


This article was downloaded by: [Princeton University]

On: 12 July 2015, At: 16:56

Publisher: Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: 5 Howick Place, London, SW1P 1WG



[Click for updates](#)

Cognition and Emotion

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/pcem20>

Human preferences are biased towards associative information

Sabrina Trapp^a, Amitai Shenhav^b, Sebastian Bitzer^a & Moshe Bar^c

^a Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany

^b Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA

^c Gonda Center for Brain Research, Bar-Ilan University, Ramat Gan, Israel

Published online: 10 Oct 2014.

To cite this article: Sabrina Trapp, Amitai Shenhav, Sebastian Bitzer & Moshe Bar (2015) Human preferences are biased towards associative information, *Cognition and Emotion*, 29:6, 1054-1068, DOI: [10.1080/02699931.2014.966064](https://doi.org/10.1080/02699931.2014.966064)

To link to this article: <http://dx.doi.org/10.1080/02699931.2014.966064>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden. Terms & Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

Human preferences are biased towards associative information

Sabrina Trapp¹, Amitai Shenhav², Sebastian Bitzer¹, and Moshe Bar³

¹Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany

²Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA

³Gonda Center for Brain Research, Bar-Ilan University, Ramat Gan, Israel

There is ample evidence that the brain generates predictions that help interpret sensory input. To build such predictions the brain capitalizes upon learned statistical regularities and associations (e.g., “A” is followed by “B”; “C” appears together with “D”). The centrality of predictions to mental activities gave rise to the hypothesis that associative information with predictive value is perceived as intrinsically valuable. Such value would ensure that this information is proactively searched for, thereby promoting certainty and stability in our environment. We therefore tested here whether, all else being equal, participants would prefer stimuli that contained more rather than less associative information. In Experiments 1 and 2 we used novel, meaningless visual shapes and showed that participants preferred associative shapes over shapes that had not been associated with other shapes during training. In Experiment 3 we used pictures of real-world objects and again demonstrated a preference for stimuli that elicit stronger associations. These results support our proposal that predictive information is affectively tagged, and enhance our understanding of the formation of everyday preferences.

Keywords: Prediction; Preference; Perception; Affect; Statistical learning; Novelty.

Our perceptual system is faced with a huge computational challenge: it has to identify features and objects based on noisy and often ambiguous sensory input. So how is perception accomplished so quickly and efficiently? In cognitive neuroscience, a growing consensus is that perception is supported by previous experiences and prior expectations, which might help resolve ambiguity

and fill-in absent information in a top-down manner (e.g., Bar et al., 2006; Summerfield & Koechlin, 2008). Information that is expected offers considerable advantages for the nervous system: it enables faster, more efficient processing (Chun & Jiang, 1998; Turk-Browne & Scholl, 2009), reduces neural responses (Alink, Schwiedrzik, Kohler, Singer, & Muckli, 2010; Todorovic,

Correspondence should be addressed to: Dr. Sabrina Trapp, Max Planck Institute for Human Cognitive and Brain Sciences, Stephanstrasse 1a, D-04103 Leipzig, Germany. E-mail: strapp@cbs.mpg.de

We thank E. Aminoff, J. Boshyan and K. Devaney for assistance with Experiment 3. This study was supported by a Max Planck postdoctoral fellowship (Dr med. Anneliese & DSc Dieter Pontius Foundation) (ST), C.V. Starr Foundation postdoctoral fellowship (AS) and Israeli Center of Research Excellence in Cognition [grant number 51/11] (MB).

van Ede, Maris, & de Lange, 2011) and sharpens representations (Kok, Jehee, & de Lange, 2012). Given these advantages, there might be a link between (neutral) information that can be predicted and positive reward. Such an *affective tag* could encourage a proactive search for predictable information, thereby warranting a higher degree of certainty and stability of an organisms' ecological niche. There is some preliminary support for this idea. In a study by Ogawa and Watanabe (2011), participants performed a contextual cueing task in which they searched for visual targets, the locations of which were either predicted or not predicted by the positioning of surrounding distractors. The speed with which participants located the targets increased across trials more rapidly for predictive displays than for non-predictive displays. Participants were subsequently asked to rate the "goodness" of visual displays. The authors found that participants gave higher ratings to predictive displays compared with both non-predictive and novel displays. However, a potential drawback of this study is that predictability was confounded with task difficulty—whenever information was predictive, it also facilitated task performance. Furthermore, "goodness" does not directly map to preference, but rather could be more reflective of how helpful participants found the display to be. Thus, it cannot be unequivocally determined whether participants' ratings reflected a preference for predictive displays or for information that involved task facilitation.

In our natural environment, objects, events and actions never occur in isolation, but rather with other entities that tend to share the same context. For example, there is a high probability that a refrigerator co-occurs with a stove, a microwave and dishes. The brain can exploit such context information (for a review see Bar, 2004), as predictable object co-occurrences might be utilized in constraining expectations to only those objects that have been regularly associated with each other in the past. Consequently, it has been proposed that associations provide the building blocks for predictions (Bar, 2007). There is indirect support that associative information, as a proxy for predictability, is itself hedonically marked. In a functional magnetic resonance imaging (fMRI) study it was reported

that visual preference judgements elicit activation in the parahippocampal cortex (Yue, Vessel, & Biederman, 2007). This area has repeatedly been implicated in the analysis of contextual associations (for a recent review see Aminoff, Kveraga, & Bar, 2013). Although it cannot be concluded from this study, its results support the hypothesis that preference towards a stimulus is related to associative processing. Another indirect link comes from studies that demonstrated that mood improves after processing information with broader associative span (Brunye et al., 2013; Mason & Bar, 2012).

Here we test directly the hypothesis that associative information, as a precursor of predictive processing, possesses a positive value. To this end, we assessed human preference judgements in three experiments. In Experiments 1 and 2, participants were trained on novel, meaningless visual shapes that were either regularly associated with other such shapes, or not. In Experiment 3, we assessed subjective preferences for images of real-world objects that had been independently normed for varying levels of associativity. We predicted that even without any apparent benefits, participants would prefer associative information over information that had no or only weak associations.

EXPERIMENT 1

Method

Participants

Fourteen participants (mean age = 24.7 years, SD = 2.4; 8 females) who were recruited from the Max Planck Institute for Human Cognitive and Brain Science database gave written informed consent and were financially compensated for their participation.

Stimuli

We used 64 arbitrary, novel visual shapes, taken from a previous study of visual associative processing (Aminoff, Gronau, & Bar, 2007). The shapes were devoid of any semantic meaning or any apparent affective value. Of these, 32 shapes were assigned to the associative condition and the other 32 were assigned to the non-associative

condition. These two sets were counterbalanced across participants.

Procedure

Experiment 1 consisted of two parts. In the first part, participants learned associations for arbitrary visual shapes. In the second part, participants were asked to choose one of two shapes that they prefer, where one of those shapes had been associated with other shapes before (as part of fixed groupings during the learning task), and the other had only occurred in random groupings (see Figure 1).

Visual association learning task. Participants were told that they were participating in a co-occurrence detection task. Four shapes were presented, forming an imaginary square around a

fixation cross at the centre of the screen, and participants had to indicate whether those shapes had co-occurred together in previous trials (button 1) or not (button 2). The shapes were presented for 3s and the trials were separated by a 1.5s interstimulus interval. Emphasis was given on accuracy, though participants were encouraged not to take too much time to decide, because trials automatically ended 3s after stimulus onset, if no response was given. Importantly, participants did not receive any feedback about their performance to avoid tagging successful trials with a reward-like outcome that might eventually influence their preference judgement. The task lasted approximately 25 minutes.

Fifty percent of the quadruplets were presented as fixed groupings, i.e., associated quadruplets

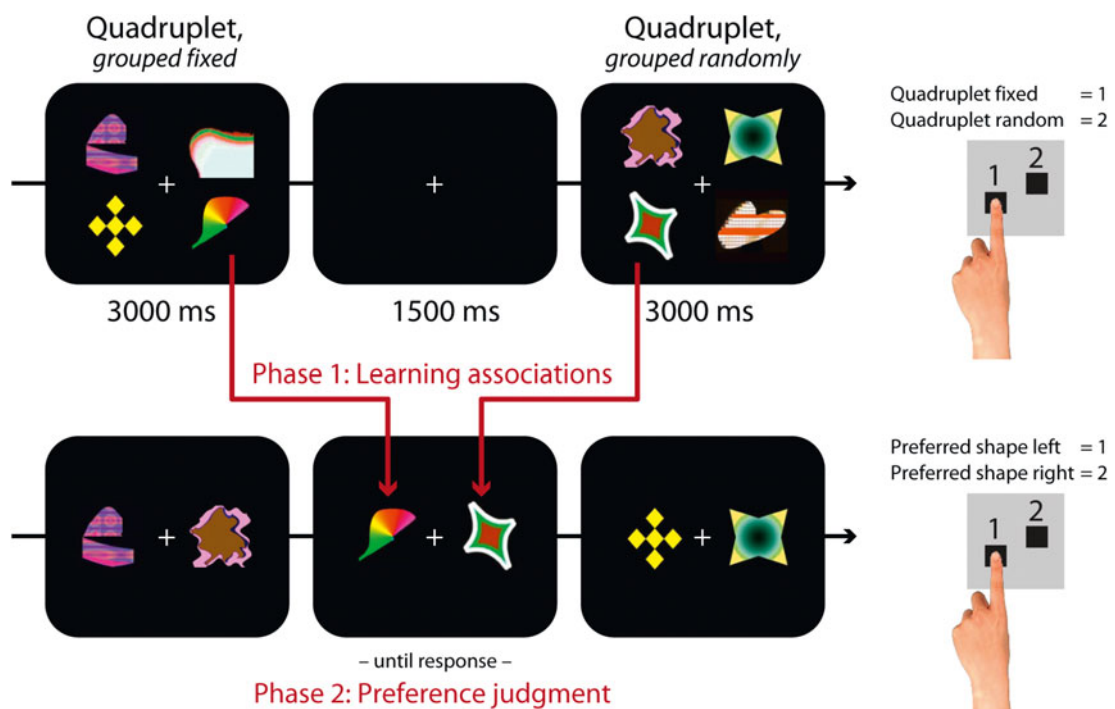


Figure 1. Experiment 1 consisted of two parts. In the first part (top), participants learned associations for arbitrary visual shapes. Fifty percent of those shapes were always grouped in the same quadruplet, and presented on the same spatial position, and 50% were randomly grouped together. The training consisted of 320 trials in total. There were 20 repetitions for each single shape in the non-associative condition, and 20 repetitions for each single shape in the associative condition. In the second part (bottom), participants were asked to choose which of two shapes they prefer, where one of those shapes had previously been associated with other shapes, and the other had only occurred in random groupings.

(four shapes of an associated quadruplet always appeared together and at the same spatial position on-screen), and the other 50% of quadruplets were randomly grouped shapes. The conditions (associative, non-associative) were randomly intermixed across the experiment. The experiment consisted of 320 trials in total. It is important to note that the repetition frequency for associated and non-associated shapes was matched, circumventing judgements based on mere familiarity of the stimuli. There were 20 repetitions for each single shape in the non-associative condition, and 20 repetitions for each single shape in an assigned quadruplet in the associative condition. Furthermore, the assignment of which shapes participated in the associated and which in the non-associated quadruplets was counterbalanced across participants.

Preference task. Immediately after the learning part, participants were given an unanticipated 2-alternative forced-choice preference task. They were told that it was unrelated to the first task, and serves to evaluate a stimulus set for future projects. Participants viewed 32 displays of shape pairs. On each trial, one shape was presented on the left, and the other shape was presented on the right side of the screen. One of the two shapes came from the associative condition, the other from the non-associative condition. The presence of associative/non-associative shapes on the right/left of the screen was randomised. Participants were instructed to press one of two buttons, indicating which shape they preferred. The pair was presented on the screen until a response was given. We encouraged participants to provide a rapid, gut-level response. Task duration was 5 minutes on average.

The experiments were presented on a 19" computer screen with the Presentation software version 14.7 (Neurobehavioral Systems, Inc).

Results

First, to ensure that associations were properly learned, we evaluated participants' success rate in

identifying repeated, associated quadruplets in the first part of the experiment. Note that the task was a trial-by-trial learning task, i.e., participants could only decide whether a quadruplet is repeated or not after at least the second trial.

One participant was excluded from analysis because the responses were always given past the response window (3s, see Procedure).

To evaluate learning success, we calculated the percentage of "associative" responses for four different time bins of the learning task (each containing 40 trials), separately for the associative and non-associative condition. The percentage of "associative" responses (and SD) for the fixed associated quadruplets for the four bins, in order, was 65.2% (12.6%), 90.4% (12.4%), 95.4% (14.5%) and 89.2% (16.9%). For non-associative quadruplets, it was 26.2% (20.2%), 30.0% (16.9%), 41.0% (28.6%) and 54.6% (26.9%).

To test our main hypothesis that participants would prefer shapes that carry more associative information, we evaluated participants' choices for shapes that had been presented in fixed vs. random quadruplets in the preference task. Across participants, the associative shape was preferred over the non-associated shape on an average of 22.5 out of the 32 choices (70.4%), a significant departure from indifference (Figure 2, left; Wilcoxon Signed-Rank, $p < .002$).

As noted above, participants displayed an increasing tendency to respond with "associative" for non-associative information towards the end of the experiment (i.e., false alarms). The reason for this bias particularly towards the end of the experiment could be attributed to proactive interference due to the inevitable repeated coupling of individual shapes or stimulus locations in the non-associative condition, based on random sampling from the 32 shapes. This might have led participants to wrongly indicate that the quadruplet was "associative". If true, this would make our results even more conservative: had participants tagged or stored this (non-associative) shape under "associative", it would decrease the likelihood that they would come to prefer the (objectively) more associative condition. Additionally, it is important




Figure 2. Results for the preference judgement task. The left bar shows the data from Experiment 1, in which participants were instructed to detect regularly associated shapes in the initial learning task. The right bar shows data from Experiment 2, in which participants were instead instructed to detect novel shape configurations during the learning task. After the learning task, subjects in both experiments were asked to choose which of two presented shapes they preferred: one of those shapes had been repeatedly associated with other shapes (in consistent groupings) in the preceding learning task, and the other had only occurred in random groupings. The bars show the overall preference for shapes that have been repeatedly associated with others. The dotted line represents choices based on indifference (50%). Error bars represent standard errors of the mean.

to note that there was never any feedback on participants' performance, and thus no information on the adequacy of participant's strategies or their correctness, something which might influence subsequent preference judgements.

We hypothesised that regularly and predictably associated shapes (i.e., where each shape of a quadruplet was always associated with the same other three shapes and always presented at the same location on-screen) would be preferred over shapes that had been coupled with other shapes at random, and in random spatial locations. Indeed, we found a strong bias towards the associative shapes. Although these findings support our hypothesis, one potential concern is that the

associations in this task are not learned incidentally. Participants are explicitly instructed to pay attention to and to learn these associative groupings. To rule out the possibility that participants preferred the shapes in the associative groupings because detecting them satisfied experimenter instructions, we conducted the same experiment again, but with a different task instruction that took the focus away from associativity. Specifically, participants were instructed to detect novel configurations of shapes. We reasoned that if we still find preferences to be biased towards associative shapes, such a bias would be independent of the task instructions.

EXPERIMENT 2

Method

Participants

Fifteen new participants (mean age = 23.6 years, SD = 2.3; 9 females) who were recruited from the Max Planck Institute for Human Cognitive and Brain Science database gave written informed consent and were financially compensated for their participation.

Stimuli

The stimuli were identical to Experiment 1.

Procedure

The procedure and task were identical to Experiment 1, except that task instruction for the first learning task was to detect novel shape configurations. Accordingly, participants had to indicate whether shape groupings were novel (button 1) or not (button 2).

Analysis

We performed a logistic regression analysis to confirm preference for associative shapes across Experiments 1 and 2 and to investigate a potential effect of task instruction which differed in the two experiments. The logistic regression predicted choice of the associative shapes in each trial with the identity of the experiment as a binary factor.

Regression coefficients for the intercept (general tendency to choose the associative shape) and the experiment factor (deviation from general tendency induced by Experiment 2) were computed together with their p -values using the Matlab Statistics Toolbox (function `mnrfit`).

Results

We again calculated the percentage of “associative” responses for four different time bins of the learning task (again, 40 trials each), separately for the associative and non-associative condition. The percentage (SDs) in these four bins was 69.0% (7.7%), 94.7% (6.1%), 98.7% (2.7%) and 95.5% (6.4%) for the associative condition. For non-associative quadruplets, these were 32.2% (13.1%), 45.3% (16.6%), 49.2% (28.4%) and 57.2% (26.3%).

We then evaluated participants’ preferences. Across participants, the associative patterns were preferred over non-associative patterns on an average of 20.5 out of the 32 choices (64.1%), a significant departure from indifference (Figure 2, right; Wilcoxon Signed-Rank, $p < .005$).

In Experiment 2, there were even more wrong “associative” responses for the non-associative condition, but most presumably for exactly the same reason as described in Experiment 1. First of all, the familiarity or frequency of occurrence was matched exactly between associative and non-associative shapes. Under the instruction to detect novel configurations, and after seeing a familiar shape in the non-associative condition, participants might have been biased and decided that the whole configuration is not novel. Second, the inevitable repeated coupling of individual shapes or stimulus locations in the non-associative condition, based on random sampling from the 32 shapes and four locations, might have allowed for a *per chance* coupling in the non-associative condition. Again, these possibilities would make our results even more conservative.

Experiments 1 and 2 support our hypothesis that participants prefer shapes that are more associative (logistic regression coefficient value of 0.86, $p < .001$) and thereby offer a higher predictive

benefit over shapes that were equally familiar but carried no predictive advantage. This bias was found for participants being instructed to pay attention to the associative and to pay attention to the non-associative quadruplets during the initial acquisition phase. There was a trend towards slightly reduced preference of the associative shapes when subjects were instructed to attend to novel shape configurations (logistic regression coefficient value of -0.28 , $p = .048$), but the preference for associative shapes was still clearly present and significant. The shapes were arbitrary and meaningless, familiarity was matched and any other features were controlled by careful counterbalancing. The stimuli had not been used in a context in which the associative information provided a straightforward advantage for task performance. The learning task also provided no feedback, which subjects might have used as a proxy for likability in their subsequent preference judgement. However, overall lower performance in the previous cover-task in the non-associative condition suggests participants confused non-associative groupings with associative groupings (i.e., false alarms), and may therefore have experienced more uncertainty when faced with randomly grouped shapes. We cannot exclude the possibility that this negatively influenced their preference for those random patterns. Therefore, in the subsequent experiment, we investigated participant’s preference judgements for pictures of real-world objects that were rated independently for their degree of associativity. We reasoned that without any task at hand, preference judgements would not be influenced by putative facilitative effects that might accompany performance.

EXPERIMENT 3

In the third study, we aimed to address issues raised in Experiments 1 and 2, that although participants never received any feedback about their performance, higher response accuracy for the associative shapes during the training phase might have influenced preference judgements. We used

pictures of real-world objects, as such stimuli had not been used previously in a specific task context, in which one or another category would allow easier or more certain decisions. Specifically, we asked participants to choose their preferred object out of a series of randomly paired everyday objects, e.g., cowboy boots and Christmas tree. The advantage of using real objects is that their associativity does not have to be learned in a separate task, but already exists due to previous exposure in our natural environment. On the other hand, not teaching participants the associations in a controlled environment, as in Experiments 1 and 2, leaves more room for variability in individual pre-existing associations and valence.

Method

Participants

Thirty-eight participants¹ who were recruited from the Boston area gave written informed consent and were financially compensated for their participation.

Stimuli

We used stimuli that had been normed for associativity in a previous unpublished pilot study. This norming consisted of having 15 raters view 300 objects for 250ms each and then giving them 2s to name the first object association that came to mind. Based on these responses, the objects were split into three groups or conditions with varying level of associativity (see Figure 3 for example stimuli). Objects in the “Strong Context: Many Associations” (*Associative-Many*) condition elicited seven or more associations across multiple contexts, but at least 80% of associations belonged to one context. Objects in the “Strong Context: Few Associations” (*Associative-Few*) condition elicited no more than three associations, at least 60% of participants named the same association, and all associations belonged to the same context. For

example, cowboy boots might be associated with a cowboy, a horse and a gun, whereas a flyswatter is strongly associated with only a fly. Objects in the third condition (*Weakly-Associative*) elicited associations belonging to many different contexts and the majority of participants did not name any object associated with the target object within the 2s response window. A blanket, for example, may be associated with a large number of other objects, but these associations are relatively weak and inconsistent.

Additional low-level controls (colour, spatial frequency and semantic category). Object groups were equated for colorfulness and spatial frequency distributions using analyses of variance to compare average colour intensities [red, green, blue (RGB)] and average power spectra across conditions. There were no significant differences in colorfulness ($F_{2,165}$ (R,G,B) = 0.07, 0.58, 1.23, all p -values $\geq .30$) or in spatial frequency ($F_{2,165} = 0.72$, $p = .49$) between the three groups of objects. Two independent raters also confirmed that the objects within different groups were approximately equated along each of the following classifications: manmade, living, indoors vs. outdoors, curved vs. angular, symmetrical, edible and manipulable.

Additional high-level controls (valence and image quality). The objects used in the preference task were normalised for valence by randomly presenting each object to 20 independent raters for 1500ms each and asking them to rate the object on a scale of 1 (very unpleasant) to 7 (very pleasant). These objects were additionally rated for subjective image quality by eight independent raters who were asked to judge whether each image was suitable for inclusion in a product catalogue on a scale of 1 (discard) to 7 (publish immediately), with 4 being “publish with some edits”. The results of these normalising efforts are reported below.

¹Due to a coding error in the demographic data across the choice study and the valence and image quality norming studies associated with Experiment 3, exact distributions are unavailable for each separately. However, across the individuals who participated in one of these three studies, approximately 70% were female, and the approximate mean age was 21.8 years.



Figure 3. Example stimuli for Experiment 3, one representative image per condition. Left: Associative-Many, middle: Associative-Few, right: Weakly-Associative.

Procedure

Participants viewed a series of object pairs, one on each side of the screen, and were instructed to press a button indicating which object they preferred. Each pair consisted of two objects from different conditions that were presented together for 1s, and participants had 2s to respond. We presented 56 objects for the Associative-Many, Associative-Few and Weakly-Associative condition (168 objects total). Each participant saw every object only once, for a total of 84 trials (object pairs) per session. The objects in each pair, and the position of each on the screen were randomly assigned. Stimulus presentation and response collection for this task and the norming studies were performed using Psychtoolbox (Brainard, 1997) running in Matlab (Mathworks, Inc.).

Analysis

Primary analyses tested whether participants indicated an above-chance (>50%) preference for the more associative item of a given pair (i.e., Associative-Many rather than Associative-Few or Weakly-Associative, or Associative-Few rather than Weakly-Associative), and whether this could be accounted for by non-associative qualities of the images. The latter analyses employed a within-

subject logistic regression (with coefficients pooled at the group level), predicting binary choices (left vs. right object) based on: (1) whether the object on the right was in the more associative category of the pair (a binary indicator variable); (2) the difference in normed valence between the two objects (right vs. left); and (3) the difference in normed image quality between the two objects. These regressions also included an intercept to account for left/right biases. Regressions were performed within participants, and the resulting within-subject regression coefficients were pooled at the group level. Subsequent analyses examined differences across types of image pairs. Similarly to our hypothesis in Experiments 1 and 2, we predicted that participants would prefer Associative-Many objects, relative to the other two conditions, although we did not have a strong a-priori prediction as to whether we would also find a difference between Associative-Few and Weakly-Associative objects.

Results

Across participants, the more associative object on the screen was preferred on 53.5% of trials (Wilcoxon Signed-Rank, $p < .0001$). 76% of participants selected the more associative object of a pair on more than half of their choice trials. This

finding was confirmed with a within-subject logistic regression predicting the participant's preference between a given pair of objects based on whether that object was in the more associative category of the pair. Consistent with the choice proportions just described, this analysis found associativity to be a significant predictor of choice (average $\beta = 0.29$, $SE = 0.066$, $t_{37} = 4.4$, $p < .0001$).

While our object categories were equated for low-level features (e.g., colour, spatial frequency; see Method) we sought to also rule out potential high-level confounds that could have otherwise driven preferences for more associative objects. In particular, we suspected that two potent factors may influence preference independent of associativity: affective valence and subjective image quality. We therefore normalised for these two factors using independent groups of raters ($N = 20$ and $N = 8$, respectively), and found differences in average ratings of valence ($F_{2,38} = 10.34$, $p < .0005$) and image quality ($F_{2,14} = 4.58$, $p < .05$) across the object categories: average valence and subjective image quality ratings were slightly higher for objects in the Associative-Many ($M_{\text{valence}} = 4.6$, $M_{\text{quality}} = 4.0$) and Weakly-Associative ($M_{\text{valence}} = 4.5$, $M_{\text{quality}} = 4.0$) conditions than objects in the Associative-Few condition ($M_{\text{valence}} = 4.3$, $M_{\text{quality}} = 3.8$). Indeed, including these in the logistic regression described earlier identified a significant effect of relative valence on choice (average $\beta = 0.54$, $SE = 0.070$, $t_{37} = 7.7$, $p < .001$). The effect of relative image quality was non-significant (average $\beta = 0.03$, $SE = 0.043$, $t_{37} = 0.81$, $p = 0.42$). However, even when simultaneously controlling for both relative valence and relative image quality of the two objects, we found that preferences continue to be predicted by whether the chosen object was more associative than the unchosen object (average $\beta = 0.27$, $SE = 0.073$, $t_{37} = 3.7$, $p < .001$).

A set of post-hoc analyses examined each type of object pair separately. We found that participants showed an above-chance preference for Associative-Many objects over Associative-Few objects ($M = 59.0\%$, Wilcoxon Signed-Rank, $p < .0001$) and over Weakly-Associative objects ($M = 54.6\%$, $p < .005$). Unexpectedly, for the third type of object pair (Weakly-Associative vs. Associative-

Few) we found a small but significant preference for Weakly-Associative objects ($M = 53.2\%$, $p = 0.03$).

We did not have a strong a-priori prediction as to whether associative strength alone would be sufficient for driving preferences (and therefore Associative-Few would be preferred to Weakly-Associative) or if the number of associations would be a critical feature as well. This specific result may imply that the breadth of associations is more important than their strength. Along these lines, it may be that restricted associative activation elicited by Associative-Few objects (e.g., flyswatter might be only associated with a fly) is akin to rumination, which is prominent in depression, where thoughts are associative, but have a very narrow focus (Bar, 2009). However, given that Weakly-Associative objects were given the lowest valence ratings and valence was a significant predictor of choice (either due to an influence of associativity on valence, or for other reasons), we decided to perform the same test after residualising choices on valence ratings. Following this correction for differences in valence ratings, the preference for Weakly-Associative over Associative-Few objects disappeared ($M = 49.9\%$, $p = 0.86$), suggesting that subjective preferences in our study were influenced by the availability of many associations, but that participants were otherwise indifferent between categories with fewer associations.

In summary, Experiment 3 joins Experiments 1 and 2 in supporting the hypothesis that associative information is preferred by human observers. This, in turn, supports our proposal that information with predictive value is tagged affectively and might thereby encourage proactive search for it in the environment.

GENERAL DISCUSSION

We hypothesised that associative items are assigned a hedonic value because of their advantage for predictions. The data from the three experiments reported here support this hypothesis by showing that when given an option to choose between two items, people tend to prefer the one that is strongly and predictably associated with other items.

The predictive brain likes things that promote predictions

Since the early days of psychology, associative learning and prediction have been topics of extensive study. In Pavlovian learning, for example, a cue (conditioned stimulus) is repeatedly presented with a food stimulus (unconditioned stimulus). Those experiments demonstrated that the subjects generate a preparatory response when exposed to the conditioned stimulus, indicating expectation of the unconditioned stimulus. Numerous recent studies in neuroscience have investigated neural correlates of Pavlovian and instrumental reinforcement learning mechanisms. It has been reported in a plethora of studies that the brain signals when an expected reward is omitted or an unexpected reward is given (e.g., Bayer & Glimcher, 2005; Fiorillo, Tobler, & Schultz, 2003; Schultz, Dayan, & Montague, 1997). Those signals were referred to as negative prediction error (when an expected reward is omitted) and positive prediction errors (when an unexpected reward is given). There is emerging evidence that the concept of prediction error and computational ideas from reinforcement learning are also useful to describe visual perception (e.g., den Ouden, Friston, Daw, McIntosh, & Stephan, 2009; Iglesias et al., 2013). For example, predictive coding models posit that the brain is hierarchically organised and predicts sensory input via feedback connections, while only deviations or PE are being fed forward (Friston, 2005; Rao & Ballard, 1999).

However, the question whether violations from sensory expectations are experienced as aversive or, conversely, whether fulfilled sensory expectations are experienced as rewarding or hedonically satisfying has been rather neglected so far. A potential reason for not addressing the affective component of such putative predictive processing might be that predictive coding models in the context of reinforcement learning and vision are highly abstracted schemes to describe an organisms' behaviour. For example, Friston, Daunizeau, and Kiebel (2009) have illustrated computationally that an artificial agent's behaviour does not need to optimise reward as assumed by reinforcement learning, but can be sufficiently described by

surprise minimisation. While surprise minimisation may provide an appropriate description of human behaviour at an abstract level, it is important to understand how such a process is supported by the meso-limbic reward system, including areas such as the amygdala that signals events of positive and negative valence (e.g., Anders, Eippert, Weiskopf, & Veit, 2008; Paton, Belova, Morrison, & Salzman, 2006). In other words, it is important to know how such reward-machinery in a real biological infrastructure reflects and mirrors abstract principles of surprise minimisation.

There is evidence that predictability is linked to affective reactions. For example, Herry et al. (2007) investigated whether (temporally) unpredictable as compared to predictable neutral sensory information elicits aversive reactions in humans and mice. The authors constructed artificial, neutral sound pulse sequences that only differed in their temporal predictability. They found that the unpredictable stream generated anxiety-like responses and increased amygdala activity in both mice and humans. Conversely, it is an intriguing question whether a *fulfilled* sensory prediction—in a more neutral context—is associated with a reward signal. Although such reward signal will presumably be subtle and difficult to probe for empirically, a first step towards this end is to ask participants for their subjective preference for stimuli that offered sensory predictability by holding all other factors constant. In this study, we examined whether associative (visual) information, as a proxy for predictability, is preferred over non-associative information. Indeed, our data show a preference for arbitrary shapes or objects that facilitate such re-activation of related, associated information, indicating that such information is itself hedonically marked, and thereby influences preference judgements. In future studies, such a preference bias might be probed for with physiological measurements. For example, Harmon-Jones and Allen (2001) found that more familiar stimuli, which were also rated as more likable, evoked higher activity in the zygomaticus major—a muscle implicated in smiling—suggesting increased positive affect. Similarly, Winkielman and Cacioppo (2001) found that positively

evaluated, higher processing fluency was also accompanied by greater zygomatic activity. Thus, demonstrating such neurophysiological modulations during exposure of associative vs. non-associative patterns would expand the implications of the current hypothesis.

We have shown evidence for a preference for associative information that can facilitate predictions, and discussed this in the context of a “predictive brain”, whose primary goal is the minimisation of surprise or the processing of unexpected information. However, these findings also bear relevance to another topic in the literature, where similar brain networks mediate processing of reward and novelty (Bunzeck, Guitart-Masip, Dolan, & Düzel, 2011; Guitart-Masip, Bunzeck, Stephan, Dolan, & Düzel, 2010; Wittmann, Bunzeck, Dolan, & Düzel, 2007). It has been shown that (reward-independent) novelty of stimulus material increases neural activity in reward-relevant brain regions, such as the striatum (Wittmann, Daw, Seymour, & Dolan, 2008). Based on this overlap between brain circuitries underlying reward and novelty processing, it has been argued that the processing of novel information is intrinsically rewarding. This suggestion accords with reinforcement learning models that have conjectured a “novelty bonus” in order to explain and encourage exploratory behaviour (Dayan & Sejnowski, 1996; Kakade & Dayan, 2002). However, the findings could also be interpreted as reflecting our strive to reduce uncertainty by looking for predictability, and do not necessarily imply that novel information is experienced as “rewarding” and preferred over familiar information. In contrast, it has been reported in many studies with a range of different stimuli and in a variety of species that likability increases through mere exposure of information, i.e., familiarity of stimuli (Hill, 1978; Zajonc, 1980).

Mood and associative processing

It has been suggested that there is a reciprocal relation between associations and mood (Bar, 2009). While there is already evidence that positive mood facilitates divergent, broadly associative thinking

patterns (Fredrickson, 2004; Isen, Johnson, Mertz, & Robinson, 1985), it was argued that this relationship also exists in the reverse direction, whereby broad associations result in positive mood. This idea is indirectly supported by the phenomenon of rumination, a core symptom of depression, which can be conceptualised as restricted, narrow associative processing. More evidence for this proposal comes from a study in which the associativity of stimulus material was manipulated (Mason & Bar, 2012). The authors measured participants’ mood after they processed word lists with narrowly constrained associative relations vs. word lists with associative relations that progressively broadened. The data show that participants’ mood was relatively better after periods of broad vs. constrained processing (see also Brunye et al., 2013). The present study adds a potential explanation for this phenomenon by offering evidence that associative information is intrinsically rewarding, presumably because of the predictability it offers.

The present findings raise the question of what is meant by associative information being intrinsically rewarding. It may be that the link between affect and associative processing is “hard-wired” in the brain, as it ultimately serves the pressure of maintaining a certain degree of stability and constancy of our ecological niche. Shenhav, Barrett, and Bar (2013) provided evidence for such a relationship by having participants view objects varying orthogonally in affective and associative properties while undergoing fMRI. Even though the participants performed an object recognition task unrelated to either object property, the authors found that increasing positive affective valence and greater associativity each led to increased activity in the medial orbitofrontal cortex. This region has been implicated in processes related to value-based decision-making (Hampton, Bossaerts, & O’Doherty, 2006; Kable & Glimcher, 2007; Knutson, Taylor, Kaufman, Peterson, & Glover 2005). The results from Shenhav et al. (2013) suggest that affect and associativity are intimately intertwined in the brain. Still, the question of how exactly it is that spreading associative information elicits affective reactions remains an important avenue for future research.

Implications for aesthetic research

The concept of processing *fluency* is a predominant explanatory tool in empirical aesthetics. Reber, Schwarz, and Winkielman (2004) proposed that aesthetic pleasure is a function of the perceiver's processing experience: the more fluently perceivers can process a piece of art, the more positively it is evaluated. *Perceptual fluency* is defined as the ease of identifying the physical properties of the stimulus, and is influenced by variables such as perceptual priming, presentation duration, repetition or figure-to-ground contrast. *Conceptual fluency* is defined as the ease of mental operations concerned with a stimulus' meaning and its relation to semantic knowledge structures. For example, Temme (1992) found that additional information, as a proxy for better conceptual accessibility, increased the enjoyment of seventeenth and nineteenth century Dutch paintings, particularly for participants who lacked formal education in art. In another study, Reber, Winkielman, and Schwartz (1998) primed pictures with a matching or an unrelated prime. When the pictures were preceded by matching primes, participants' liked them significantly more than pictures preceded by unrelated primes. These and related findings were explained by *conceptual fluency* which increases with additional information and ultimately enables "cognitive mastering" of the artwork (Leder, Belke, Oeberst, & Augustin, 2004). However, these results are also compatible with an associative account: here, additional information might simply increase the breadth of associations, which allow embedding the artwork in a well-known semantic network with a predictable structure, resulting in greater positive affect, which is hence attributed to the artwork (affect-as-information, Forgas, 1995; Schwarz & Clore, 1983). However, the present explanation and the fluency-account are not mutually exclusive, but rather complement each other: as "fluency" refers to the perceivers' processing experience, one would still like to know what makes processing of one stimulus more "fluent" than the processing of another. We here offer associativity or predictability as a factor that has, to the best of our knowledge, not yet been reported in the aesthetic literature, as the focus

has been on properties of single stimuli, and not their interrelatedness. Taken together, it seems that the propensity to activate associations is an important factor that influences preference decisions and might therefore be considered in future studies aiming to explain aesthetic judgements.

Concluding remarks and outlook

The present study demonstrates that associative information is preferred over non-associative information. Given that associative information can provide a platform for building predictions (Bar, 2007; Turk-Browne, Scholl, Johnson, & Chun, 2010), these data strongly suggest a link between preference and predictability. Such a link fits nicely in the framework of a "predictive brain", which has been suggested to be primarily driven by the imperative to minimise surprise (Friston, 2009). Without a doubt, the interplay between affective reactions to expected and unexpected, surprising input is not linear, and presumably modulated by many factors. For example, in the context of an art gallery or a piece of music, violations of expectations might be a highly desirable feature and experienced as aesthetically pleasant (Schafer, Overy, & Nelson, 2013). Individual factors, like level of anxiety or tendency for thrill-seeking behaviour might further act as modulators in the context of evaluating prediction errors and stimulus predictability. These questions point to an exciting novel research agenda that has to be unpacked in future work.

Manuscript received 28 April 2014

Revised manuscript received 30 July 2014

Manuscript accepted 12 September 2014

First published online 9 October 2014

REFERENCES

- Alink, A., Schwiedrzik, C. M., Kohler, A., Singer, W., & Muckli, L. (2010). Stimulus predictability reduces responses in primary visual cortex. *Journal of Neuroscience*, 30, 2960–2966. doi:10.1523/JNEUROSCI.3730-10.2010

- Aminoff, E., Gronau, N., & Bar, M. (2007). The parahippocampal cortex mediates spatial and non-spatial associations. *Cerebral Cortex*, 17, 1493–1503. doi:10.1093/cercor/bhl078
- Aminoff, E. M., Kveraga, K., & Bar, M. (2013). The role of the parahippocampal cortex in cognition. *Trends in Cognitive Sciences*, 17, 379–390. doi:10.1016/j.tics.2013.06.009
- Anders, S., Eippert, F., Weiskopf, N., & Veit, R. (2008). The human amygdala is sensitive to the valence of pictures and sounds irrespective of arousal: An fMRI study. *Social Cognitive and Affective Neuroscience*, 3, 233–243. doi:10.1093/scan/nsn017
- Bar, M. (2004). Visual objects in context. *Nature Reviews: Neuroscience*, 5, 617–629. doi:10.1038/nrn1476
- Bar, M. (2007). The proactive brain: Using analogies and associations to generate predictions. *Trends in Cognitive Sciences*, 11, 280–289. doi:10.1016/j.tics.2007.05.005
- Bar, M. (2009). A cognitive neuroscience hypothesis of mood and depression. *Trends in Cognitive Sciences*, 13, 456–463. doi:10.1016/j.tics.2009.08.009
- Bar, M., Kassam, K. S., Ghuman, A. S., Boshyan, J., Schmid, A. M., Dale, A. M., & Halgren, E. (2006). Top-down facilitation of visual recognition. *Proceedings of the National Academy of Sciences of the United States of America*, 103, 449–454. doi:10.1073/pnas.0507062103
- Bayer, H. M., & Glimcher, P. W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, 47(1), 129–141. doi:10.1016/j.neuron.2005.05.020
- Brunye, T. T., Gagnon, S. A., Paczynski, M., Shenhav, A., Mahoney, C. R., & Taylor, H. A. (2013). Happiness by association: Breadth of free association influences affective states. *Cognition*, 127(1), 93–98. doi:10.1016/j.cognition.2012.11.015
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 443–446. doi:10.1163/156856897X00357
- Bunzeck, N., Guitart-Masip, M., Dolan, R. J., & Düzel, E. (2011). Contextual novelty modulates the neural dynamics of reward anticipation. *Journal of Neuroscience*, 31, 12816–12822. doi:10.1523/JNEUROSCI.0461-11.2011
- Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, 36(1), 28–71. doi:10.1006/cogp.1998.0681
- Dayan, P., & Sejnowski, T. J. (1996). Exploration bonuses and dual control. *Machine Learning*, 25, 5–22.
- Den Ouden, H. E., Friston, K. J., Daw, N. D., McIntosh, A. R., & Stephan, K. E. (2009). A dual role for prediction error in associative learning. *Cerebral Cortex*, 19, 1175–1185. doi:10.1093/cercor/bhn161
- Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science*, 299, 1898–1902. doi:10.1126/science.1077349
- Forgas, J. P. (1995). Mood and judgment: The affect infusion model (AIM). *Psychological Bulletin*, 117(1), 39–66. doi:10.1037/0033-2909.117.1.39
- Fredrickson, B. L. (2004). The broaden-and-build theory of positive emotions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 359, 1367–1377. doi:10.1098/rstb.2004.1512
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360, 815–836. doi:10.1098/rstb.2005.1622
- Friston, K. (2009). The free-energy principle: A rough guide to the brain? *Trends in Cognitive Sciences*, 13, 293–301. doi:10.1016/j.tics.2009.04.005
- Friston, K. J., Daunizeau, J., & Kiebel, S. J. (2009). Reinforcement Learning or Active Inference? *PLOS One*, 4(7), e6421–529. doi:10.1371/journal.pone.0006421
- Guitart-Masip, M., Bunzeck, N., Stephan, K. E., Dolan, R. J., & Düzel, E. (2010). Contextual novelty changes reward representations in the striatum. *Journal of Neuroscience*, 30, 1721–1726. doi:10.1523/JNEUROSCI.5331-09.2010
- Hampton, A. N., Bossaerts, P., & O'Doherty, J. P. (2006). The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. *Journal of Neuroscience*, 26, 8360–8367. doi:10.1523/JNEUROSCI.1010-06.2006
- Harmon-Jones, E., & Allen, J. J. B. (2001). The role of affect in the mere exposure effect: Evidence from psychophysiological and individual differences approaches. *Personality and Social Psychology Bulletin*, 27, 889–898. doi:10.1177/0146167201277011
- Herry, C., Bach, D. R., Esposito, F., Di Salle, F., Perrig, W. J., Scheffler, K., & Seifritz, E. (2007). Processing of temporal unpredictability in human and animal amygdala. *Journal of Neuroscience*, 27, 5958–5966. doi:10.1523/JNEUROSCI.5218-06.2007
- Hill, W. F. (1978). Effects of mere exposure on preferences in nonhuman mammals. *Psychological Bulletin*, 85(6), 1177–1198. doi:10.1037/0033-2909.85.6.1177

- Iglesias, S., Mathys, C., Brodersen, K. H., Kasper, L., Piccirelli, M., den Ouden, H. E., & Stephan, K. E. (2013). Hierarchical prediction errors in midbrain and basal forebrain during sensory learning. *Neuron*, 80, 519–530. doi:10.1016/j.neuron.2013.09.009
- Isen, A. M., Johnson, M. M., Mertz, E., & Robinson, G. F. (1985). The influence of positive affect on the unusualness of word associations. *Journal of Personality and Social Psychology*, 48, 1413–1426. doi:10.1037/0022-3514.48.6.1413
- Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, 10, 1625–1633. doi:10.1038/nn2007
- Kakade, S., & Dayan, P. (2002). Dopamine: Generalization and bonuses. *Neural Networks*, 15, 549–559.
- Knutson, B., Taylor, J., Kaufman, M., Peterson, R., & Glover, G. (2005). Distributed neural representation of expected value. *Journal of Neuroscience*, 25, 4806–4812. doi:10.1523/JNEUROSCI.0642-05.2005
- Kok, P., Jehee, J. F., & de Lange, F. P. (2012). Less is more: Expectation sharpens representations in the primary visual cortex. *Neuron*, 75, 265–270. doi:10.1016/j.neuron.2012.04.034
- Leder, H., Belke, B., Oeberst, A., & Augustin, D. (2004). A model of aesthetic appreciation and aesthetic judgments. *British Journal of Psychology*, 95, 489–508. doi:10.1348/0007126042369811
- Mason, M. F., & Bar, M. (2012). The effect of mental progression on mood. *Journal of Experimental Psychology: General*, 141, 217–221. doi:10.1037/a0025035
- McClure, S. M., Berns, G. S., & Montague, P. R. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron*, 38, 339–346. doi:10.1016/S0896-6273(03)00154-5
- Ogawa, H., & Watanabe, K. (2011). Implicit learning increases preference for predictive visual display. *Attention, Perception, & Psychophysics*, 73(6), 1815–1822. doi:10.3758/s13414-010-0041-2
- Paton, J. J., Belova, M. A., Morrison, S. E., & Salzman, C. D. (2006). The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature*, 439, 865–870. doi:10.1038/nature04490
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87. doi:10.1038/4580
- Reber, R., Winkielman, P., & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science*, 9, 45–48. doi:10.1111/1467-9280.00008
- Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? *Personality and Social Psychology Review*, 8, 364–382. doi:10.1207/s15327957pspr0804_3
- Schaefer, R. S., Overy, K., & Nelson, P. (2013). Affect and non-uniform characteristics of predictive processing in musical behaviour. *Behavioral and Brain Sciences*, 36(3), 226–227. doi:10.1017/S0140525X12002373
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275, 1593–1599. doi:10.1126/science.275.5306.1593
- Schwarz, N., & Clore, G. L. (1983). Mood, misattribution and judgment of well-being. Informative and directive functions of affective states. *Journal of Personality and Social Psychology*, 45, 513–523. doi:10.1037/0022-3514.45.3.513
- Shenhav, A., Barrett, L. F., & Bar, M. (2013). Affective value and associative processing share a cortical substrate. *Cognitive, Affective, & Behavioral Neuroscience*, 13(1), 46–59. doi:10.3758/s13415-012-0128-4
- Summerfield, C., & Koechlin, E. A. (2008). Neural representation of prior information during perceptual inference. *Neuron*, 59, 336–347. doi:10.1016/j.neuron.2008.05.021
- Temme, J. E. (1992). Amount and kind of information in museums: Its effect on visitors satisfaction and appreciation of art. *Visual Arts Research*, 18(2), 74–81.
- Todorovic, A., van Ede, F., Maris, E., & de Lange, F. P. (2011). Prior expectation mediates neural adaptation to repeated sounds in the auditory cortex: An MEG study. *Journal of Neuroscience*, 31, 9118–9123. doi:10.1523/JNEUROSCI.1425-11.2011
- 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. doi:10.1037/0096-1523.35.1.195
- 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, 11177–11187. doi:10.1523/jneurosci.0858-10.2010
- Winkielman, P., & Cacioppo, J. T. (2001). Mind at ease puts a smile on the face: Psychophysiological evidence that processing facilitation elicits positive

- affect. *Journal of Personality and Social Psychology*, 81, 989–1000. doi:[10.1037/0022-3514.81.6.989](https://doi.org/10.1037/0022-3514.81.6.989)
- Wittmann, B. C., Bunzeck, N., Dolan, R. J., & Düzel, E. (2007). Anticipation of novelty recruits reward system and hippocampus while promoting recollection. *NeuroImage*, 38, 194–202. doi:[10.1016/j.neuroimage.2007.06.038](https://doi.org/10.1016/j.neuroimage.2007.06.038)
- Wittmann, B. C., Daw, N. D., Seymour, B., & Dolan, R. J. (2008). Striatal activity underlies novelty-based choice in humans. *Neuron*, 58, 967–973. doi:[10.1016/j.neuron.2008.04.027](https://doi.org/10.1016/j.neuron.2008.04.027)
- Yue, X., Vessel, E. A., & Biederman, I. (2007). The neural basis of scene preferences. *NeuroReport*, 18, 525–529. doi:[10.1097/WNR.0b013e328091c1f9](https://doi.org/10.1097/WNR.0b013e328091c1f9)
- Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist*, 35, 151–175. doi:[10.1037/0003-066X.35.2.151](https://doi.org/10.1037/0003-066X.35.2.151)