas L. Griffiths. Serial reproduction reveals the geometry of visuospatial representations. *Proceedings of the National Academy of Sciences* Mar 2021, 118 (13) e2012938118

Ferguson, T.S. (1983). Bayesian density estimation by mixtures of normal distributions. *Recent advances in statistics*, 24, 287—302.

Field 1993. Field DJ, Hayes A, Hess RF. Contour integration by the human visual system: evidence for a local “association field.” *Vision Res* 33: 173–193, 1993.

Froyen, V., Feldman, J., & Singh, M. (2015). Bayesian hierarchical grouping: Perceptual grouping as mixture estimation. *Psychological Review*, 122(4), 575–597.

Geisler 2001. Geisler WS, Perry JS, Super BJ, Gallogly DP. Edge co-occurrence in natural images predicts contour grouping performance. *Vision Res* 41: 711–724, 2001.

Geisler 2008. Geisler WS. Visual perception and the statistical properties of natural scenes. *Annu Rev Psychol* 59: 167–192, 2008.

Geman, S., & Geman, D. (1984). Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, (6), 721—741.

Griffiths TL, Kalish ML (2005) A bayesian view of language evolution by iterated learning in 1033 *Proceedings of the 27th Annual Meeting of the Cognitive Science Society*. Vol. 27.

Hartigan, P.M., Algorithm as 217: computation of the dip statistic to test for unimodality, *J. R. Stat. Soc. Ser. C (Appl. Stat.)* 34 (3) (1985) 320–325 .

Hess 1999. Hess R, Field D. Integration of contours: new insights. *Trends Cogn Sci* 3: 480–486, 1999.

Hemmer, P., & Steyvers, M. (2009). A Bayesian Account of Reconstructive Memory. *Topics in cognitive science*, 1 1, 189-202 .

Huang, L., & Pashler, H. (2007). A Boolean map theory of visual attention. *Psychological Review*, 114, 599–631.

Huttenlocher, J., Hedges, L. V., & Vevea, J. L. (2000). Why do categories affect stimulus judgment? *Journal of Experimental Psychology: General*, 129(2), 220–241.
<https://doi.org/10.1037/0096-3445.129.2.220>

Im, H.Y., & Chong, S.C. (2014). Mean size as a unit of visual working memory. *Perception*, 43(7), 663-676.

Jiang Y, Olson IR, Chun MM. Organization of visual short-term memory. *Journal of Experimental Psychology: Learning Memory and Cognition*. 2000;26:683–702.

Kempe, V., Gauvrit, N., & Forsyth, D. (2015). Structure emerges faster during cultural transmission in children than in adults. *Cognition*, 136, 247—254.

Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: an experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences of the United States of America*, 105(31), 10681–10686.

Kirby, S. (2001). Spontaneous evolution of linguistic structure: An iterated learning model of the emergence of regularity and irregularity. *Journal of Evolutionary Computation*, 5, 102–110.

Miller G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 81–97.

Mesoudi, A. (2007). Using the methods of experimental social psychology to study cultural evolution. *Journal of Social, Evolutionary, and Cultural Psychology*, 1, 35–38.

Orhan, A.E., & Jacobs, R.A. (2013). A probabilistic clustering theory of the organization of visual short-term memory. *Psychological review*, 120(2), 297—328.

Orhan, A.E., & Jacobs, R.A. (2014a). Toward ecologically realistic theories in visual short-term memory research. *Attention, Perception, & Psychophysics*, 76(7), 2158—2170.

Orhan, A.E., & Jacobs, R.A. (2014b). Are performance limitations in visual short-term memory tasks due to capacity limitations or model mismatch. *Manuscript under review*.

Orhan, A.E., Sims, C.R., Jacobs, R.A., & Knill, D.C. (2014). The adaptive nature of visual working memory. *Current Directions in Psychological Science*, 23(3), 164—170.

Peterson DJ, Berryhill ME. The Gestalt principle of similarity benefits visual working memory. *Psychonomic Bulletin & Review*. 2013 Dec;20(6):1282-1289.

Quinlan PT, Wilton RN. Grouping by Proximity or Similarity? Competition between the Gestalt Principles in Vision. *Perception*. 1998;27(4):417-430.

Sanborn, A., & Griffiths, T.L. (2007). Markov chain Monte Carlo with people. *Advances in neural information processing systems*, (pp. 1265—1272).

Sigman 2001. Sigman M, Cecchi GA, Gilbert CD, Magnasco MO. On a common circle: natural scenes and Gestalt rules. *Proc Natl Acad Sci USA* 98: 1935–1940, 2001.

Sims, C.R., Jacobs, R.A., & Knill, D.C. (2012). An ideal observer analysis of visual working memory, *Psychological review*, 119(4), 807—830.

M. V. Srinivasan, S. B. Laughlin, and A. Dubs, "Predictive coding: a fresh view of lateral inhibition," *Proc. R. Soc. London Ser. B* 216, 427-459 (1982)

Switkes, E., Mayer, M.J., & Sloan, J.A. (1978). Spatial frequency analysis of the visual environment: anisotropy and the carpentered environment hypothesis. *Vision Research*, (13), 1393—1399.

Wertheimer M. Untersuchungen zur Lehre von der Gestalt, II. *Psychologische Forschung*. 1923; 4:301–350. (Translated extract reprinted as “Laws of organization in perceptual forms.” In W. D. Ellis (Ed.), (1938). *A source book of Gestalt psychology* (pp. 71–94). London, U. K.: Routledge & Kegan Paul Ltd.).

Woodman GF, Vecera SP, Luck SJ. Perceptual organization influences visual working memory. *Psychonomic Bulletin & Review*. 2003;10:80–87.

Wagemans, Johan, Feldman, Jacob, Gepshtain, Sergei, et al.. "A century of Gestalt psychology in visual perception: II. Conceptual and theoretical foundations." *Psychological Bulletin*, 138, no. 6 (2012) American Psychological Association: 1218-1252.

Xu J, Griffiths TL (2010) A rational analysis of the effects of memory biases on serial reproduction. *Cognitive Psychology* 60(2):107–126.

Xu Y. Encoding color and shape from different parts of an object in visual short-term memory. *Perception & Psychophysics*. 2002a;64:1260–1280.

Xu Y, Chun MM. Visual grouping in the human parietal cortex. *Proceedings of the National Academy of Sciences, USA*. 2007;104:18766–18771.

