Reward impacts visual statistical learning

Su Hyoun Park[1], Leeland L. Rogers[1], Matthew R. Johnson[2], and Timothy J. Vickery[1]

1. Department of Psychological and Brain Sciences, University of Delaware
2. Department of Psychology, University of Nebraska-Lincoln

Author Note

This research was supported by

National Sciences Foundation grants BCS 1558535 and OIA 1632849

Please Address Correspondence to:

Su Hyoun Park

Department of Psychological and Brain Sciences

University of Delaware

108 Wolf Hall

Newark, DE 19716

e-mail: suhyounp@udel.edu

Abstract

Humans automatically detect and remember regularities in the visual environment—a type of learning termed visual statistical learning (VSL). Many aspects of learning from reward resemble statistical learning in respects, yet whether and how reward learning impacts VSL is largely unexamined. In two studies, we found reward contingencies affect VSL, with high-value associated with stronger behavioral and neural signatures of such learning than low-value images. In experiment 1, subjects learned values (high or low) of images through a trial-and-error risky choice task. Unbeknownst to them, images were paired as four types—High-High, High-Low, Low-High, and Low-Low. In a subsequent recognition task and reward memory task, subjects chose the more familiar of two pairs (a target and a foil) and recalled the value of images. We found better recognition when the first images of pairs have high-values, with High-High pairs showing the highest recognition rate. In experiment 2, we provide evidence that brain responses were affected by both value and statistical contingencies. When we compared responses between the high-value first image and the low-value first image, greater activation in regions that included inferior frontal gyrus, anterior cingulate gyrus, hippocampus, orbitofrontal cortex, among other regions were found. These findings were driven by the interaction between statistically structured information and reward—the same value contrast yielded no regions for second-image contrasts and for singletons. Our results suggest that high-value first images were involved in greater attentional engagement, potentially enabling better memory for statistically learned pairs and reward information than low-value first images.

Keywords: reward, visual statistical learning, reward motivation, memory, fMRI

INTRODUCTION

Reward motivation impacts human cognition in many contexts (Haber & Knutson, 2010). Value is linked to stimuli that are critical for individuals' survival (e.g., primary reward; water or food), but learned associations between reward and neutral stimuli can also shape one's behavior (e.g., secondary reward; money; Daw & Doya, 2006). There is vast literature demonstrating how secondary cues, especially monetary reward, guide an individual's cognitive processes such as memory, attention, and decision making. Higher associated value facilitates stimulus-reward memory association (Adcock et al., 2006), and features and objects that are associated with higher value capture more attention than those with low- or no rewards (e.g., Anderson, 2013; Theeuwes & Belopolsky, 2012). Individuals' decision-making tends to optimize action so that rewards are maximized and losses minimized (Tversky & Kahneman, 1979). However, the relationship between learning and reward is typically studied in the context of learning rewarding associations, specifically, or memory of individual stimuli that are explicitly or implicitly associated to reward (Miendlarzewska et al., 2016). In the present study, we examine how learning explicitly about rewarding associations modulates the undirected and uncued learning of visual statistical associations.

Visual statistical learning (VSL) is a type of learning that reflects automatic and unsupervised extraction of statistical contingencies by the visual system (Fiser & Aslin, 2001, 2002). Prior studies suggest that humans may, in part, accomplish efficient processing of complex visual environments by learning and exploiting knowledge of visual regularities (Fiser & Aslin, 2001, 2002; Turk-Browne et al., 2005). In two early VSL studies, Fiser and Aslin (2001, 2002) found that when particular visual items co-occurred with others, subsequent recognition rates of those regularities were above chance, even though those regularities were

task-irrelevant, no instructions to remember the associations were given, and the associations were not cued. A typical VSL paradigm takes place in the context of passive viewing or simple cover tasks. How VSL occurs in the context of different task demands and contexts, as it must occur in everyday life, is underexplored. Since intentional seeking and learning about rewards is so foundational to behavior, it is natural to ask how learning about reward might impact incidental learning of regularities.

What mechanisms might drive impacts of reward on VSL? Given its known relationship with both reward learning and VSL, selective attention might play a key role. Selective attention is drawn to stimuli associated with high reward compared to low or no reward (e.g., Anderson et al., 2013; Theeuwes & Belopolsky, 2012) and is also required to process statistically structured information (Baker et al., 2004; Turk-Browne et al., 2005). Baker et al. (2004) found that visual regularities were not learned in the absence of selective attention. Turk-Browne et al. (2005) found the learning of regularities occurred only with an attended color stream when participants were exposed to an interleaved stream that composed of attended- and unattended-color (but see Musz et al., 2015). Selective attention may play a critical role in learning visual regularities, with the degree of selective attention modulating the strength of learning.

Brain imaging studies provide further evidence of the possible role of selective attention in VSL. The lateral occipital cortex (LOC) is known for its role in object perception (Grill-Spector et al., 2001; James et al., 2003), but previous studies also showed greater LOC activation to attended relative to distractor or ignored objects (Vuilleumier et al., 2005; Woolgar et al., 2015). Considering prior evidence that attention modulates response patterns in higher visual areas (e.g., Murray & Wojciulik, 2004), LOC may be related to attentional processing during object perception (see also Stokes et al., 2009). In previous studies, greater LOC activity was

shown when exposed to visual regularity (Turk-Browne et al., 2009) and rewarded items (Anderson, 2017). As these findings imply that reward has an influence on the allocation of attention (Theeuwes & Belopolsky, 2012), and that VSL is enhanced by increased selective attention (Turk-Browne et al., 2005), we predict that reward might impact VSL via enhanced attentional processing of higher-reward items in comparison to low-reward items.

We further predict that the effect of reward might be especially potent when the high reward item is in the first position in a temporally presented pair sequence. In VSL, the position of an item in a stereotyped sequence seems to determine the neural response profile to that item. In one of the earliest neural studies of VSL, researchers used face and scene stimuli to investigate implicit perceptual anticipation in the hippocampus (Turk-Browne et al., 2010). They examined anticipatory responses in the hippocampus while participants made a categorical response to face or a scene pictures that appeared one at a time. Unbeknownst to participants, each run was constructed from four pairs of images and four single random images. The right anterior hippocampus and medial temporal lobe showed enhanced responses when the first picture of a pair appeared (i.e. predicting the stimuli) as compared to novel singletons, which suggest that during the acquisition of statistical regularities, the first item of the structured information plays its role in predicting and evaluating subsequent items. Therefore, when reward is embedded in VSL sequences, reward may evoke different responses according to the position of the structured information it is associated with. We also predicted that higher reward that is specifically associated with early items in a temporal sequence would aid visual statistical learning. If attentional processing is involved in this interaction, greater activations may be found in brain regions in frontal and parietal areas, such as inferior frontal gyrus, precentral gyrus, and anterior cingulate gyrus, that are known for their roles in attentional capture (Beck & Vickery,

2019; Corbetta & Shulman, 2002; Fockert et al., 2004), in addition to the LOC. To clearly see the interaction between varying rewards (i.e., high vs. low) and the position of an item in a structured sequence, we used pairs presented in temporal succession to instantiate statistical regularities, but pairs were constructed with different reward variations (i.e., High-High, High-Low, Low-High, and Low-Low).

To our knowledge, Rogers et al. (2016) is the only work to examine the relationship between monetary reward and VSL directly. Despite finding evidence of visual statistical learning, the amount of reward associated with stimuli and sequences did not affect the strength of VSL in their studies, suggesting that reward processing and VSL were operating independently. However, the manipulation of reward, in that case, may have been too subtle for participants to process reward contingencies in a VSL paradigm. Therefore, to motivate learning and enhance participants' performance, we employed a risky choice task (e.g., Clark et al., 2009), which is more likely to lead to in-depth processing of reward information. With this manipulation, we expected that participants would be more engaged in the task, and enhanced learning would be observed for both reward and statistical information.

In the present study, we examined how reward modulates VSL with behavioral and neural approaches. We asked how reward variations affect the learning of statistical regularities and probed the underlying neural mechanisms of our finding that reward associations do, in fact, shape VSL. In Experiment 1, we found higher recognition rates for pairs when the first image of a pair had a high-value, which suggested the high value of the first item in a pair enhances learning (or low-reward impairs learning). Neural evidence from Experiment 2 supported this finding such that when the high reward was associated with the first item, as compared to when the low reward was associated with the first item, greater activation was observed in inferior

frontal gyrus, left anterior cingulate gyrus, middle temporal gyrus, hippocampus, OFC, and LOC, among other regions. This difference was not observed for second items of a pair, nor for singleton items unassociated with other images. This may be interpreted as comparatively more attentional resources being recruited to high reward items, leading to enhanced learning in both value information and visual regularity. Thus, reward may play a role similar to selective attention in VSL, or it may affect VSL by shaping selective attention.

EXPERIMENT 1

The aim of Experiment 1 was to examine the influence of learned value on VSL by embedding different amounts of reward into structured pairs (i.e., High-High, High-Low, Low-High, and Low-Low reward pairs) that always co-occurred temporally in a sequence of decisions. After subjects learned the value in a temporally structured sequence, we tested recognition for each type of pair, allowing us to examine how the high- or low-reward association might interact with the location of reward (i.e., first or second) in structured pairs.

Method

Participants

All procedures were approved by the University of Delaware Institutional Review Board. Thirty-three University of Delaware students who were 18-40 years of age (23 Female) participated for course credit or cash. At the last phase of Experiment 1, participants' memory for the image value was measured. Pilot data suggested that reward memory recognition judged by the third phase was almost always above chance levels. Since it was crucial for participants to have a memory of reward associated with constituent items to judge reward effects on VSL, we

established exclusion criteria based on last-phase performance. Two participants were excluded because they did not show above chance (50%) reward memory recognition rate.

Stimuli and Apparatus

Experiment 1 was run on Windows 10 with a 24-inch LCD monitor with a resolution of 1920 x 1080. The experiment was programmed in MATLAB with Psychophysics Toolbox v. 3 (Brainard, 1997; Kleiner et al., 2007). We used 32 fractal images as novel visual stimuli. Images were randomly assigned into structured sequences (i.e. pairs) between participants. Stimuli were 200 pixels x 200 pixels, and participants sat approximately 57 cm from the monitor (images subtended approximately 5° of visual angle).

Procedure

The experiment consisted of three phases. Participants performed 1) a learning phase followed by 2) a surprise pair recognition phase. In the last phase, they completed 3) a reward memory test, which asked participants to explicitly recall the value of each image (i.e., high or low; 2AFC). Before the experiment began, participants were given instructions about the learning task. However, no information was provided to participants about the subsequent memory-test phases prior to completing the learning phase.

During the learning phase (Fig 1A), images were presented at the center of the screen, sequentially. Participants were instructed to do a risky choice task, in which they learned the values (high or low) of fractal images through trial-and-error. For each image, participants needed to make a choice (phrased as a “gamble”) of “Yes” or “No.” If they chose "Yes” (press the Z button on the keyboard), they had a 50% chance of winning nothing (0 points) and a 50% chance of winning points. Importantly, "high-reward" images were associated with a 50% chance

to win 10 points, while "low-reward" images were associated with a 50% chance to win 2 points. If they chose "No" (pressed the M button on the keyboard), they always got 1 point and, importantly, were able to see what they could have gained (i.e., 0, 2, or 10) if they chose "Yes" on that trial. This way, they were still able to learn 1) the associated value (if 2 or 10 points were assigned on that trial) and 2) whether they won by not choosing “Yes” on that trial (if 0 was assigned on that trial). Participants were told that the points added up over time and they would get money based on their point totals. At the end of the experiment, the points were converted to maximum of $10 (i.e., total points (maximum of 3200) were divided by 320). Subjects were informed beforehand that points would be converted to money at the end of the experiment, but not of the exact conversion rate. If they could not choose within 2 seconds, it was counted as "Miss."

Unbeknownst to participants, we paired images so that some images always predicted other images on the following trial. This led to four types of pairings (High-High, High-Low, Low-High, and Low-Low) (Fig. 1B). All structured pairs were pseudo-randomized within the stream such that no immediate repetition of a pair (e.g., ABAB) or two sets of pairs (e.g., ABEFABEF) could occur. The 32 fractal images (16 pairs) were repeated four times within each block. With a total of 5 blocks, each image/pair appeared a total of 20 times. The 16 pairs were equally divided into four of each of the pairing conditions.

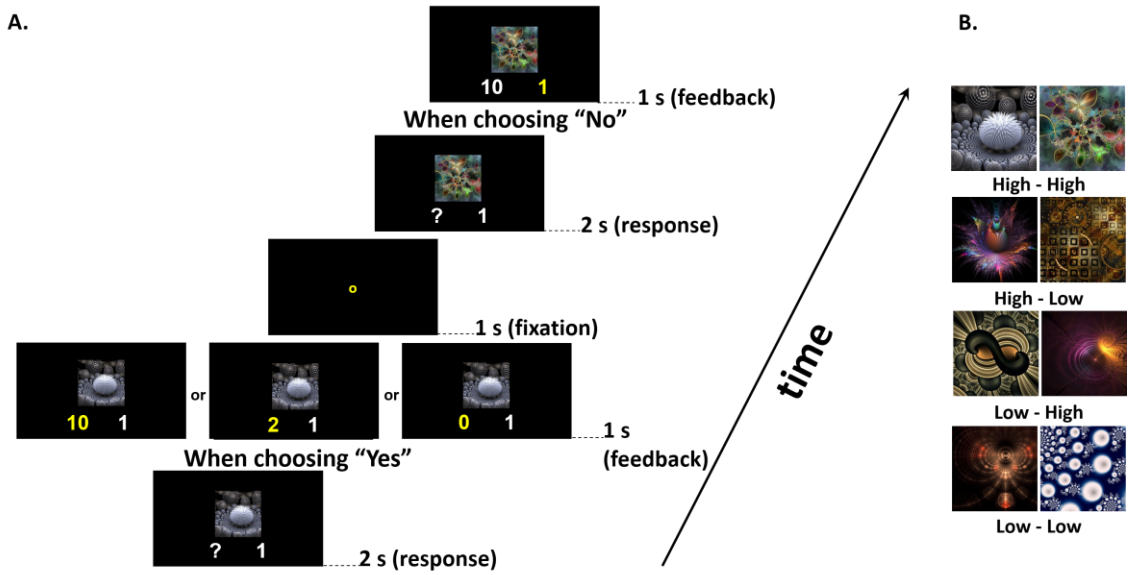

Figure 1. A) General procedure of the learning phase in Experiment 1. B) Pairs were equally divided into four reward variations (High-High, High-Low, Low-High, and Low-Low).

Following the learning phase, the recognition phase began. Participants were given on-screen instructions before they began the recognition phase. This phase involved a two-alternative forced-choice task in which participants were asked to choose which of two two-image sequences was more familiar (Fig. 2A). One of the sequences was a sequence of a target pair, and the other one was a sequence of a foil pair. The target pair was a structured pair that was presented multiple times during the learning phase (e.g., AB, CD, EF, etc.). Foil pairs were recombined from pairs constructed from using the first image of one target pair and the second image of another target pair (e.g., AD). Each target and foil pair were presented four times during the test phase. We constrained each target pair type (in terms of reward) to match with all types of foil pairs (e.g., High-High (target) vs. High-High (foil); High-Low (foil); Low-High (foil); Low-Low(foil)) in each presentation. No feedback was given during this phase, and participants had unlimited time to respond.

After the recognition phase, participants were asked to remember the value of all images that they saw during the learning phase and choose whether they had high or low-values in a two-alternative forced-choice paradigm. All 32 images were presented one by one in a random order (Fig. 2B), with no time constraints and no feedback provided.

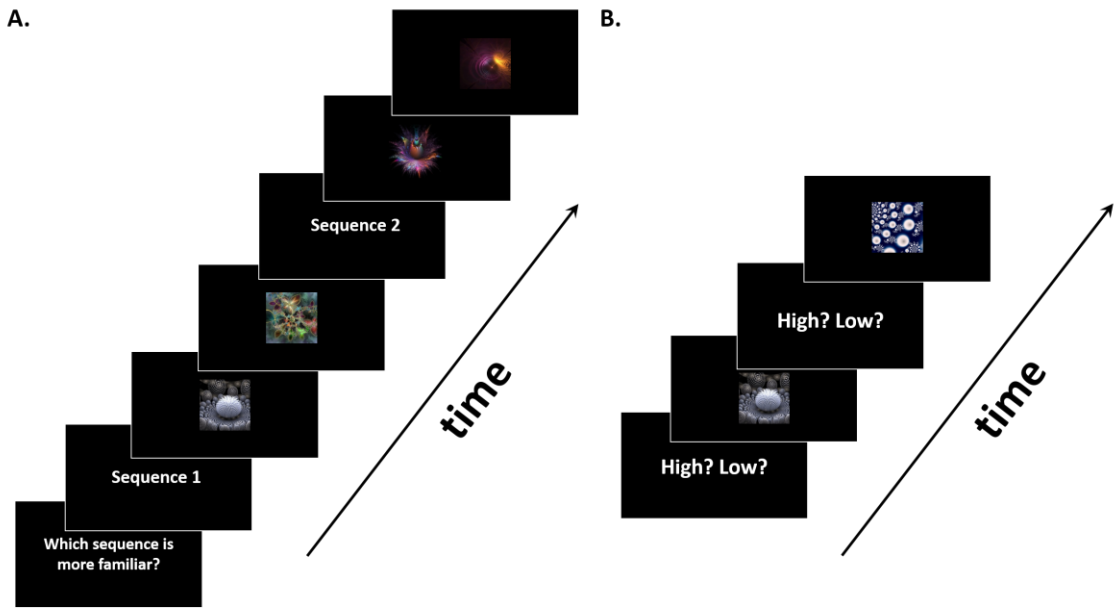

Figure 2. General procedure of the memory tests. (A) Example of the recognition test. (B) Example of the reward memory test.

RESULTS

A two-way repeated measures ANOVA (value of image x block) on risky choice proportion (i.e., choosing yes) showed a significant main effect of value of image, $F(1, 30) = 48.04$, $p < .001$, $\eta p^2 = .616$ (but not with blocks, $F < 1$) and an interaction between them, $F(4, 120) = 12.48$, $p < .001$, $\eta p^2 = .3$. Proportion of making a risky choice to high-value images gradually increased across blocks, and the opposite was observed with low-value images (Fig 3).

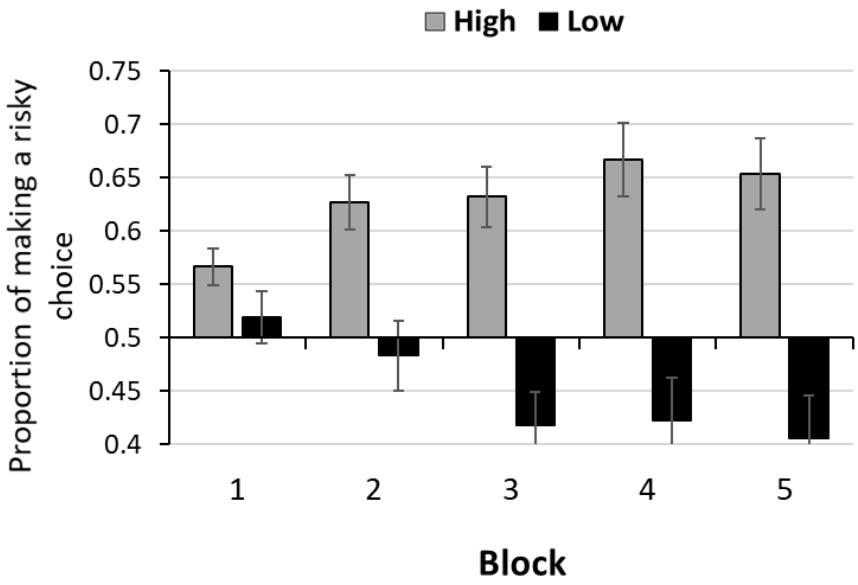

Figure 3. Proportion of making a risky choice throughout blocks. In this and all other figures, error bars represent standard error of the mean.

In regards to the recognition phase, a one-sample t-test against chance (50%) yielded significant learning only for the High-High condition, t(30)= 2.71, p =.01, d=.49. In addition, with a 2 (value of first image, high or low) x 2 (value of second image, high or low) repeated measures ANOVA, we only found a significant main effect of the first image such that there was better recognition when the first image of a pair was a "High" image, F(1, 30)=6.41, p = .01, ηp^2 = .17 (Fig. 4). To ensure that results were not impacted by foil pair value, we conducted a 2 (value of first image, high or low) x 2 (value of second image, high or low) repeated measures ANOVA based on foil type, which resulted in no significant main effects and no interaction of foil value on recognition accuracy, all F <1.

In the last reward memory phase, the mean proportion correct of reward memory was 0.75 (SD: 0.01, t(30)=12.49, p < .001, Cohen's d=2.24 ; one-sample t-test against chance (50%)). A 2 (first image of a pair or second image of a pair) x 2 (high reward or low reward) repeated

measures ANOVA did not reveal any significant main effects or interaction between image position and value (all p>.5) in the last reward memory phase.

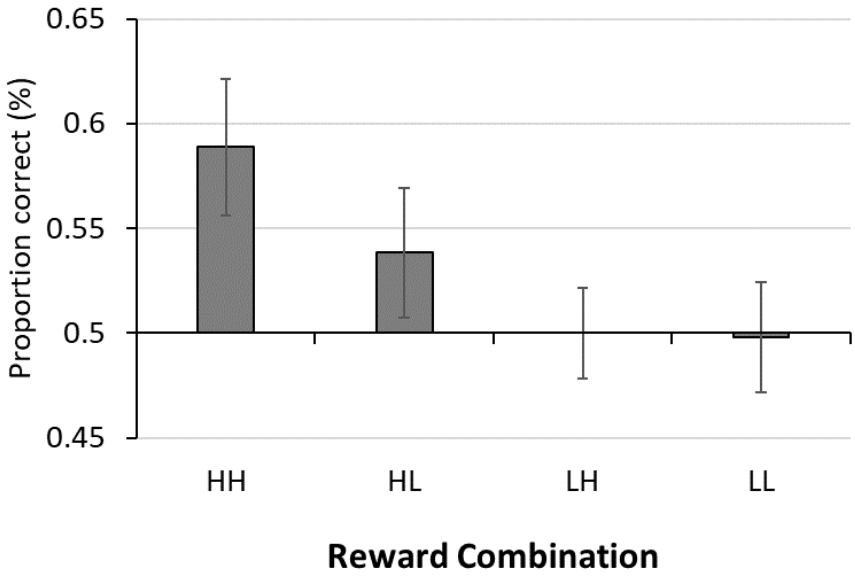

Figure 4. Accuracy at choosing target pairs over foil pair in four reward variations.

DISCUSSION

Previous research found no differences in VSL amongst no-, low-, or high-reward conditions (Rogers et al., 2016). However, previous efforts did not explicitly draw attention to value during exposure to statistical associations. In the current study, using a risky choice task, the subjects' task was to learn the value of images, which drew attention explicitly to reward during exposure. Under these constraints, we found better recognition for pairs when the first image of the pair was a high-reward image.

A number of mechanisms might explain this finding, with variations in attention caused by associated value being one candidate. In VSL, the first item of structured pairs plays an important role in predicting and evaluating subsequent outcomes during the acquisition of statistical regularities (Turk-Browne et al., 2010). Reward could impact VSL by drawing intense

attention to the high-reward image that was located in the first position of a pair. As we do not see any benefit for pairs where the high-reward image appeared second (i.e., Low-High pairs), we speculated that value information might interact with VSL because attention is engaged with greater frequency and/or intensity when the first image of a pair is associated with high-reward, in advance of the predictable second image. This in turn enables learning of the association. On the other hand, if the first image of a pair did not receive such priority (i.e., the low value first image), VSL may not be fully engaged. In Experiment 2, we conducted an fMRI study to uncover the neural correlates of the interaction between reward (i.e., high vs. low) and the position of an image in a structured sequence (first vs. second vs. non-structured information), and how this interaction is modulated by attentional processing.

EXPERIMENT 2

Experiment 1 showed better recognition for pairs when the first image of a pair was associated with high-value. To investigate the neural basis of the findings in Experiment 1, we measured brain responses to visual images that were associated with both varying levels of reward and sequential contingencies, using event-related fMRI. We examined the neural activation of the first and the second image in pairs, and how it differed according to the amount of reward (high vs. low). We also compared images with structural information (i.e., pairs) and without such information (i.e., singletons) in each of high and low-value (e.g., high paired images vs. high singleton; low paired images vs. low singleton), and asked how the varying level of reward affected the processing of statistically structured information.

METHOD

Participants

Thirty University of Delaware students who were 18-40 years of age each participated in one 2-hour long experimental session (mean age: 21.6; 22 females). One participant did not show above chance levels of learning in the last reward memory phase, so that participant was excluded. All participants were right-handed, reported having normal color vision, and were compensated $20/hour. All procedures were approved by the University of Delaware Institutional Review Board.

Stimuli, Apparatus, and Procedure

Stimuli were the same as Experiment 1 except that we added 16 more fractal images as singletons, so a total of 48 fractal images were used in Experiment 2 (32 images assigned to 16 pairs, and 16 singleton images). There were 4 runs of the risky choice task (i.e., learning phase) and in each run, a new set of 4 pairs and 4 singletons were presented, with each repeating six times within the block. We chose six repetitions based on prior studies showing evidence of learning even with a small number of repetitions (Turk-Browne et al., 2010). The added singletons allowed us to directly compare the differences in neural activity for images that contained statistical structure information and images that do not.

The procedure was similar to Experiment 1, but participants performed the risky choice task inside of the scanner. We included jittered intervals between 1) the choice phase and feedback phase of the trial and 2) the feedback phase of the trial and the next image presentation (Fig. 5). Jittered intervals consisted of 2s, 3s, 4s, or 5s., and they were evenly divided across conditions and randomly presented. During the risky choice task phase, participants responded with an MRI-compatible button box. The rules of the task were identical to Experiment 1. Following the four learning runs, a passive viewing run was performed. In this run, all 48 images were presented one more time, with each of the 16 pairs presented in pair-wise order and 16

singletons randomly presented in between pairs. Participants were asked to focus on each image but otherwise passively view them. Each image was presented for 1 s followed by a jittered interval [2s, 3s, 4s, or 5s] [1]. After all runs, participants completed the recognition test and reward memory test outside of the scanner. The procedures for the recognition and reward memory tests were the same as Experiment 1, and all 48 images (including singletons) were shown in the reward memory test. The incentive was provided based on the points participants earned during the risky choice task, and points were converted to a maximum of $15 (i.e., total points (maximum of 900) were divided by 60). Participants were informed at the beginning of the experiment of the possible reward and that points would be converted to cash rewards.

[1] Due to time constraints, twenty-one participants performed one run of passive viewing, and eight participants performed three runs of passive viewing.

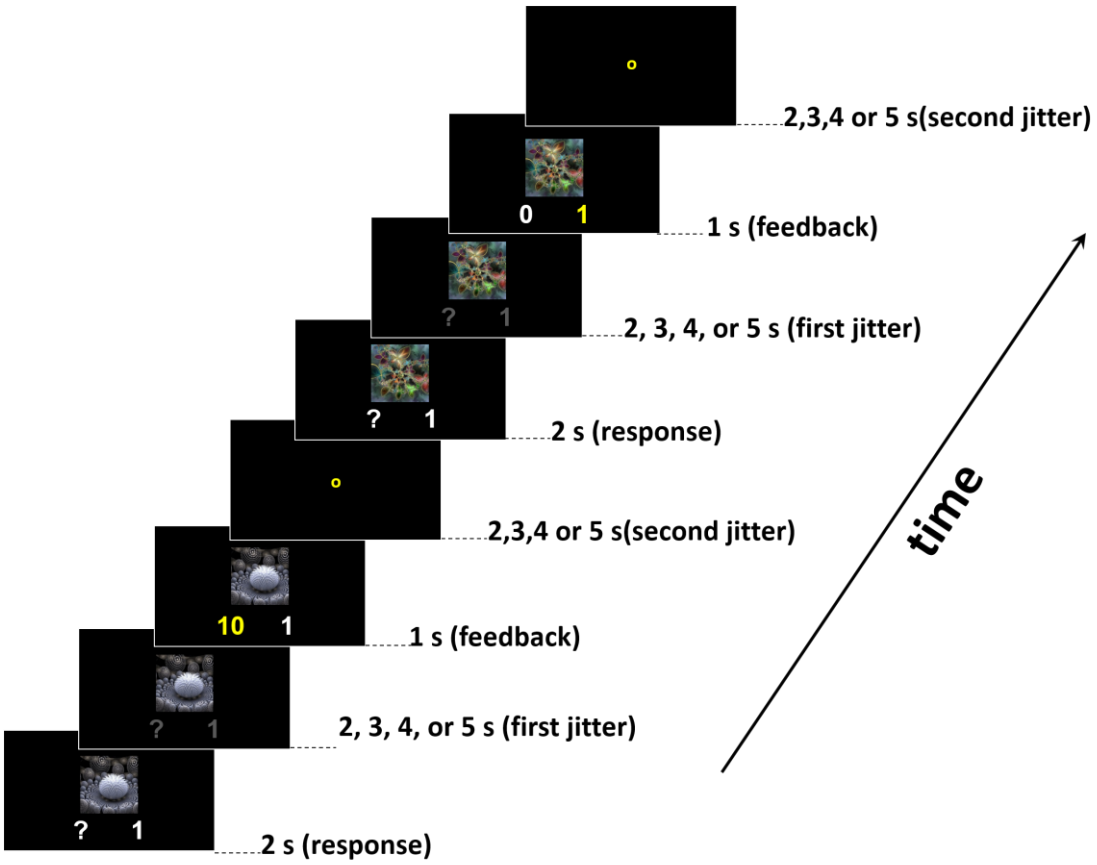

Figure 5. General procedure of the learning phase in Experiment 2.

Data acquisition

Neuroimaging data were acquired on a 3T Siemens Prisma system using a 64- channel head/neck coil. One high-resolution T1-weighted MPRAGE structural image was collected (0.7 mm isotropic voxels). Functional scans consisted of a T2*-weighted Siemens Multiband (multiband factor of 8) EPI sequence with 80 slices acquired in an interleaved manner, and with an oblique axial orientation (approximately 25° from anterior commissure/posterior commissure line). The in-plane resolution was 2.0 mm x 2.0 mm, and slice thickness was 2.0 mm with no skip (TR=1 s, TE = 32 ms, flip angle 61°), resulting in isotropic voxels. Each learning run consisted of 784 volumes and lasted 13 minutes and 4 seconds, and the passive viewing run contained 237 volumes and lasted 3 minutes and 57 seconds.

Structural and Functioning Processing

Data analyses were performed using fMRIB Software Library (FSL, www.fmrib.ox.ac.uk/fsl) version 5.0.9, FMRI Expert Analysis Tool (FEAT) version 6.0 (Jenkinson et al., 2012), and the AFNI software package (Cox, 1996). For structural scans, we first performed skull-stripping by using BET (Smith, 2002), and then registered to a standard MNI152 2-mm template. For functional runs, data were first de-obliqued (AFNI's 3dWarp) and re-oriented to match the standard template (fslreorient2std). Then, data were motion corrected, smoothed (8 mm FWHM Gaussian kernel), and high-pass temporal filtered with a 100s cutoff.

At the first-level analysis of the risky choice task phase, a total of 116 runs (4 runs, 29 participants) were modeled using a standard GLM approach. Fifteen explanatory variables (EVs) were set up: HH-First, HH-Second, HL-First, HL-Second, LH-First, LH-Second, LL-First, LL-Second, High-Singleton, Low-Singleton, Choice-Yes-Win, Choice-Yes-Lose, Choice-No-Win, Choice-No-Lose, and the first presentation of each image as a regressor of no interest. The first presentation of all images was not included in the reward/location variables, because there had been no opportunity to learn either associated value or statistical contingency. For the passive viewing task, a total of 45 runs (1 run: 21 participants; 3 runs: 8 participants) were modeled using a standard GLM. Ten explanatory variables (EVs) were set up: HH-First, HH-Second, HL-First, HL-Second, LH-First, LH-Second, LL-First, LL-Second, High-Singleton, Low-Singleton. Regressors were unit-height boxcar functions that modeled the appearance of image (2 s duration) or the response / outcome (2 s duration), and were convolved with a double-gamma canonical hemodynamic response function. A second-level, fixed-effect analysis was then used to combine across four learning runs within each subject for the learning phase and up to three passive viewing runs. Finally, a third-level mixed-effects analysis was used to combine

participants' data. Third-level results were cluster-corrected for multiple comparisons using Randomise, FSL's nonparametric permutation testing tool (Jenkinson et al., 2012), with 5000 permutations and threshold free cluster enhancement (TFCE).

Our primary interest was uncovering any effect uniquely driven by the High-First images (H1) compared to the Low-First (L1) images, to uncover activity associated with attention-guided or prioritized processing coinciding with reward and order. We also ran contrasts to investigate any differences between high/low-value images that appeared with or without statistical structure (e.g., H1 or H2 > High Singleton (Hsin); L1 or L2 > Low Singleton (Lsin), and vice versa). This approach allowed us to explore the potential for reward to influence statistically structured or unstructured images (i.e., pairs vs singletons), as the additional associative information bound to structured images (or lack thereof for singletons) may predict learning based on their learned status as a high or low reward image. For the passive viewing phase, we focused on whether there is any relationship between reward contingencies and serial position even when the risky choice task was removed. If so, it would suggest that reward-associated structured or unstructured images continue to be represented uniquely outside of reward-related contexts.

RESULTS

Behavior data

We analyzed participants' choices (i.e., yes or no) for the risky choice task for each time presentation (1st to 6th) collapsed over runs. As shown in Fig 6, a two-way repeated measures ANOVA (value of image x number of presentation) on risky choice proportion (i.e., choosing yes) showed a significant main effect of value, $F(1, 28) = 51.21$, $p < .001$, $\eta p^2 = .647$, and a trend (but not significant) of main effect of the number of presentation, $F(1, 28) = 2.22$, p

= .055, ηp^2 = .074. A significant interaction between value of image (high or low) and the number of presentation (1st to 6th) was found, F (5, 140) = 29.46, p < .001, ηp^2 =.513. Proportion of making a risky choice was equally high for both high-value and low-value images at the first presentation, but across the second to sixth presentation, the proportion of making a risky choice on high-value images gradually increased, and the opposite was observed with low-value images.

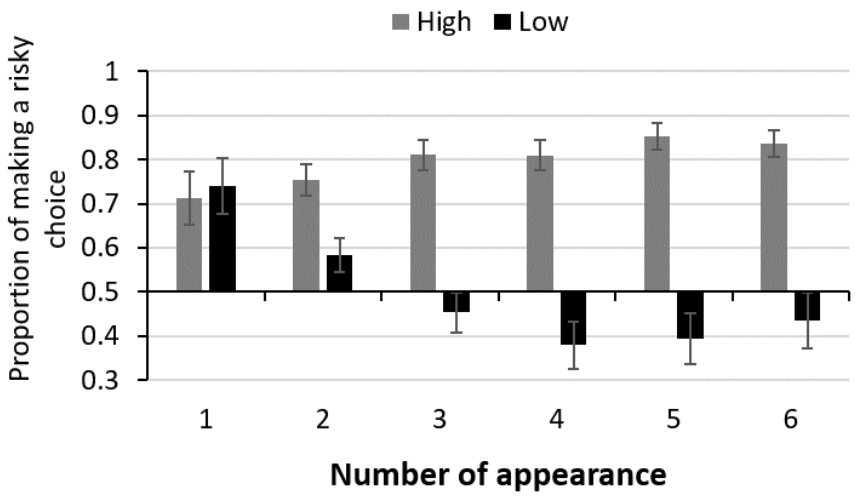

Figure 6. Proportion of making a risky choice by the number of presentations, split by value.

Participants completed two memory tasks after scanning: the recognition test and the reward memory test. For the recognition test, a one-sample t-test of recognition accuracy against chance (50%) yielded significant learning for all pair conditions; High-High: t(28)=3.3, p=.002, d=.62, High-Low: t(28)=3.26, p=.002, d=.6; Low-High: t(28)=3.1, p=.004, d=.57; Low-Low: t(28)=2.46, p=.02, d=.45 (Fig. 7). A 2 (value of first image, high or low) x 2 (value of second image, high or low) repeated measures ANOVA did not show any significant main effects nor an interaction (all p>.5). In addition, a 2 (value of first image, high or low) x 2 (value of second image, high or low) repeated measures ANOVA did not reveal any main effects nor an

interaction of foil type (i.e., foil pairs of High-High, High-Low, Low-High, and Low-Low conditions) (F <1). Although the behavioral results of Experiment 1 did not completely replicate, this is likely due to design differences. In the final passive viewing task during scanning, pairs of all reward variations were presented in pairwise order; thus, the overall recognition rate may be increased across the board and eliminate the differences between reward variations. This will be discussed further in the general discussion section.

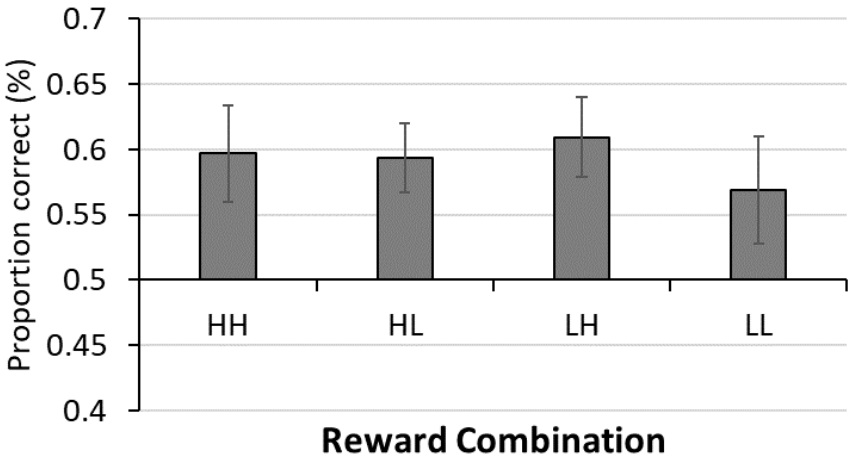

Figure 7. Accuracy at choosing target pairs over foil pair in four reward variations.

In the last reward memory phase, the mean proportion correct was 0.79 (SD: 0.13, t(28)=11.75, p < .001, d=2.18 ; one-sample t-test against chance, 50%). When we divided the results into the image type (the first, second images for pairs and singletons) and the reward type (high and low images), a repeated measures ANOVA did not show any significant main effect nor interaction (all p>.2; Fig. 8).

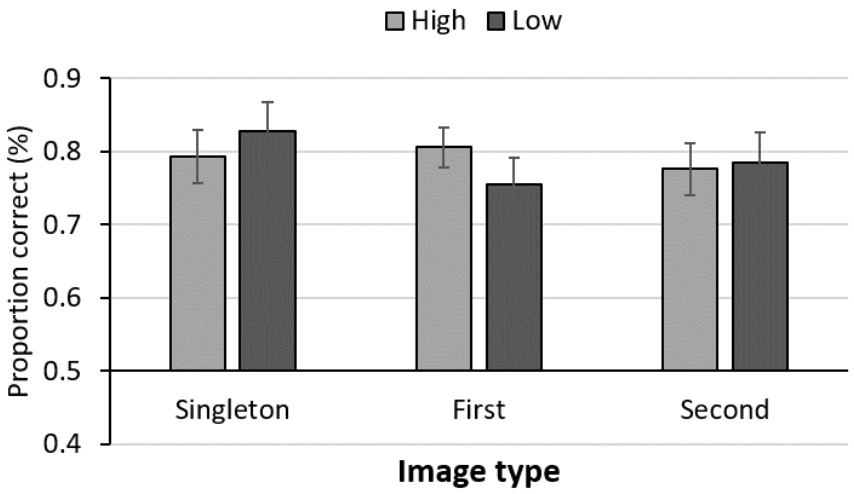

Figure 8. Accuracy at choosing reward value (high or low).

fMRI data

Learning phase

There were two main approaches we chose to explore the potential interaction of reward and VSL. First, based on our findings in Experiment 1, we were interested in uncovering any differences in neural responses between the high-value first image (H1) and the low-value first image (L1) in pairs, to examine whether differences would be consistent with differences in attentional engagement. Secondly, we were interested in contrasting any such observations with differences that might arise in response to high-value second images (H2) vs. low-value second images (L2), and high-value singletons (Hsin) vs. low-value singletons (Lsin), to ask whether structure modulated this response.

The contrast of the high-value first images versus the low-value first images (i.e., H1 > L1) yielded significant clusters in inferior frontal gyrus (IFG), middle temporal gyrus, superior temporal gyrus, parahippocampal gyrus, temporal fusiform cortex, hippocampus, amygdala, thalamus, orbitofrontal cortex (OFC), lateral occipital cortex (LOC) (all bilaterally) as well as

right putamen, left anterior cingulate gyrus (ACC), and left paracingulate gyrus (Table 1 and Figure 9). To examine whether these results were driven solely by the value difference (i.e., high vs. low), we contrasted the activity provoked by the high-value second images with that in response to the low-value second images (i.e., H2>L2 and L2>H2) but no significant difference was observed. There was also no significant difference between high-value singletons and low-value singletons (i.e., Hsin>Lsin and Lsin>Hsin). Additionally, a statistical comparison of the interactions between 1) (H1-Hsin) and (L1-Lsin), and 2) (H1-L1) and (H2-L2) were measured. The contrast of (H1-Hsin) > (L1-Lsin) yielded greater activation in the postcentral gyrus, precentral gyrus, middle temporal gyrus, superior temporal gyrus, hippocampus, amygdala, and other regions (Table 1 and Figure 10). The lack of any observable difference between high-value singleton and low-value singleton, and the significant interaction in many regions, supports the conclusion that H1 > L1 outcomes are not driven solely by the value difference, but rather an interaction between statistical regularity and value differences. With the contrasts of (H1-L1) and (H2-L2) and (H2 – Hsin) and (L2 – Lsin), no significant clusters were observed.

Considering these activations in conjunction with our results from Experiment 1, these results suggest an interaction of value processing and statistical regularity, such that high-value first images (i.e., predictive images) in particular provoke deeper processing and greater attentional engagement than low-value predictive images. The greater activation in the IFG, ACC, and LOC support our hypothesis that attention plays an important role in enhanced processing of the high-value first images.

Table 1. Results of contrast with H1>L1 and (H1-Hsin) > (L1-Lsin). In this and all other tables, clusters with five or fewer voxels were not reported.

Anatomical Label	**Hemisphere**	**Cluster size (voxel)**	**P value (TFCE)**	**Peak MNI, mm**
1. High value first image (H1) > Low value first image (L1)				
Middle Temporal Gyrus, Superior Temporal Gyrus, Parahippocampal Gyrus, Temporal Fusiform Cortex, Hippocampus, Amygdala, Thalamus, Orbitofrontal Cortex	Left	9220	0.009	-54, -28, -4
Parahippocampal Gyrus, Temporal Fusiform Cortex, Hippocampus, Amygdala, Thalamus, Accumbens	Right	2360	0.014	26, -34, -16
Middle Temporal Gyrus, Superior Temporal Gyrus, Supramarginal Gyrus, Planum Temporale, Parietal Operculum	Right	890	0.034	72, -32, 2
Orbitofrontal Cortex, Inferior Frontal Gyrus	Right	122	0.041	28, 20, -22
Frontal Pole	Left	54	0.04	-16, 50, 42
Cingulate Gyrus	Left	23	0.041	-2, -12, 34
Occipital Fusiform Gyrus	Right	13	0.046	30, -68, 0
Lateral Occipital Cortex	Left	8	0.048	-28, -88, 20
Paracingulate Gyrus	Left	6	0.05	-4, 46, 8
2. (H1 - High Singleton) > (L1 - Low Singleton)				
Postcentral Gyrus, Precentral Gyrus	Right	203	0.033	26, -26, 68
Planum Temporale, Superior Temporal Gyrus	Right	145	0.041	50, -32, 14

Middle Temporal Gyrus, Superior Temporal Gyrus	Left	35	0.042	-66, -28, 0
Hippocampus, Amygdala	Left	22	0.037	-20, -10, -20
Planum Temporale	Right	10	0.048	64, -10, 2

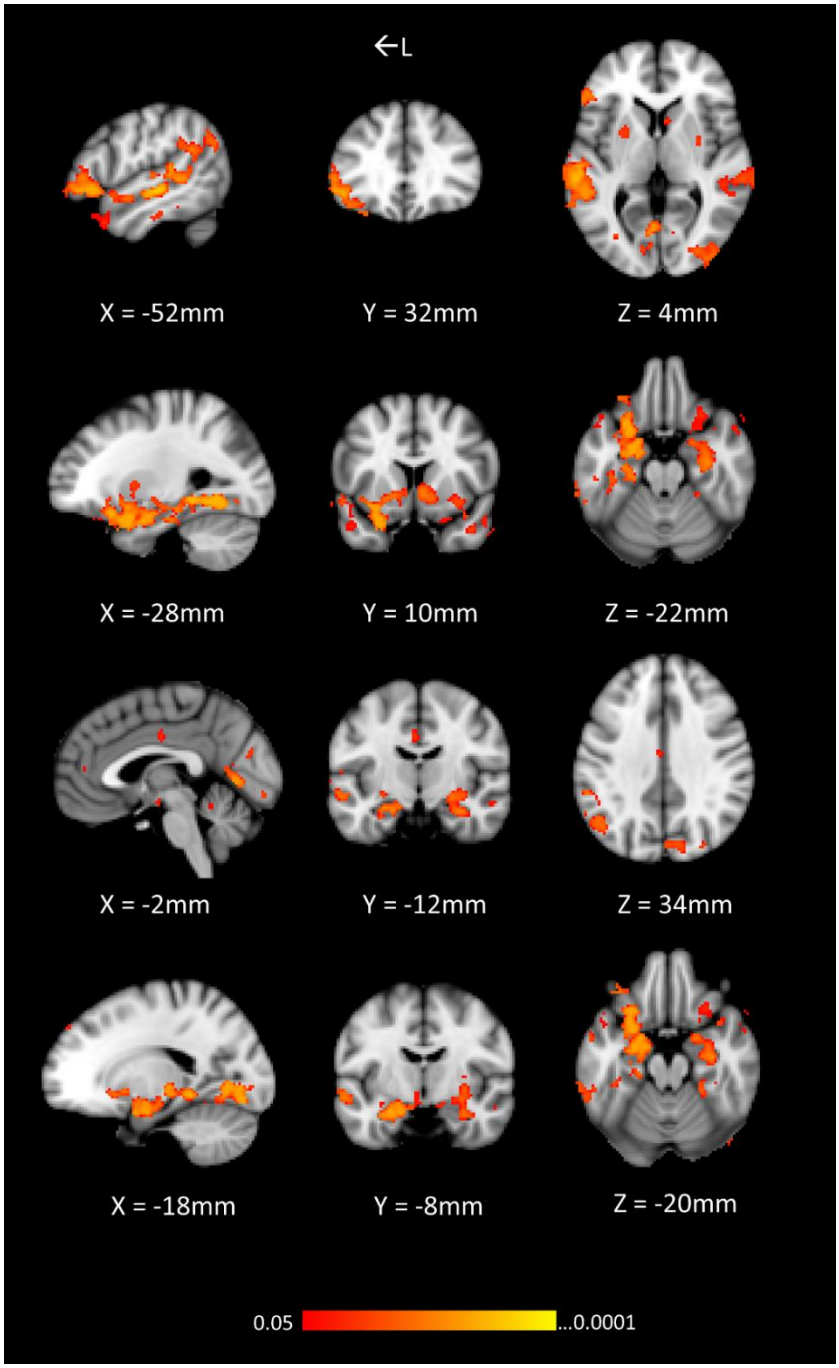

Figure 9. The H1 > L1 contrast yielded clusters that included IFG, middle temporal gyrus, superior temporal gyrus, parahippocampal gyrus, temporal fusiform cortex, hippocampus, amygdala, thalamus, OFC, LOC (all bilaterally) as well as right putamen, left ACC, and left paracingulate gyrus. From top to bottom, coordinates are centered on left IFG, left OFC, left ACC, and left hippocampus. In this and all other figures, coordinates are in MNI standard space.

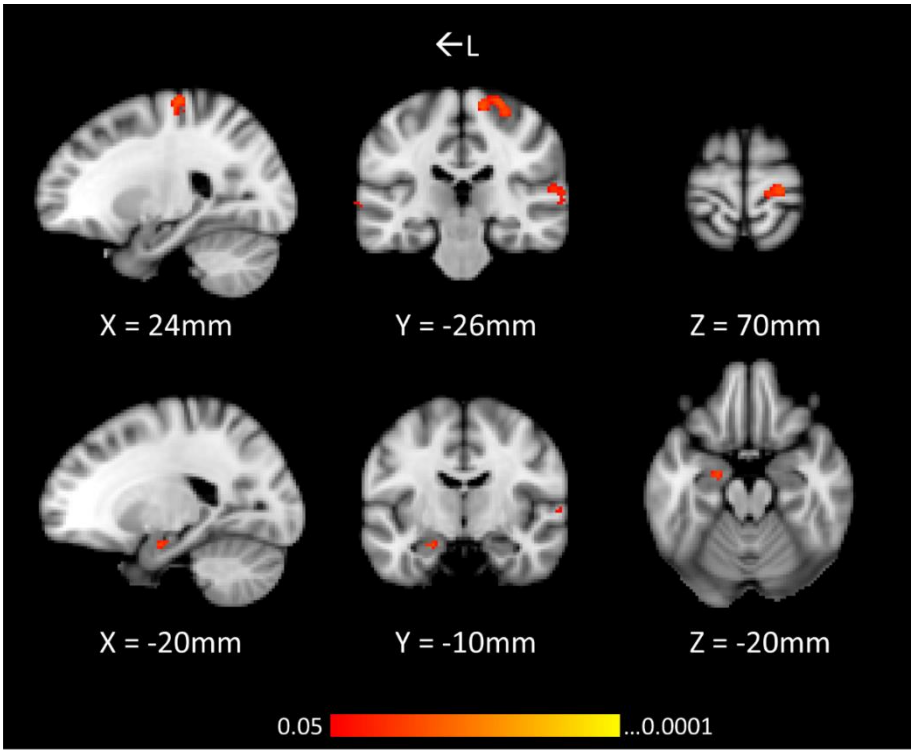

Figure 10. The (H1-Hsin) > (L1-Lsin) contrast yielded significant clusters in the postcentral gyrus, precentral gyrus, middle temporal gyrus, superior temporal gyrus, hippocampus, amygdala. Top row coordinates are centered on right precentral gyrus, and bottom row coordinates are centered on left hippocampus.

Following up on these results, we examined how statistical regularities modulate responses, keeping value constant. We examined four contrasts: 1) H1 vs. Hsin, 2) H2 vs. Hsin, 3) L1 vs. Lsin, and 4) L2 vs. Lsin. We observed significant clusters for 1) Lsin > L1 and 2) Lsin> L2. The contrast of Lsin > L1 showed greater activation in middle temporal gyrus, hippocampus, amygdala, putamen, LOC, and other regions (Table 2 and Figure 11A), and the contrast of Lsin > L2 resulted in clusters in similar areas (Table 2 and Figure 11B). Comparisons between high-value paired images and high-value singletons did not yield any significant differences.

Table 2. Results of contrast with Lsin>L1 and Lsin>L2

Anatomical Label	**Hemisphere**	**Cluster size (voxel)**	**P value (TFCE)**	**Peak MNI, mm**
1. Low-value singleton (Lsin) > Low-value first image (L1)				
Middle Temporal Gyrus, Superior Temporal Gyrus, Caudate, Parahippocampal Gyrus, Temporal Fusiform Cortex, Hippocampus, Amygdala, Cingulate Gyrus, Thalamus, Putamen, Postcentral Gyrus, Planum Temporale, Precentral Gyrus, Lateral Occipital Cortex, Ventricle, Right cerebellum, Accumbens	Both (unless stated specifically)	44022	0.009	10, -34, 58
Precuneous Cortex	Left	146	0.036	-22, -48, 26
Lingual Gyrus	Left	32	0.046	-6, -70, 2
Cerebellum	Left	32	0.046	-4, -46, -12
Intracalcarine Cortex ,Precuneous Cortex	Right	8	0.05	6, -66, 12
2. Low-value singleton (Lsin) > Low-value second image (L2)				
Middle Temporal Gyrus, Superior Temporal Gyrus, Precuneous Cortex, Supramarginal Gyrus	Right	1903	0.02	22, -52, 30
Caudate, Pallidum, Putamen	Right	670	0.021	16, 4, 26
Ventricle	Both	145	0.041	0, 4, 6
Precuneous Cortex	Left	63	0.038	-24, -52, 26

Orbitofrontal Cortex, Frontal Pole, Caudate, Putamen	Right	52	0.042	22, 32, -6
Temporal Occipital Fusiform Cortex	Left	48	0.048	-40, -44, -10
Planum Polare	Right	24	0.044	42, -24, -4
Inferior Temporal Gyrus, Lateral Occipital Cortex	Left	6	0.049	-48, -66, -14

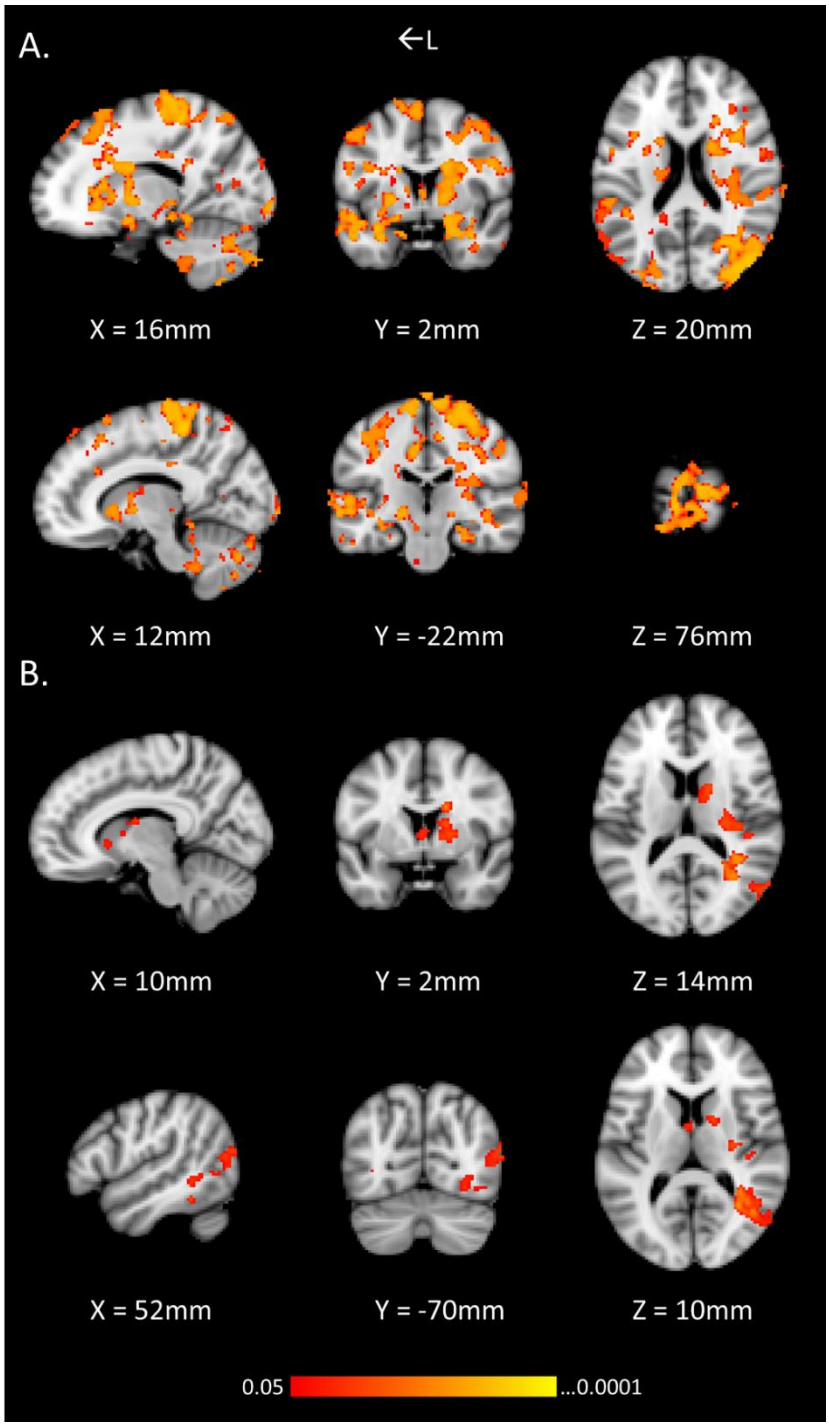

Figure 11. A) The contrast of Lsin > L1 showed greater activation in middle temporal gyrus, hippocampus, amygdala, putamen, and LOC. B) The contrast of Lsin > L2 also showed significant activations in middle temporal gyrus, hippocampus, inferior temporal gyrus, amygdala, putamen, and LOC. From top to bottom row, coordinates are centered on right caudate, right precentral, right caudate, and right LOC.

We were not able to find any significant differences with contrasts scrutinizing the first images of paired images > singletons (see Turk-Browne et al., 2010). With our design, however, paired images contained not only statistical structure but also reward information, and the interaction between these two variables may drive a different pattern of results. Rather, we found that low-value singletons showed greater activity than low-value predictive (L1) images in areas recognized for playing a role in processing reward information (e.g., caudate, putamen, hippocampus). These results suggest that our (H1-L1)>(Hsin-Lsin) interaction may have been driven predominantly by differences in the way that L1 images are processed compared to low-value images that are non-predictive.

Passive viewing phase

During the passive viewing phase, participants were not required to perform any task other than to focus on each image as it goes by. We were interested in seeing whether any reward/structure related findings from the risky choice task phase would extend into other contexts (i.e., a context where participants are no longer making a choice or actively earning reward). However, we were unable to find similar patterns of activity with contrasts we ran with the risky choice task. We suspect that any failure to observe patterns of activity similar to that found for the learning phase is possibly due to a lack of power from having time to collect data from a single of the passive viewing task for most participants.

DISCUSSION

In Experiment 2, we measured brain responses to visual images that were associated with both varying levels of reward and statistical contingencies. We found that the high-value first image (i.e., H1) led to greater activity in areas including IFG, left ACC, LOC, fusiform gyrus,

orbitofrontal cortex (OFC), accumbens, precuneous cortex, parahippocampal gyrus, middle temporal gyrus, amygdala, hippocampus, and putamen as compared to the low-value first image (i.e., L1). These findings suggest that H1, in comparison to L1, led to greater attentional engagement (Beck & Vickery, 2019; Murray & Wojciulik, 2004b; Stokes et al., 2009), and may enhance associative learning thusly. The contrasts of (H1-Hsin) > (L1-Lsin) yielded greater activations in the precentral gyrus, middle temporal gyrus, hippocampus, and amygdala, which supports the possibility that the differences between the high-value first image and the low-value first image are not driven solely by the value difference, but by an interaction of predictiveness and value. As no difference was found between the high-value singletons and the low-value singletons, the greater activation in precentral gyrus in this interaction also supports our hypothesis that the first image of the pair that was associated with high value received the attentional priority in comparison to that was associated with the low value (Fockert et al., 2004).

For contrasts comparing first images and singletons (e.g., H1>Hsin and L1>Lsin), we were not able to replicate the findings of Turk-Browne et al. (2010). We speculate that embedded reward information possibly altered learning in such a way that made it unique versus when VSL occurs in the absence of reward. Rather, we found evidence that when an image contains both statistical structure and reward information, the reward learning may be predicted by the type of statistical structure it is associated with. Specifically, we found that low-value predictive images (i.e., L1) provoked less activity than non-predictive low-value singletons. In contrast, there was no difference in high-value comparisons between paired images and singletons. This suggests that the predictive nature of a stimulus may specifically down-regulate responses to low-value images, and thus that attention was less guided/prioritized to the low-value first image than the high-value first image.

GENERAL DISCUSSION

Across two experiments, we provided behavioral and neural evidence that reward alters visual statistical learning. In Experiment 1, better recognition of pairs when the first image of a pair was associated with high-value was observed, and this effect was especially pronounced for High-High pairs. Neural evidence supports this finding, such that when the first image of a pair was associated with high-value, in comparison to the first image being associated with low-value, greater BOLD response was observed in brain areas that have been known to play an essential role in associative learning (e.g., hippocampus, precuneous cortex, parahippocampal gyrus; Turk-Browne et al., 2010), and value processing (e.g., OFC, accumbens, and caudate; Baliki et al., 2013; Kringelbach & Rolls, 2004). These findings are not driven solely by the value difference, but rather the interaction between statistically structured information and reward – with additional analyses, we showed that this value difference was specific to predictive items, suggesting that predictiveness and value interact in determining neural responses to images.

Attentional processing may be involved in this interaction, as we found greater activations in LOC and frontal and parietal areas, such as inferior frontal gyrus, precentral gyrus, and anterior cingulate gyrus, all regions whose activity is known to scale with attentional processing (Beck & Vickery, 2019; Corbetta & Shulman, 2002; Fockert et al., 2004), when the H1 was compared to L1. We did not observe similar differences with comparisons of H2 vs. L2, or Hsin vs. Lsin, which implies that not all high-value images led to attentional prioritization compared to low-value images. Rather, a combination of predictiveness and reward value was crucial in provoking this response. In VSL, the predictive image plays an important role, provoking anticipatory responses (Turk-Browne et al., 2010). Our findings suggest that VSL occurs differentially as a function of the magnitude of reward associated with the first image. In

other words, when different amounts of reward are embedded in visual regularity, the rewards may interact with VSL in a way that it only impacts on the first position of the structured sequence.

In Experiment 2, the processing of reward information elicited different patterns with behavior and neural approaches. Our behavioral results of learning of reward associations (i.e., the reward memory test) showed no difference in reward memory as a function of the structured information, which means that predictive structure (i.e., paired vs. singleton) did not impact the recognition of reward information. However, our neural evidence reveals that reward-related responses were differentiated based on which structural position the reward was embedded in. Differences in the binding of reward based on stimulus-stimulus predictiveness may be too subtle for behavioral methods to uncover, although differences in the design of our experiments may also play a role in discrepant findings between behavioral and neural methods.

Different results of the recognition phase between Experiment 1 and 2 may also be derived from the design of our experiments. We found above chance levels of learning in all conditions (i.e., High-High, High-Low, Low-High, and Low-Low pairs) in Experiment 2, unlike Experiment 1 that showed above chance learning only for High-High and High-Low pairs. We speculate that this difference was due to the passive viewing task, where all pairs appeared one to three times across all participants. In this phase, participants were not required to make any choice, which means they did not have to process information related to reward variation. Hence, there is a possibility of the difference in recognition rates between reward variations might be washed out since participants were exposed to all pairs in a reward-absent environment. Prior work provided evidence that regularity information can be learned with a small amount of exposure (Turk-Browne & Scholl, 2009), and we noticed that the overall proportion correct

responses were higher in Experiment 2 than Experiment 1. Prior work and the trend of our finding support our speculation that the overall recognition was boosted because of the addition of the passive viewing task. In addition, in Experiment 2, each run introduced a new set of images, and this aspect of the design may have introduced differences in VSL, as well.

To our knowledge, this is the first work to provide evidence of behavioral and neural responses being modulated by the interaction of reward and VSL. As mentioned above, Rogers et al. (2016) first explored the interaction between reward and VSL, but the reward variations (i.e., no-, low-, or high-reward) did not affect the learning of regularities. In our work, by using a risky choice task, we enhanced participants' engagement to the task and value, and were able to observe an effect of reward on VSL. This implies that robust engagement with value information may be necessary to induce interactions with the learning of visual regularities. In conjunction with other recent results highlighting the importance of task during exposure shaping VSL (Vickery et al., 2018), the current study highlights the need to carefully consider context during exposure to regularities, and how those contexts shape incidental learning.

With respect to the passive viewing task, we were unable to observe a similar pattern of activity as that found in the risky choice task phase. We suspect that this failure is possibly due to a lack of power, due to our only having time to collect data from a single run of the passive viewing task for most participants. Another possible explanation for lack of such a finding is that the effect of reward in VSL may only arise within the context of tasks that draw attention to value, like our risky choice task. Therefore, simply viewing the sequence of images may not yield that same neural responses as actively making a risky choice on each image. Further studies of how the interaction between reward and VSL may affect the later representation of memory even outside of a reward-related context may be needed.

The present study provides evidence that VSL is altered by reward. When a high reward is embedded in the first location of a statistically structured pair, it aids learning: a result we found support for in neural evidence. Several brain areas that reflect attentional capture, reward processing, associative learning, and the intermixed effect among them support the notion that reward contingencies affect VSL. These findings highlight the fact that reward may play a role similar to selective attention in VSL, such that the more the image can guide attentional resources, the better it can convey the reward information, and ultimately, facilitate visual statistical learning.

References

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. E. (2006). Reward-motivated learning: mesolimbic activation precedes memory formation. *Neuron*, *50*(3), 507–517.

Anderson, B. (2013). A value-driven mechanism of attentional selection. *Journal of Vision*, *13*(2013), 1–16. https://doi.org/10.1167/13.3.7.doi

Anderson, B A, Laurent, P. A., & Yantis, S. (2013). Reward predictions bias attentional selection. *Front Hum Neurosci*, *7*(June), 262. https://doi.org/10.3389/fnhum.2013.00262

Anderson, Brian A. (2017). Reward processing in the value-driven attention network: Reward signals tracking cue identity and location. *Social Cognitive and Affective Neuroscience*, *12*(3), 461–467. https://doi.org/10.1093/scan/nsw141

Baker, C. I., Olson, C. R., & Behrmann, M. (2004). Visual Statistical Learning Role of Attention and Perceptual Grouping in Visual Statistical Learning. *Psychological Science*, *15*(7), 1–8.

Baliki, M. N., Mansour, A., Baria, A. T., Huang, L., Berger, S. E., Fields, H. L., & Apkarian, A. V. (2013). Parceling Human Accumbens into Putative Core and Shell Dissociates Encoding of Values for Reward and Pain. *Journal of Neuroscience*, *33*(41), 16383–16393. https://doi.org/10.1523/jneurosci.1731-13.2013

Beck, V. M., & Vickery, T. J. (2019). Multiple states in visual working memory: Evidence from oculomotor capture by memory-matching distractors. *Psychonomic Bulletin & Review*, *26*(4), 1340–1346.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436.

https://doi.org/10.1163/156856897X00357

Clark, L., Lawrence, A. J., Astley-Jones, F., & Gray, N. (2009). Gambling near-misses enhance motivation to gamble and recruit win-related brain circuitry. *Neuron*, *61*(3), 481–490.

Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*(3), 201–215.

Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, *29*(3), 162–173.

Daw, N. D., & Doya, K. (2006). The computational neurobiology of learning and reward. *Current Opinion in Neurobiology*, *16*(2), 199–204. https://doi.org/10.1016/j.conb.2006.03.006

Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, *12*(6), 499–504. https://doi.org/10.1111/1467-9280.00392

Fiser, J., & Aslin, R. N. (2002). Statistical learning of higher-order temporal structure from visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(3), 458–467. https://doi.org/10.1037//0278-7393.28.3.458

Fockert, J. de, Rees, G., Frith, C., & Lavie, N. (2004). Neural correlates of attentional capture in visual search. *Journal of Cognitive Neuroscience*, *16*(5), 751–759.

Grill-Spector, K., Kourtzi, Z., & Kanwisher, N. (2001). The lateral occipital complex and its role in object recognition. *Vision Research*, *41*(10–11), 1409–1422.

Haber, S. N., & Knutson, B. (2010). The reward circuit: linking primate anatomy and human

imaging. *Neuropsychopharmacology*, *35*(1), 4–26.

James, T. W., Culham, J., Humphrey, G. K., Milner, A. D., & Goodale, M. A. (2003). Ventral occipital lesions impair object recognition but not object-directed grasping: an fMRI study. *Brain*, *126*(11), 2463–2475.

Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., & Smith, S. M. (2012). Fsl. *Neuroimage*, *62*(2), 782–790.

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception*, *36*(14), 1.

Kringelbach, M. L., & Rolls, E. T. (2004). The functional neuroanatomy of the human orbitofrontal cortex: Evidence from neuroimaging and neuropsychology. *Progress in Neurobiology*, *72*(5), 341–372. https://doi.org/10.1016/j.pneurobio.2004.03.006

Miendlarzewska, E. A., Bavelier, D., & Schwartz, S. (2016). Influence of reward motivation on human declarative memory. *Neuroscience & Biobehavioral Reviews*, *61*, 156–176.

Murray, S. O., & Wojciulik, E. (2004a). Attention increases neural selectivity in the human lateral occipital complex. *Nature Neuroscience*, *7*(1), 70–74. https://doi.org/10.1038/nn1161

Murray, S. O., & Wojciulik, E. (2004b). Attention increases neural selectivity in the human lateral occipital complex. *Nature Neuroscience*, *7*(1), 70–74.

Musz, E., Weber, M. J., & Thompson-Schill, S. L. (2015). Visual statistical learning is not reliably modulated by selective attention to isolated events. *Attention, Perception, & Psychophysics*, *77*(1), 78–96.

Rogers, L. L., Friedman, K. G., & Vickery, T. J. (2016). No Apparent Influence of Reward upon

Visual Statistical Learning. *Frontiers in Psychology*, *7*, 1687.

Smith, S. M. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, *17*(3), 143–155.

Stokes, M., Thompson, R., Nobre, A. C., & Duncan, J. (2009). Shape-specific preparatory activity mediates attention to targets in human visual cortex. *Proceedings of the National Academy of Sciences*, *106*(46), 19569–19574.

Theeuwes, J., & Belopolsky, A. V. (2012). Reward grabs the eye: Oculomotor capture by rewarding stimuli. *Vision Research*, *74*, 80–85. https://doi.org/10.1016/j.visres.2012.07.024

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

Turk-Browne, N. B., & Scholl, B. J. (2009). Flexible visual statistical learning: transfer across space and time. *Journal of Experimental Psychology. Human Perception and Performance*, *35*(1), 195–202. https://doi.org/10.1037/0096-1523.35.1.195

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. *Dx.Doi.Org*, *21*(10), 1934–1945. https://doi.org/10.1162/jocn.2009.21131

Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit Perceptual Anticipation Triggered by Statistical Learning. *Journal of Neuroscience*, *30*(33), 11177–11187. https://doi.org/10.1523/JNEUROSCI.0858-10.2010

Tversky, A., & Kahneman, D. (1979). Prospect theory: An analysis of decision under risk.

Econometrica, *47*(2), 263–291.

Vickery, T. J., Park, S. H., Gupta, J., & Berryhill, M. E. (2018). Tasks determine what is learned in visual statistical learning. *Psychonomic Bulletin & Review*, *25*(5), 1847–1854.

Vuilleumier, P., Schwartz, S., Duhoux, S., Dolan, R. J., & Driver, J. (2005). Selective attention modulates neural substrates of repetition priming and "implicit" visual memory: suppressions and enhancements revealed by FMRI. *Journal of Cognitive Neuroscience*, *17*(8), 1245–1260.

Woolgar, A., Williams, M. A., & Rich, A. N. (2015). Attention enhances multi-voxel representation of novel objects in frontal, parietal and visual cortices. *Neuroimage*, *109*, 429–437.