Routledge
Taylor & Francis Group

Visual Cognition

ISSN: 1350-6285 (Print) 1464-0716 (Online) Journal homepage: http://www.tandfonline.com/loi/pvis20

Object representations are biased toward each other through statistical learning

Ru Qi Yu & Jiaying Zhao

VISUAL COGNITION, 2018
VOL. 26, NO. 4, 253–267
https://doi.org/10.1080/13506285.2018.1435596

Object representations are biased toward each other through statistical learning

Ru Qi Yu[a] and Jiaying Zhao[a,b]

[a]Department of Psychology, University of British Columbia, Vancouver, Canada; [b]Institute for Resources, Environment and Sustainability, University of British Columbia, Vancouver, Canada

ABSTRACT

The visual system is remarkably efficient at extracting regularities from the environment through statistical learning. While such extraction has extensive consequences on cognition, it is unclear how statistical learning shapes the representations of the individual objects that comprise the regularities. Here we examine how statistical learning alters object representations. In three experiments, participants were exposed to either random arrays containing objects in a random order, or structured arrays containing object pairs where two objects appeared next to each other in fixed spatial or temporal configurations. After exposure, one object in each pair was briefly presented and participants judged the location or the orientation of the object without seeing the other object in the pair. We found that when an object reliably appeared next to another object in space, it was judged as being closer to the other object in space even though the other object was never presented (Experiments 1 and 2). Likewise, when an object reliably preceded another object in time, its orientation was biased toward the orientation of the other object even though the other object was never presented (Experiment 3). These results demonstrated that statistical learning fundamentally shapes how individual objects are represented in visual memory, by biasing the representation of one object toward its co-occurring partner. Importantly, participants in all experiments were not explicitly aware of the regularities. Thus, the bias in object representations was implicit. The current study reveals a novel impact of statistical learning on object representation: spatially co-occurring objects are represented as being closer in space, and temporally co-occurring objects are represented as having more similar features.

ARTICLE HISTORY
Received 20 July 2017
Accepted 27 January 2018

KEYWORDS
Statistical learning; implicit bias; visual memory; object representations; visual features

A remarkable ability of the visual system is the rapid extraction of the stable aspects in the environment. Statistical learning involves the automatic detection of regularities in terms of how objects co-occur in space or over time (Fiser & Aslin, 2001; Perruchet & Pacton, 2006; Saffran, Aslin, & Newport, 1996; Turk-Browne, Scholl, Chun, & Johnson, 2009; Zhao, Al-Aidroos, & Turk-Browne, 2013). Statistical learning is an automatic process (Turk-Browne, Jungé, & Scholl, 2005), operating in multiple sensory modalities and over many types of stimuli, including tones (Saffran, Johnson, Aslin, & Newport, 1999), haptic input (Conway & Christiansen, 2005), shapes (Fiser & Aslin, 2001), colours (Turk-Browne, Isola, Scholl, & Treat, 2008), and lines (Zhao, Ngo, McKendrick, & Turk-Browne, 2011).

Exposure to regularities has extensive consequences on cognition. For example, the presence of regularities draws attention implicitly (Yu & Zhao, 2015; Zhao et al., 2013; Zhao, Cosman, Vatterott, Gupta, & Vecera, 2014), alters the spatial scale of attention (Zhao & Luo, 2017), interferes with ensemble processing (Zhao et al., 2011), facilitates object label learning (Graf Estes, Evans, Alibali, & & Saffran, 2007) and object categorization (Turk-Browne, Scholl, Johnson, & Chun, 2010), and increases visual short-term memory capacity (Brady, Konkle, & Alvarez, 2009; Umemoto, Scolari, Vogel, & Awh, 2010). Moreover, statistical learning occurs implicitly, in that participants remain unaware of the regularities (Conway & Christiansen, 2005; Fiser & Aslin, 2001; Perruchet & Pacton, 2006; Saffran et al., 1996; Turk-Browne et al., 2009; Turk-Browne, Jungé, & Scholl, 2005; Zhao et al., 2013).

While past work on statistical learning focused on the impact of regularities on various cognitive processes, it is unclear how statistical learning shapes the representations of individual objects that co-occur in space or over time. Previous studies have proposed several factors that influence the fidelity of

CONTACT Ru Qi Yu ruqiyu@psych.ubc.ca Department of Psychology, University of British Columbia, Vancouver, BC, Canada V6T 1Z4

memory representations of objects that are spatially or temporally associated. These factors include spatial or temporal grouping by distance or by categories (Huttenlocher, Hedges, & Duncan, 1991; McNamara, 1986). For example, when two objects were presented closer in time, the recognition of these objects was facilitated and their spatial locations were also recalled as being closer (McNamara, Halpin, & Hardy, 1992). When two objects were presented at the same time, their locations were recalled as being closer in space than when the objects were presented in isolation (Recker & Plumert, 2009). Likewise, when two objects were categorically related, their locations were recalled as being closer in space (Recker & Plumert, 2009). Moreover, a sequence of spatially and temporally related objects tends to be judged as the same object (Wallis & Bülthoff, 2001). A recent study demonstrates that an object is perceived as having a higher value if it is temporally associated with another object that predicted monetary rewards than if it is associated with an object that did not predict monetary rewards, and this change in value through association is predicted by hippocampal activities (Wimmer & Shohamy, 2012). This study shows that the feature of an object (e.g., monetary value) becomes more similar to that of another object simply through temporal association. Taken together, these findings suggest that spatial or temporal association between objects systematically alters the memory representation of these objects.

It is proposed that the change in representation occurs because of a grouping or clustering mechanism (McNamara, 1986; Orhan & Jacobs, 2013; Schapiro, Kustner, & Turk-Browne, 2012). For example, the spatial or temporal proximity grouped the objects as one representational unit such that the recall of an object in the group is biased toward other objects in the group (McNamara et al., 1992; Recker & Plumert, 2009).

In the context of statistical learning, individual objects are associated with each other not by spatial or temporal proximity, but by joint or transitional probabilities. That is, the spatial distance or the temporal interval between two objects is constant, but the joint or transitional probability is higher between two objects in a pair than two objects not in a pair. For example, in an AB pair, B always follows A in a continuous temporal sequence (co-occurring over time) or B always appears next to A in a spatial array (co-occurring in space). After B, any other pair can follow, and therefore the transitional probability between B and another object is lower. It is the reliable probability between the two objects that serves as a grouping cue for a pair.

Given the previous findings, we hypothesize that the statistical association can also shape the representation of individual objects through a grouping mechanism. This hypothesis is motivated by recent studies that showed a chunking effect of statistical learning, that is, two co-occurring objects in space or time may be grouped as one unit. For example, co-occurring colour dots in a spatial scene are compressed in working memory (Brady et al., 2009), and their quantity is under-estimated consistent with a feature-grouping mechanism (Zhao & Yu, 2016). Such chunking effect emerges rapidly and unintentionally, even after one exposure to the co-occurring objects (Batterink, 2017). One proposed consequence of chunking is that the two objects may be represented as being closer in space or more similar to each other. This is based on past findings mentioned above, and also studies that show children and adults judge the objects from the same spatial group as being closer together than they really are, a systematic bias toward the centre of the group (Hund & Plumert, 2002, 2003; Hund, Plumert, & Benney, 2002). The bias is shown in other features such as size, where the recall of an individual object in a group is biased toward the group average (Brady & Alvarez, 2011). An influential formal model of visual representation suggests that individual items in an array are represented as a probability distribution over possible clustering or partitions of all items (probabilistic clustering theory; Orhan & Jacobs, 2013). This account proposes that representations of objects in the same cluster share parameters and are therefore dependent. This account supports our hypothesis that two co-occurring objects may be represented as being closer in space or as having more similar features.

The goal of the current study is to examine how statistical learning of object co-occurrences (i.e., regularities) shapes the memory representations of individual objects. There are two possible ways where statistical learning alters the representation of an individual object. One possibility is that the representation of one object is biased toward the representation of its co-occurring partner, which is consistent with our hypothesis. By biasing toward we mean that the location of one object is recalled as being closer to its co-occurring partner, or the feature of one object is recalled as being similar to its co-occurring partner. In a recent study, the cortical activation patterns in the medial temporal lobe of two temporally co-occurring objects became more correlated after statistical learning, compared to the patterns of two objects that were not temporally associated (Schapiro et al., 2012). This means that statistical learning increased the similarity in the neural representations of two co-occurring objects. The current study explores whether statistical learning also increases the similarity in object representations in working memory, in addition to the increased neural similarity. The other possibility is that co-occurring objects may be represented as being more distinct from each other, since successful learning of regularities depends on the successful discrimination between individual objects (e.g., Stager & Werker, 1997). Both human and animal imaging studies have suggested that stronger discrimination between two stimuli renders the neural representations of these stimuli more dissimilar (Faber, Joerges, & Menzel, 1999; Li, Howard, Parrish, & Gottfried, 2008).

To seek behavioural evidence of how statistical learning shapes object representations, we conducted three experiments. In each experiment, participants were first exposed to either structured arrays (containing object pairs) or random arrays, while performing a cover task. After exposure, one object within each pair was briefly presented and participants judged the location of the object (Experiments 1 and 2), or the orientation of a line (Experiment 3). We examined the fidelity of participants' judgments of the objects as a way to measure the extent and the direction of change in the object representations as a result of learning. Specifically, we predict that if statistical learning biases their representations toward each other, one object should be judged as being closer to the other object in space even in the absence of the other object (Experiments 1 and 2). Alternatively, if learning the object pair leads to more distinct representations, one object should be judged as being further away from the other object. Likewise, if statistical learning biases the line representations toward each other, one line should be judged as being more oriented toward the other line (Experiment 3). Alternatively, if learning leads to more distinct representations, one line should be judged as being oriented away from the other line. Alternatively, if learning leads to more distinct representations, one line should be judged as being oriented away from the other line.

Experiment 1

The goal of the experiment was to examine how regularities in terms of object co-occurrences in space alter the representations of individual objects.

Participants

Eighty undergraduate students (58 female, mean age = 19.9 years, *SD* = 1.9) from University of British Columbia participated for course credit. Participants in all experiments had normal or corrected-to-normal vision and provided informed consent. All experiments have been approved by the University of British Columbia Behavioral Research Ethics Board. No participants were excluded from the experiments (except for Experiment 3). The sample sizes in the current experiments were estimated from the sample size and effect size from our previous experiments (Zhao et al., 2013). Data analysis was performed only when all data were collected.

Stimuli

Stimuli consisted of eight black shapes, each subtending 2.3° of visual angle. The shapes were randomly assigned for every participant to four "pairs" grouped in fixed horizontal, vertical, and diagonal configurations (Figure 1a). For every trial, the four pairs were presented in a 4×4 invisible grid (subtending $12.7° \times 12.7°$), with the constraint that each pair neighboured at least one of the other pairs. This constraint ensured that statistical learning could not solely be determined by spatial segmentation cues other than co-occurrence.

Apparatus

In all experiments, participants were seated 50 cm from a computer monitor (refresh rate = 60 Hz). Stimuli were presented using MATLAB (Mathworks) and the Psychophysics Toolbox (http://psychtoolbox.org).

Procedure

Participants were randomly assigned to either the *structured* condition or the *random* condition ($N = 40$

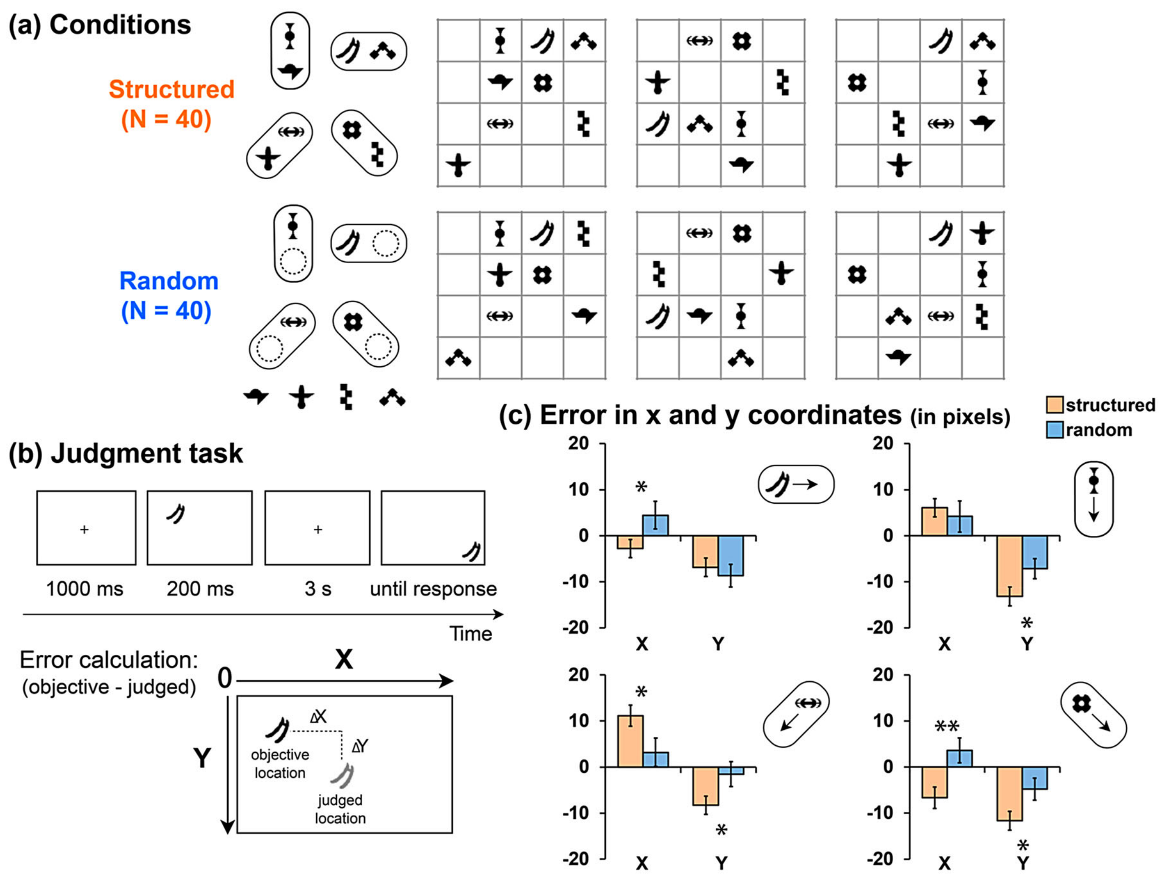

Figure 1. Stimuli, tasks, and results from Experiment 1. (a) In the structured condition, eight shapes were grouped into four pairs, which were assigned to horizontal, vertical, and diagonal configurations. In the random condition, one shape in each pair remained in the same position but the other shape appeared randomly in one of the other three pairs. In each trial, the four pairs were presented in an invisible 4×4 grid. Three sample trials are shown in each condition. (b) In the judgment task, the shape that always remained in the same position within each pair was briefly presented on the screen, followed by a 3 s interval. Participants reported the position of the shape using the cursor. The error was calculated by subtracting the judged position from the objective position of the shape, measured in pixels along the X and Y screen coordinates. (c) Each graph shows the judgment errors in each pair. (Error bars represent ± 1 SEM; $*p < .05$; $**p < .01$).

in each). There were three phases in each condition: exposure, judgment, and test. During the exposure phase in the structured condition, the four pairs were presented in the grid for every trial. In the random condition, one shape always remained in the same position within each pair but the other shape appeared randomly in any of the four pairs. Specifically, the left shape in the horizontal pair, the top shape in the vertical pair, the top right shape in one diagonal pair, and the top left shape in the other diagonal pair remained in the same position within its respective pair, while the remaining four shapes appeared randomly in the remaining four positions (Figure 1a). This ensured that any bias observed of the shape that remained in the same position within a pair would be driven only by the reliable co-occurrence of the two shapes in the pair, and not by changes in spatial locations. In both conditions, each shape was repeated 200 times, resulting in a total of 200 trials. In each trial, the array was presented for 1500 ms followed by an inter-stimulus interval (ISI) of 1000 ms. Of all trials, 20% contained a duplicate shape: one shape in a randomly selected pair was changed to match the other shape from the pair. The duplicates, which existed in both conditions, were used in support of a cover task (duplicate detection task) during the exposure phase. Participants were asked to detect whether there were two identical shapes in each trial by pressing the "z" key for no or the "/" key for yes (key assignment counterbalanced across participants). The purpose of the cover task was to hide the purpose of the experiment and to ensure incidental encoding of regularities.

The judgment phase immediately followed the exposure phase in both conditions. For each trial, the shape that always remained in the same relative position within its pair appeared in the grid, as in the exposure phase. This includes the left shape in the horizontal pair, the top shape in the vertical pair, and the top right shape and the top left shape in the diagonal pairs. Thus, only four distinct shapes were presented. We did not test the remaining four shapes in the four pairs, because in the random condition for each trial they could randomly appear with any of the four shapes that were fixed in the same relative position in the pair, whereas in the structured condition the remaining four shapes always appeared in the same relative position in the pair. Testing only one member in the pair ensured that any bias we found in the judgment task was not caused by the relative position of shapes in the structured condition, but rather by the reliable pairing between two shapes. In other words, we wanted to make sure that the bias arises not because the shape always appeared in the fixed position within a pair, but rather because the shape always appeared with another shape (i.e., reliable co-occurrences). Therefore, we did not test the remaining shape in each pair.

The matrix contained a total of 16 locations in a 4×4 matrix. Given the spatial configurations of the pairs, the shape to the presented in a pair in the judgment phase could not appear in all 16 locations (e.g., the left member of a horizontal pair could never appear in the right-most column in the matrix). Consequently, for the horizontal pair and the vertical pair, the shape could appear in 12 valid locations in the grid. For the two diagonal pairs, the shape could appear in nine valid locations in the grid. One shape appeared in a valid position in the grid at a time, and this was repeated twice, resulting in 84 trials in total ($[12 + 12 + 9 + 9] \times 2 = 84$). The trials were presented in a random order. In each trial, the shape appeared on the screen for 200 ms, followed by a blank interval of 3000 ms. Then the same shape appeared again on the bottom right corner of the screen, and participants used the cursor to locate the shape in the position in which it had just appeared (Figure 1b). No feedback was given on judgment accuracy.

After the judgment phase, participants in the structured condition completed the two-alternative forced-choice (2AFC) test phase. In each trial, two sets of shapes were presented for 1000 ms, one on the left side of the screen and the other on the right side. One set was the target pair from the exposure phase, and the other was a foil pair consisting of one shape from the target pair and one shape from a different pair. Participants were asked to choose which pair looked more familiar. The two shapes in the foil pair had never appeared in this spatial configuration, but the two shapes in the target pair had been repeatedly presented in the exposure phase. If participants had learned the co-occurrence between the two shapes in the target pair, they should be able to choose the target pair as more familiar. Each target pair was tested against one of the two possible foil pairs, and each target–foil combination was tested twice, creating 16 trials (order randomized). The left/right side of the target pair was counterbalanced across trials. Because all individual shapes were equally frequent in the exposure and test phases, participants could only choose the pair as more familiar if they had learned which shapes co-occurred.

After the test, a debriefing session was conducted, where participants were asked if they noticed any shapes that always appeared with one another in the first part of the experiment, and if they noticed any patterns or regularities in the shapes. For those who responded yes, we further asked them to identify which shapes. Participants were asked other questions, including how confident they were in the test phase and what they thought the purpose of the study was.

Results

During exposure, mean accuracy in the duplicate detection task was 93.3% ($SD = 11.0\%$) in the structured condition and 92.3% ($SD = 14.9\%$) in the random condition. Performance did not differ between the two conditions [$t(78) = 0.36$, $p = .72$, $d = .08$]. At the test phase, in the structured condition pairs were chosen over foils for 55.5% ($SD = 10.8\%$) of the time, which was reliably above chance [50%; $t(39) = 3.20$, $p = .003$, $d = .51$]. During debriefing, all participants were asked to identify any regularities in terms of how the shapes appeared in the exposure phase, and 45% of the participants incorrectly reported the shape duplicates as regularities, but no participant was able to correctly identify the specific shapes that reliably co-occurred within a pair. Moreover, most participants did not explicitly know the answer and felt like guessing in the 2AFC test phase.

This suggests that participants did not have explicit awareness of the shape pairs, but rather implicitly learned the shape pairs.

The critical question was whether the shape pairs biased the representation of individual shapes. To address this question, we calculated the error in pixels between the objective location and the judged location of each shape presented in the judgment phase. The error was derived from subtracting the judged coordinates from the objective coordinates in pixels along the X- and the Y-axes on the screen (Figure 1b). We compared the signed error of the shape in each pair between the structured and the random conditions (Figure 1c). Using this analysis, a smaller or negative error suggests a rightward bias on the X-axis, or a downward bias on the Y-axis. In contrast, a larger error suggests a leftward bias on the X-axis, or an upward bias on the Y-axis.

For the left shape in the horizontal pair, the error on the X-axis was smaller in the structured condition than in the random condition [$t(78) = 2.01$, $p = .048$, $d = .45$]. There was no significant difference in the errors on the Y-axis between the two conditions [$t(78) = 0.57$, $p = .57$, $d = .13$]. This suggests that the left shape in the horizontal pair was biased toward the right.

For the top shape in the vertical pair, the error on the Y-axis was smaller in the structured condition than in the random condition [$t(78) = 2.10$, $p = .04$, $d = .47$]. There was no significant difference in the errors on the X-axis between the two conditions [$t(78) = 0.42$, $p = .68$, $d = .09$]. This suggests that the top shape in the vertical pair was biased toward the bottom.

For the top right shape in the first diagonal pair, the error on the X-axis was larger in the structured condition than in the random condition [$t(78) = 2.07$, $p = .04$, $d = .46$]. The error on the Y-axis was smaller in the structured condition than in the random condition [$t(78) = 2.01$, $p = .048$, $d = .45$]. This suggests that the top right shape in the diagonal pair was biased toward the left and also toward the bottom.

For the top left shape in the second diagonal pair, the error was smaller in the structured condition than in the random condition [$t(78) = 2.91$, $p = .005$, $d = .65$] on the X-axis, and also on the Y-axis [$t(78) = 2.18$, $p = .03$, $d = .49$]. This suggests that the top left shape in the diagonal pair was biased toward the right and also toward the bottom.

The previous analysis using signed errors allowed us to examine the recall accuracy of the presented object in each pair. To further examine whether the object was recalled to be closer to its occurring partner, we computed the distance between the judged location of the presented shape and the location of its co-occurring partner (if it were presented) in the judgment task (Figure 2a). We found that the judged location of the presented shape was reliably closer to its co-occurring partner's location in the horizontal pair [$t(78) = 2.12$, $p = .04$, $d = 0.48$], in the vertical pair [$t(78) = 2.14$, $p = .04$, $d = 0.48$], in the first diagonal pair [$t(78) = 2.87$, $p < .01$, $d = 0.65$], and in the second diagonal pair [$t(78) = 4.09$, $p < .001$, $d = 0.91$] in the structured condition than in the random condition. An average distance was computed across the four types of pairs, and distance was again reliably closer in the structured condition than the random condition [$t(78) = 5.04$, $p < .001$, $d = 1.13$]. To examine whether the 2AFC test performance predicted the judgment of the locations, we correlated the accuracy of test performance with the overall average distance in the structured condition. The correlation was not reliable [$r(38) = 0.07$, $p = .68$], suggesting that the test performance was not associated with the magnitude of the bias.

These results consistently showed that if a shape co-occurred with another shape in space, the shape was subsequently represented as being closer to the other shape with which it had co-occurred previously. This suggests that the representation of one shape is biased toward its co-occurring partner.

Experiment 2

The goal of this experiment was to generalize the findings of Experiment 1 from shape pairs to colour pairs by testing how regularities in terms of colour co-occurrences alter the representations of individual objects.

Participants

Fifty-six new undergraduate students (42 female, mean age = 20.3 years, $SD = 1.6$) from University of British Columbia participated for course credit.

Stimuli

The stimuli consisted of 10 coloured circles. The diameter of each circle subtended 1.4° of visual angle. The circles were generated from 10 distinct colours (colour

Figure 2. Distance between the judged location of the presented shape and the location of its co-occurring shape (if it were presented) for (a) Experiment 1 and (b) Experiment 2. The distance was calculated as the absolute distance between the two locations combining the X and Y coordinates. The results were graphed in pixels for each of the four pairs. The overall average distance across all four pair types was also graphed in the right-most column.

name and R/G/B values: red, 255/0/0; green, 0/255/0; blue, 0/0/255; yellow, 255/255/0; magenta, 255/0/255; cyan, 0/255/255; grey, 185/185/185; orange, 255/140/0; brown, 103/29/0; black, 0/0/0). Eight of the 10 circles were randomly assigned, without replacement and for every participant, to four colour pairs, and the remaining two single circles were not paired. As in Experiment 1, the four pairs were grouped into fixed horizontal, vertical, and diagonal configurations (Figure 3a).

For every trial, an array of circles was presented on the screen. The number of circles in the array ranged from three to 10, creating eight levels of numerosity. This ensured that both odd and even numbers were equally presented in the experiment. For odd numbers of three, five, seven, and nine, the array consisted of one, two, three, and four unique colour pairs, respectively, and one single circle. For even numbers of four, six, and eight, half of the time the array consisted of two, three, and four pairs, respectively; the other half of the time, the array consisted of one, two, and three pairs, respectively, with two additional single circles. For the number 10, the array contained all four pairs and two single circles. The array appeared in an invisible 4×4 grid (subtending $12.4° \times 12.4°$) with the constraint that each pair neighboured at least one of the other pairs or one of the single circles (Figure 3a). This constraint ensured that statistical learning could not solely be determined by spatial segmentation cues other than co-occurrence. Each level of numerosity was repeated 50 times, resulting in a total of 400 trials (order randomized for every participant). In both conditions, 20% of the arrays contained a duplicate colour: one circle in a randomly selected pair was changed to match the colour of the other circle from that pair. These duplicates, which existed in the arrays of both conditions, were used in support of a cover task (duplicate detection task).

Procedure

Participants were randomly assigned to either the structured or the random conditions ($N = 28$ in each). The three phases in each condition (exposure, judgment, and test) were identical to those in Experiment 1, except two changes. First, during the exposure

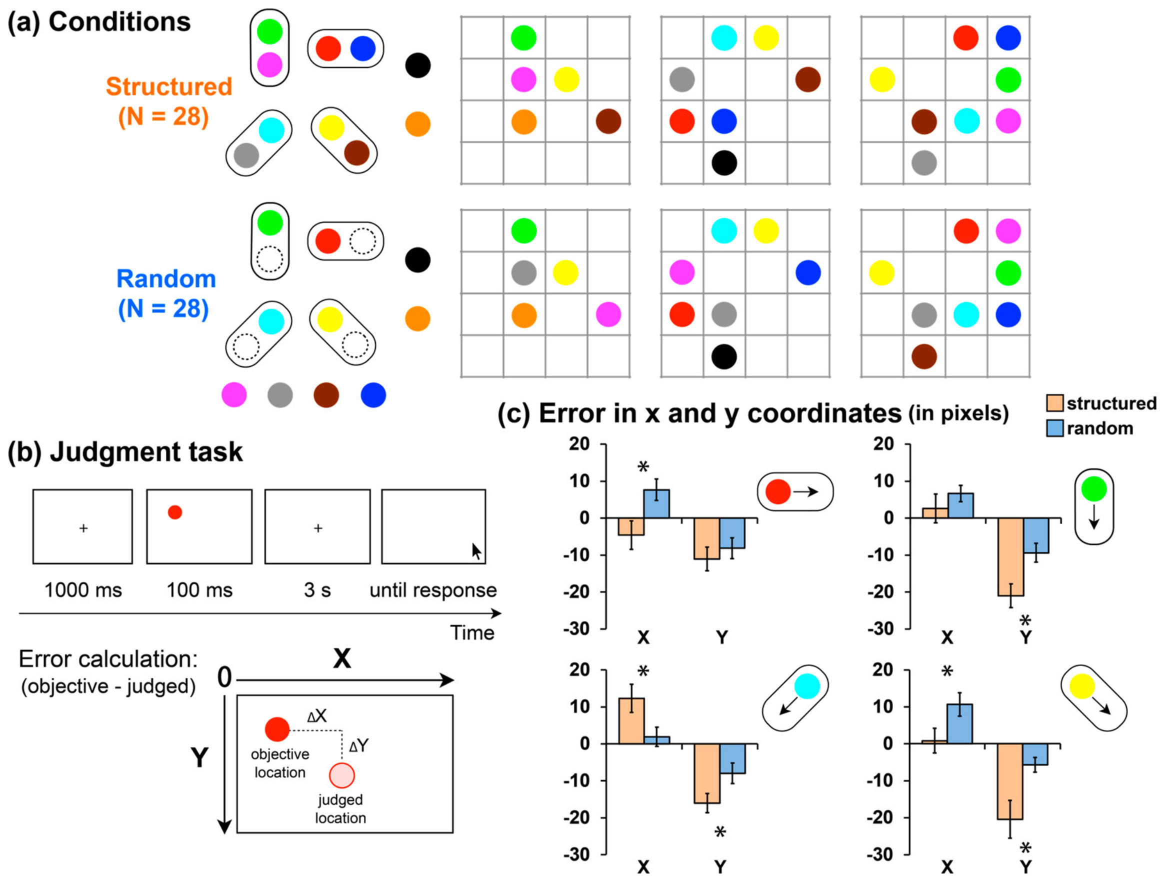

Figure 3. Stimuli, tasks, and results from Experiment 2. (a) In the structured condition, 10 coloured circles were grouped into four colour pairs, with two remaining single circles. The four pairs were assigned to horizontal, vertical, and diagonal configurations. In the random condition, one circle in each pair remained in the same position but the other circle appeared randomly in one of the other three pairs. In each trial, the four pairs and two singles were presented in an invisible 4 × 4 grid. Three sample trials are shown in each condition. (b) In the judgment task, the circle that always remained in the same position within each pair was briefly presented on the screen, followed by a 3 s interval. Participants reported the position of the circle using the cursor. The error was calculated by subtracting the judged position from the objective position of the circle, measured in pixels along the X and Y screen coordinates. (c) Each graph shows the judgment errors in each pair. (Error bars represent ± 1 SEM; *$p < .05$).

phase each array was presented for 1000 ms followed by an ISI of 1000 ms. Second, during the judgment phase, one circle from a given pair appeared on the screen for 100 ms, followed by a blank interval of 3000 ms. The cursor then appeared on the bottom right corner of the screen, and participants used the cursor to indicate the centre of the circle that they just saw (Figure 3b).

Results

During exposure, mean accuracy in the duplicate detection task was 94.9% ($SD = 4.9\%$) in the structured condition and 96.5% ($SD = 2.0\%$) in the random condition. Performance did not differ between the two conditions [$t(54) = 1.60$, $p = .12$, $d = .43$]. At the test phase, in the structured condition pairs were chosen over foils for 58.7% ($SD = 13.9\%$) of the time, which was reliably above chance [50%; $t(27) = 3.32$, $p = .003$, $d = .63$]. During debriefing, some participants reported the colour duplicates as the regularities, but no participant identified any correct colour pair. This suggests that participants implicitly learned the colour pairs during exposure.

The critical question was whether the colour pairs biased the spatial representation of the circle within the pair. As in Experiment 1, we calculated the error in pixels between the objective location and the judged location of each circle presented in the judgment phase (Figure 3b). We then compared the error in each pair between the structured and the random conditions (Figure 3c).

For the left circle in the horizontal pair, the error on the X-axis was smaller in the structured condition than in the random condition [$t(54) = 2.53$, $p = .01$, $d = .68$]. There was no significant difference in the errors on the Y-axis between the two conditions [$t(54) = 0.69$, $p = .50$, $d = .18$]. This suggests that the left circle in the horizontal pair was biased toward the right.

For the top circle in the vertical pair, the error on the Y-axis was smaller in the structured condition than in the random condition [$t(54) = 2.23$, $p = .03$, $d = .59$]. There was no significant difference in the errors on the X-axis between the two conditions [$t(54) = 1.09$, $p = .28$, $d = .29$]. This suggests that the top circle in the vertical pair was biased toward the bottom.

For the top right circle in the first diagonal pair, the error on the X-axis was larger in the structured condition than in the random condition [$t(54) = 2.26$, $p = .03$, $d = .60$]. The error on the Y-axis was smaller in the structured condition than in the random condition [$t(54) = 2.13$, $p = .04$, $d = .57$]. This suggests that the top right circle in the diagonal pair was biased toward the left and also toward the bottom.

For the top left circle in the second diagonal pair, the error was smaller in the structured condition than in the random condition [$t(54) = 2.14$, $p = .04$, $d = .57$] on the X-axis, and also on the Y-axis [$t(54) = 2.68$, $p = .01$, $d = .72$]. This suggests that the top left circle in the diagonal pair was biased toward the right and also toward the bottom.

Again, to provide further support for our hypothesis that the representations of co-occurring circles are biased toward each other, we computed the distance between the judged location of the presented circle and the location of its co-occurring partner (if it were presented) in the judgment task (Figure 2b). We found that the judged location of the presented circle was reliably closer to its co-occurring partner's location in the horizontal pair [$t(54) = 2.39$, $p = .02$, $d = 0.65$], in the vertical pair [$t(54) = 2.39$, $p = .01$, $d = 0.64$], in the first diagonal pair [$t(54) = 3.05$, $p < .01$, $d = 0.82$], and in the second diagonal pair [$t(54) = 3.32$, $p = .001$, $d = 0.90$], in the structured condition than in the random condition. An average distance was computed across the four pairs, and distance was again reliably closer in the structured condition than the random condition [$t(54) = 4.12$, $p < .001$, $d = 1.11$]. To examine whether the 2AFC test performance predicted the judgment of the locations, we correlated the accuracy of test performance with the overall average distance in the structured condition. The correlation was not reliable [$r(26) = 0.10$, $p = .61$], suggesting that the test performance was not associated with the magnitude of the bias.

These results consistently showed that if one colour co-occurred with another colour in space, its location was subsequently judged as being closer to that of the other colour with which it had co-occurred previously. The results fully replicate those in Experiment 1, demonstrating that the representation of one object is biased toward its co-occurring partner.

Experiment 3

The first two experiments showed that one object is represented as being closer to another object if they co-occur in space. Is this finding specific to spatial representations? The final experiment aimed to generalize this effect to a different type of representations, by testing how regularities alter orientation representations.

Participants

Eighty new undergraduate students (61 female, mean age = 19.7 years, $SD = 2.1$) from University of British Columbia participated for course credit.

Stimuli

The stimuli consisted of six lines (each subtending 2.4°) varying in colour and orientation dimensions. There were six colour values (red/green/blue/yellow/orange/purple) and six orientation values (15°/45°/75°/105°/135°/165°). In the structured condition, the six colours and the six orientations were randomly assigned into three pairs for each participant, and the colour and orientation feature sequences were overlaid to create a single line stream (Figure 4a). The stream was generated by pseudorandomly sequencing 80 repetitions of each pair, with the constraint that there was no back-to-back repetition of the same pair. In the random condition, the six orientations were assigned into three pairs for each participant as in the structured condition. However on the colour dimension, the first colour remained in the same first position in each pair, but the second colour randomly appeared in the second position in any of the other three pairs (Figure 4a). In other words, the only difference between the structured and the random conditions was that the second

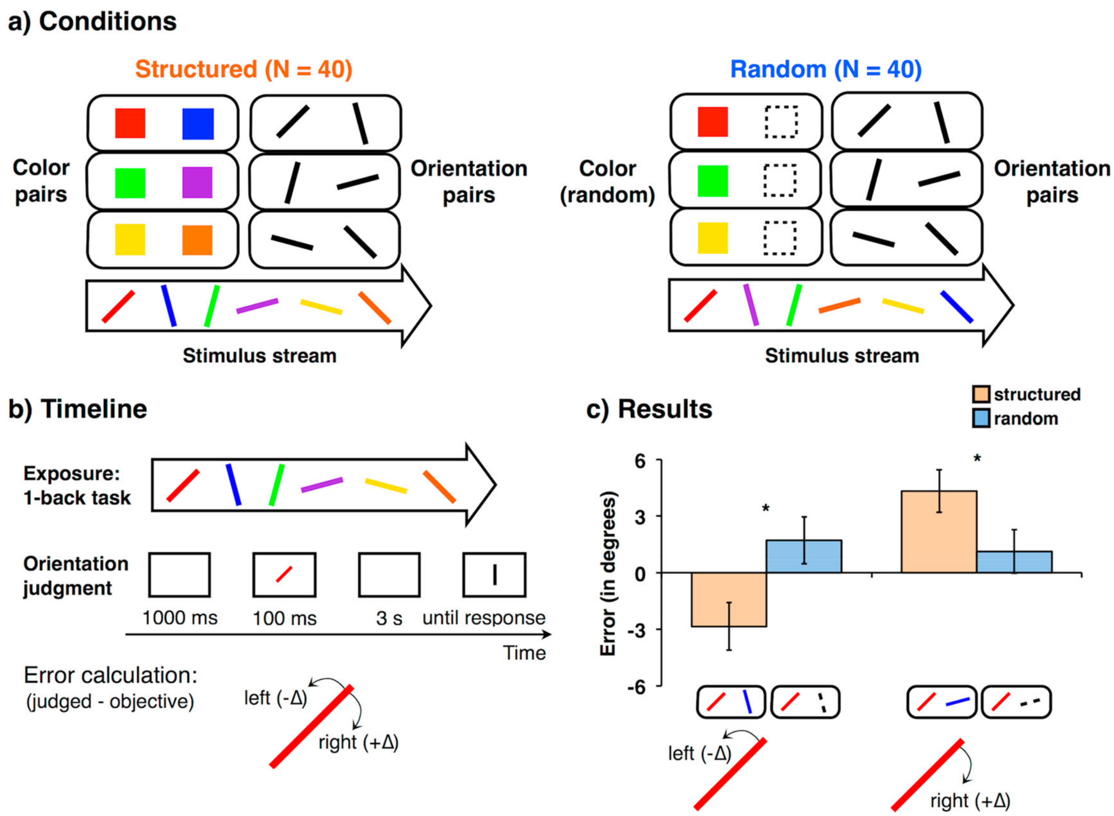

Figure 4. Stimuli, tasks, and results from Experiment 3. (a) In the structured condition, the colours and the orientations of the lines were grouped into three temporal pairs, but in the random condition only the orientations of the lines were grouped into pairs. The colour of the first line in each pair was the same as that in the structured condition, and the colour of the second line appeared randomly. (b) In the exposure phase, participants viewed a temporal sequence containing the line pairs, and performed a 1-back task. In the judgment phase, the first line in each pair appeared briefly, followed by a 3 s interval. Participants reported the orientation of the line by adjusting a black line. The error was calculated by subtracting the judged orientation from the objective orientation of the line, measured in degrees. (c) The graph shows the judgment errors in the orientation of the first line in a pair in the structure and the random conditions. The trials were divided into two groups: in one group, the second line in a pair was tilted to the left of the first line; in the other group, the second line was tilted to the right of the first line. (Error bars represent ± 1 SEM; $*p \leq .05$).

colour reliably followed the first colour within each pair in the structured condition, whereas the second colour was random in the random condition.

Procedure

Participants were randomly assigned to either the structured or the random conditions ($N = 40$ in each). There were three phases in each condition: exposure, judgment, and test. During the exposure phase in both conditions, the lines appeared one at a time at fixation for 750 ms, followed by an ISI of 750 ms. Of all trials, 20% contained a duplicate line identical to the previous line on both colour and orientation dimensions. For every trial starting the second trial, participants performed a cover task (1-back task) where they indicated whether the current line was the same as the previous line by pressing the "z" key for no or the "/" key for yes (key assignment counter-balanced across participants).

During the judgment phase in both conditions, only the line that remained in the first position within each pair was used, and therefore three distinct lines were presented in their original colours and orientations as in the exposure phase. Each of the three lines was repeated 10 times, creating 30 trials in total. The trials were presented in a random order. In each trial, one line was presented at fixation for 100 ms, followed by a blank interval of 3000 ms. A black line was then presented at fixation in vertical orientation, and participants used the cursor to rotate the black line to match the orientation in which the coloured line had just appeared (Figure 4b). No feedback was given on judgment accuracy.

After the judgment phase, participants in the structured condition completed the two-alternative forced-choice test phase. In each trial, participants viewed two sequences of two coloured lines presented at fixation and separated by a 1000 ms pause. Each line appeared for 750 ms followed by a 750 ms ISI. Participants judged whether the first or second sequence seemed more familiar. One sequence was a pair from the exposure phase and the other was a foil composed of two colours that never appeared sequentially. The foils were constructed by choosing the first line from each pair and the second line from the same pair but in a different colour. In other words, the pair and the foil were identical on the orientation dimension, but in the foil the colour of second line never followed the colour of the first line. Pairs were tested against each foil twice, for a total of 12 trials (equating frequency of pairs and foils at test). The order of trials was randomized, and whether the pair or foil appeared first was counterbalanced.

Results

During exposure, mean accuracy in the 1-back task was 83.4% ($SD = 21.4\%$) in the structured condition and 84.8% ($SD = 18.4\%$) in the random condition. Performance did not differ between the two conditions [$t(78) = 0.31$, $p = .76$, $d = .07$]. At the test phase, in the structured condition pairs were chosen over foils for 62.3% ($SD = 19.7\%$) of the time, which was reliably above chance [50%; $t(39) = 3.95$, $p < .001$, $d = .62$]. During debriefing, a few participants reported the line duplicates as the regularities, but no participant identified any correct colour pair. This suggests that participants implicitly learned the colour pairs during exposure.

The critical question was whether the colour pairs biased orientation judgments. To address this question, we calculated the signed error in degrees between the objective orientation and the judged orientation of each line presented in the judgment phase (Figure 4b). If the judged line was tilted to the left of the objective orientation, the error was coded as negative. If the judged line was tilted to the right of the objective orientation, the error was coded as positive. The errors between the structured and the random conditions were shown in Figure 4c.

For each participant, the trials were divided into two groups: in one group, the second line in the pair was tilted to right of the first line; and in the other group, the second line was tilted to left of the first line. This allowed us to examine whether there was a systematic bias in the orientation of the first line toward the orientation of the second line in a pair. Because the lines were randomly grouped into pairs for each participant, for some participants they may not see both groups of trials, and some pairs would contain perpendicular lines just by chance. When the two lines were perpendicular to each other, it was impossible to determine whether there was a left-ward bias or a right-ward bias in the judgment of the presented line to its co-occurring partner. Discarding the data for those participants was merely a result of the random assignment of lines into pairs. Thus, only a subset of the participants in each condition was included in the analysis. Specifically, 32 participants in the structured condition and 26 in the random condition were included in the group where the second line in a pair was tilted to right of the first line. Only 31 participants in the structured condition and 26 in the random condition were included in the group where the second line in a pair was tilted to left of the first line.

We found that when the second line was tilted to the right of the first line in a pair, the error was marginally larger in the structured condition than in the random condition [$t(56) = 1.96$, $p = .05$, $d = .52$]. This suggests that there was a rightward bias of the first line in the structured condition. However, when the second line was tilted to the left of the first line, the error was smaller in the structured condition than in the random condition [$t(55) = 2.56$, $p = .01$, $d = .68$]. This shows that there was a leftward bias of the first line in the structured condition.

To examine whether the 2AFC test performance was indicative of the judgment performance, we correlated the accuracy of test performance with the error in the judged orientation the structured condition. When the line was tilted to the left, the correlation was not reliable [$r(29) = 0.26$, $p = .16$], suggesting that the test performance was not associated with the bias. However, when the line was tilted to the right, there was a negative correlation [$r(29) = -0.41$, $p = .02$], suggesting that the test performance was associated with the magnitude of the bias but in the wrong direction.

These results demonstrated that if one coloured line reliably followed another coloured line over time, its orientation was biased toward the orientation of the other line. Consistent with the first two experiments, the results suggest that regularities biased the representations of individual objects toward their co-occurring partners.

General discussion

This study aimed to elucidate how regularities in terms of object co-occurrences shaped the representations of individual objects. Specifically, we explored whether objects were represented toward each other as a result of statistical learning. To address this question, we first exposed observers with arrays in which two objects reliably co-occurred or with random arrays. After exposure, observers briefly viewed one object on its own and reported the location or the orientation of the object. Across three experiments, we found that if an object reliably appeared with another object in space, its location was subsequently represented as being closer to the location of the other object (Experiments 1 and 2). Likewise, if an object reliably preceded another object in time, its orientation was biased toward the orientation of the other object (Experiment 3). These results demonstrated that statistical learning biased the representations of co-occurring objects toward each other. This finding is highly consistent and robust because we observed the effect in every pair across three independent experiments using different paradigms.

The current findings contribute to the growing literature on how learning guides object representations in memory. Recent work has shown that the representation of an object in visual working memory is shaped by a number of factors, including perceptual expertise (Curby, Glazek, & Gauthier, 2009), prior exposure (Brady et al., 2009), retrieval history from short-term memory (Fan & Turk-Browne, 2013), and explicit cueing of attention (Lepsien & Nobre, 2007; Nobre et al., 2004). The current study demonstrates a learning-induced intrusion from long-term memory in working memory. That is, object regularities have been learned and stored in long-term memory during the exposure phase. In the judgment phase, one object in the pair was briefly presented followed by a 3 s delay. The observers had to hold the object in working memory during the delay, before reporting its location or orientation. The finding that the object was judged as being closer to its absent partner suggests that the prior knowledge about the pair was automatically retrieved from long-term memory, and this knowledge interfered with the online representation of the object in working memory. Building on the previous finding where statistical learning increases the similarity in neural representations of co-occurring objects (Schapiro et al., 2012), the current study shows that statistical learning also increases the similarity of object representations held in working memory.

However, a critical question remains: why is the object representation biased toward its co-occurring partner as a result of statistical learning? One explanation focuses on the implicit activation of the absent object in a pair. Exposure to the reliable co-occurrences between the two objects enhances the perceived predictability between the objects. Indeed, seeing one object in a pair automatically sets up an implicit anticipation of the other object which has not yet been shown (Turk-Browne et al., 2010). This anticipation may activate the representation of the other object in working memory, and therefore biasing the recall of the first object from working memory. For example in Experiment 2, the red circle always appeared next to the blue circle in the structured condition, but in the random condition it appeared next to the blue, purple, grey, or brown circle. Thus, seeing the red circle on its own may automatically activate the blue circle which used to appear nearby. This means that the participant is holding two representations in working memory, the red circle which was shown briefly and the absent blue circle which was activated by the red circle. The representation of the blue circle may thus interfere with the recall of the red circle.

An alternative explanation focuses on the implicit activation of the average feature in a pair. Exposure to the reliable co-occurrences between the two objects may result in implicit grouping of the two objects into one unit, creating a higher-order representation of the unit. This representation may facilitate the automatic extraction of ensemble features of the group, such as the average location of the two objects (Alvarez & Oliva, 2008; Chong & Treisman, 2005). Seeing one object of a pair may activate the unitized representation of the pair and its ensemble features. For example in Experiment 2, after exposure to the reliable pairings of the red circle and the blue circle, the average location of the two circles was automatically extracted. Thus, seeing the red circle on its own may activate the average location of the pair, biasing the representation of the location of the red circle toward the pair average. This account is consistent with the finding that the representation of an individual object within a set is automatically biased toward to the ensemble feature of the set (Brady & Alvarez, 2011). This explanation is also consistent with the probabilistic clustering theory (Orhan & Jacobs, 2013), and raises another factor for clustering, namely, the reliable co-occurrences in space or over time may lead to a joint representation of the two objects, thus introducing dependencies in the joint representation.

A caveat in interpreting the current results was the possibility that the bias in the judgment phase was the consequence of decision processes, rather than the changes in the representations of objects. While the current study does not distinguish between a bias in memory representation and a bias in decisions in the judgment phase, we do think that the bias is more likely to arise from memory representation as a result of learning. This is because of two reasons. First, our participants were not explicitly aware of the regularities or intentionally trying to learn the regularities, showing that statistical learning is largely implicit and incidental, and thus the decision process was minimal in the judgment phase. Second, participants were asked to recall the location of the object which they saw 3 s ago, the bias observed in their judgment was more likely to reflect errors in the memory representation, not from decision processes since it was an explicit decision task among alternatives.

While the current study shows the impact of regularities on object representations, a question remains regarding the extent of such impact. That is, how regularities alter the representations across multiple feature dimensions. For example, in Experiment 1 the two shapes in a pair were represented as being closer in space, but it is currently unknown whether the identities of the two shapes were altered as a result of learning. Specifically, is it the case that the representations of the co-occurring objects were biased toward each other on the shape dimension? Future studies are needed to examine the extent of the representational bias across different feature dimensions.

To conclude, the current findings are significant in several ways. We found a novel impact of statistical learning on object representations. Previous studies have found that the internal representations of individual objects are rendered more salient as their associations are learned (Barakat, Seitz, & Shams, 2013; O'Brien & Raymond, 2012; Wimmer & Shohamy, 2012). Our study extends on the previous findings beyond salience, and demonstrates that regularities in terms of object co-occurrences in space or over time can bias the representations of individual objects toward their co-occurring partners. This finding provides a new perspective on the functions of statistical learning. Statistical association not only involves the reliable co-occurrences between individual objects in space or time, but also serves as a chunking process where two individual objects become more related in their representations. Our findings are also consistent with past memory literature that demonstrates a spatial or temporal contiguity effect where the fidelity of memory representations of objects is affected by spatial or temporal associations. Our study shows that statistical associations (i.e., reliable joint or transitional probabilities) can also shape the representation of individual objects. Moreover, the current paradigm offers a new implicit measure of statistical learning. The judgment task measures the fidelity of object representations in working memory without relying on explicit awareness. The bias in object representations may serve as an implicit signal that underlies the discrimination between the pair and the foil during familiarity test, which has been used as a signature of learning in studies on statistical learning (e.g., Fiser & Aslin, 2001; Saffran et al., 1996; Turk-Browne et al., 2008). Finally, the current findings reveal a new way in which the internal representations of objects can be modified by experience. That is, the reliable co-occurrences learned previously can change the current representations of the objects.

Acknowledgments

For helpful conversations, we thank Nick Turk-Browne, Brian Scholl, Dan Yamins, Sumeyye Cakal, and the Zhao Lab.

Disclosure Statement

No potential conflict of interest was reported by the authors.

Funding

This work was supported by NSERC Discovery Grant (RGPIN-2014-05617 to JZ), the Canada Research Chairs program (to

JZ), the Leaders Opportunity Fund from the Canadian Foundation for Innovation (F14-05370 to JZ), and the NSERC Canada Graduate Scholarship Master's program and Elizabeth Young Lacey Fellowship (to RY).

References

Alvarez, G. A., & Oliva, A. (2008). The representation of simple ensemble visual features outside the focus of attention. *Psychological Science*, *19*, 392–398.

Barakat, B. K., Seitz, A. R., & Shams, L. (2013). The effect of statistical learning on internal stimulus representations: Predictable items are enhanced even when not predicted. *Cognition*, *129*, 205–211.

Batterink, L. J. (2017). Rapid statistical learning supporting word extraction from continuous speech. *Psychological Science*, *28*, 921–928. doi:10.1177/0956797617698226

Brady, T. F., & Alvarez, G. A. (2011). Hierarchical encoding in visual working memory: Ensemble statistics bias memory for individual items. *Psychological Science*, *22*, 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*, 487–502.

Chong, S. C., & Treisman, A. (2005). Statistical processing: Computing the average size in perceptual groups. *Vision Research*, *45*, 891–900.

Conway, C. M., & Christiansen, M. H. (2005). Modality-constrained statistical learning of tactile, visual, and auditory sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 24–39.

Curby, K. M., Glazek, K., & Gauthier, I. (2009). A visual short-term memory advantage for objects of expertise. *Journal of Experimental Psychology: Human Perception and Performance*, *35*, 94–107.

Faber, T., Joerges, J., & Menzel, R. (1999). Associative learning modifies neural representations of odors in the insect brain. *Nature Neuroscience*, *2*, 74–78.

Fan, J. E., & Turk-Browne, N. B. (2013). Internal attention to features in visual short-term memory guides object learning. *Cognition*, *129*, 292–308.

Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, *12*, 499–504.

Graf Estes, K., Evans, J. L., Alibali, M., & Saffran, J. R. (2007). Can infants map meaning to newly segmented words? Statistical segmentation and word learning. *Psychological Science*, *18*, 254–260.

Hund, A. M., & Plumert, J. M. (2002). Delay-induced bias in children's memory for location. *Child Development*, *73*, 829–840.

Hund, A. M., & Plumert, J. M. (2003). Does information about what objects are influence children's memory for where objects are? *Developmental Psychology*, *39*, 939–948.

Hund, A. M., Plumert, J. M., & Benney, C. J. (2002). Experiencing nearby locations together in time: The role of spatiotemporal contiguity in children's memory for location. *Journal of Experimental Child Psychology*, *82*, 200–225.

Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, *98*, 352–376.

Lepsien, J., & Nobre, A. C. (2007). Attentional modulation of object representations in working memory. *Cerebral Cortex*, *17*, 2072–2083.

Li, W., Howard, J. D., Parrish, T. B., & Gottfried, J. A. (2008). Aversive learning enhances perceptual and cortical discrimination of indiscriminable odor cues. *Science*, *319*, 1842–1845.

McNamara, T. P. (1986). Mental representations of spatial relations. *Cognitive Psychology*, *18*, 87–121.

McNamara, T. P., Halpin, J. A., & Hardy, J. K. (1992). Spatial and temporal contributions to the structure of spatial memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 555–564.

Nobre, A. C., Coull, J. T., Maquet, P., Frith, C. D., Vandenberghe, R., & Mesulam, M. M. (2004). Orienting attention to locations in perceptual versus mental representations. *Journal of Cognitive Neuroscience*, *16*, 363–373.

O'Brien, J. L., & Raymond, J. E. (2012). Learned predictiveness speeds visual processing. *Psychological Science*, *23*, 359–363.

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

Perruchet, P., & Pacton, S. (2006). Implicit learning and statistical learning: One phenomenon, two approaches. *Trends in Cognitive Sciences*, *10*, 233–238.

Recker, K. M., & Plumert, J. M. (2009). How Do opportunities to view objects together in time influence children's memory for location? *Journal of Cognition and Development*, *9*, 434–460.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*, 1926–1928.

Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, *70*, 27–52.

Schapiro, A. C., Kustner, L. V., & Turk-Browne, N. B. (2012). Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Current Biology*, *22*, 1622–1627.

Stager, C. L., & Werker, J. F. (1997). Infants listen for more phonetic detail in speech perception than in word-learning tasks. *Nature*, *388*, 381–382.

Turk-Browne, N. B., Isola, P. J., Scholl, B. J., & Treat, T. A. (2008). Multidimensional visual statistical learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*, 399–407.

Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The automaticity of visual statistical learning. *Journal of Experimental Psychology: General*, *134*, 552–564.

Turk-Browne, N. B., Scholl, B. J., Chun, M. M., & Johnson, M. K. (2009). Neural evidence of statistical learning: Efficient detection of visual regularities without awareness. *Journal of Cognitive Neuroscience*, *21*, 1934–1945.

Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit perceptual anticipation triggered by statistical learning. *The Journal of Neuroscience*, *30*, 11177–11187.

Umemoto, A., Scolari, M., Vogel, E. K., & Awh, E. (2010). Statistical learning induces discrete shifts in the allocation of working memory resources. *Journal of Experimental Psychology: Human Perception and Performance*, *36*, 1419–1429.

Wallis, G., & Bülthoff, H. H. (2001). Effects of temporal association on recognition memory. *Proceedings of the National Academy of Sciences*, *98*, 4800–4804.

Wimmer, G. E., & Shohamy, D. (2012). Preference by association: How memory mechanisms in the hippocampus bias decisions. *Science*, *338*, 270–273.

Yu, R. Q., & Zhao, J. (2015). The persistence of the attentional bias to regularities in a changing environment. *Attention, Perception and Psychophysics*, *77*, 2217–2228.

Zhao, J., Al-Aidroos, N., & Turk-Browne, N. B. (2013). Attention is spontaneously biased toward regularities. *Psychological Science*, *24*, 667–677.

Zhao, J., & Luo, Y. (2017). Statistical regularities guide the spatial scale of attention. *Attention, Perception, & Psychophysics*, *79*, 24–30.

Zhao, J., Ngo, N., McKendrick, R., & Turk-Browne, N. B. (2011). Mutual interference between statistical summary perception and statistical learning. *Psychological Science*, *22*, 1212–1219.

Zhao, J., & Yu, R. Q. (2016). Statistical regularities reduce perceived numerosity. *Cognition*, *146*, 217–222.

Zhao, L., Cosman, J. D., Vatterott, D. B., Gupta, P., & Vecera, S. P. (2014). Visual statistical learning can drive object-based attentional selection. *Attention, Perception, & Psychophysics*, *76*, 2240–2248.