Multivariate pattern analysis of event-related potentials predicts the subjective relevance of everyday objects

William Francis Turner[1], Phillip Johnston[1], Kathleen de Boer[1], Carmen Morawetz[2],

& Stefan Bode[1]*

[1] Melbourne School of Psychological Sciences, The University of Melbourne, Parkville 3010, Victoria, Australia

[2] Department of Education and Psychology, Freie Universität Berlin, Habelschwerdter Allee 45, 14195 Berlin, Germany

*Corresponding Author: Stefan Bode

Email: sbode@unimelb.edu.au

Abstract

Potentially decision-relevant stimuli have been proposed to undergo immediate semantic processing. The current study investigated whether information regarding the general desirability ('Wanting') of visually presented 'everyday' objects was rapidly and automatically processed. Participants completed a foreground task while their electroencephalogram (EEG) was recorded, and task-irrelevant images were presented in the background. Following this, participants rated the images with regards to Wanting and the potentially related attributes of Relevance, Familiarity, Aesthetic Pleasantness and Time Reference. Multivariate pattern classification was used to predict the ratings from patterns of EEG data. Prediction of Wanting and Relevance was possible between 100-150 ms following stimulus presentation. The other dimensions could not be predicted. Wanting and Relevance ratings were highly correlated and displayed similar feature weight maps. The current results suggest that the general desirability and subjective relevance of everyday objects is rapidly and automatically processed for a wide range of visual stimuli.

Keywords

1. Introduction

There is now a large body of research suggesting that many unconscious factors can influence decision outcomes (Custers & Aarts, 2010). These factors include simple biases of previous choices in arbitrary decision scenarios (Bode et al., 2014b) as well as more complex biases in financial decisions, exerted, for example, by incidental rewarding stimuli on decisions between small immediate rewards and larger delayed rewards in intertemporal choice (e.g., Kim & Zauberman, 2013; Murawski et al., 2012; Simmank, et al., 2015; Wilson & Daly, 2004; Van den Bergh, et al., 2008; Zhong & Devoe, 2010). In general, the presence of rewards, specific social situations, and (subliminal) priming has been shown to play a role in biasing decision outcomes in a variety of situations, outside the awareness of the decision-maker (Aarts et al., 2007; Custers & Aarts, 2010; Dijksterhuis & Aarts, 2010; Zedelius et al., 2014).

How these unconscious biases are generated is, however, still debated. For example, with intertemporal choices it might be plausible that any positive stimulus in the environment activates neural reward circuits, increasing the likelihood of reward seeking behaviour in a subsequent decision task. In support of this, many studies have found that images of rewarding objects, such as brand logos (Murawski et al., 2012; Zhong & Devoe, 2010), sexual cues (Van den Bergh et al., 2008; Wilson & Daly, 2004), or food and status symbols (Simmank, et al., 2015) can bias decision-makers towards immediate gratification. However, others have shown that decision-makers can also be biased towards delayed gratification (Benoit et al., 2011; Cheng et al., 2012; Peters & Büchel, 2010), suggesting a more complex mechanism.

Given the variety of cues that have been shown to bias decision-making and goals (Custers & Aarts, 2010), one possibility is that stimuli which exert such biases derive their impact from a higher-level semantic analysis. A recent study by Bode and colleagues (2014a) made an initial step towards testing this proposal by exposing participants to a variety of positively valenced stimuli (similar to those found to influence intertemporal decisions in previous studies) using a passive viewing paradigm with an attention-demanding foreground task. Following this task, participants were asked to rate all images with regards to whether they were felt to be subjectively related to the present or the future (a stimulus attribute termed 'Time Reference'). It was shown that these ratings could be predicted successfully from distributed patterns of event-related potentials (ERPs) recorded during passive viewing, despite the fact that, whilst initially viewing the images, the participants were unaware that they would be required to subsequently rate them. This finding supports the assumption that a fast and automatic analysis of semantic stimulus features takes place. Importantly, such an analysis might not only explain carry-over effects on incidental decision processes (e.g. intertemporal choice), but might also facilitate everyday decision-making for the analysed objects themselves by preparing unconscious shortcuts which streamline the decision-making process (Bode et al., 2014a; Creswell et al., 2013; Dijksterhuis et al., 2006a,b).

For this to be true, several aspects of stimuli would have to be automatically processed, integrated, and made available as a general decision-related signal to the decision-maker. Several studies using functional magnetic resonance imaging (fMRI) have investigated whether categorical stimulus attributes other than Time Reference are represented in brain activity directly following stimulus presentation. For example, O'Doherty and colleagues (2003) demonstrated that when individuals were presented with images of faces, brain activity was predictive of subsequent attractiveness ratings, even though the experimental task was to judge gender. Others have shown that the processing of

stimuli on a category level occurred rapidly following the presentation of a natural scene, even though the scene was unattended and irrelevant to the experimental task (Peelen et al., 2009). Brain activity has also been found to be systematically modulated according to individual preferences for specific stimuli, even when attention was directed away from the stimuli (Lebreton et al., 2009; Tusche et al., 2010; 2013). However, given the poor temporal resolution of fMRI, it remains unclear how rapidly the suggested processes took place. A number of studies have used electroencephalography (EEG), and its relatively superior temporal resolution, to demonstrate that semantic stimulus attributes, such as basic category information (e.g. cars vs. buildings), is rapidly reflected in brain activity following presentation of visual objects (Simanova et al., 2010; Taghizadeh-Sarabi et al., 2015; Wang et al., 2012).

Crucially, information regarding Time Reference as well as category membership is not necessarily directly relevant to object-related decision making. Therefore, the results from previous studies provide insufficient evidence that information of direct relevance for decision-making is rapidly and automatically processed following exposure to a visual stimulus. Notably, there is some evidence that people's choices for consumer products can be predicted from their brain activity recorded during passive viewing of the products (Levy et al., 2011; Telpaz et al., 2015; Tusche et al., 2010). For example, an fMRI study, which measured brain activity during passive exposure to images of cars, demonstrated that hypothetical purchase decisions could be directly predicted from patterns of brain activity in decision-related brain regions such as the medial prefrontal cortex and insula (Tusche et al., 2010). However, these results were obtained in a sample of participants who indicated a high interest in cars prior to the experiment. Thus, these stimuli were highly valenced for the selected sample. Similarly, previous studies have typically used highly valenced and optimised stimuli (e.g., Bode et al., 2014a; Levy et al., 2011; Telpaz et al., 2015), which had a high likelihood of being relevant for decision-making. For example, the stimuli used by Bode and colleagues (2014a) were positively valenced and were selected on the basis that they strongly related to the present or the future (i.e. they were optimised for the Time Reference dimension). While this approach is valid and arguably helpful for optimising the probability of detecting the neural signatures of these processes, it also bears the danger of communicating a distorted picture of stimulus processing for decision-making. In fact, it could imply that all stimuli are unconsciously processed in great semantic depth, or that detailed purchase decisions are always prepared unconsciously for incidentally encountered objects in the environment. However, the question remains as to whether the automatic extraction of decision-relevant information also takes place for more 'everyday' objects, which have not been optimised in this manner. One possibility is that in-depth semantic processing first requires the detection of relevance or a general desire towards this object, and that any further processing is abandoned if the object does not fulfil these criteria. As suggested for early visual processing (Felsen & Dan, 2005) as well as social stimuli (Zaki & Ochsner, 2009), processing of naturalistic stimuli might indeed differ substantially from optimised stimuli.

The current study addressed this question by exposing participants to a selection of object images drawn from a novel picture set, the Nencki Affective Picture System (NAPS; Marchewka et al., 2014), for which normative ratings for valence, approach/avoidance and arousal were available. Notably, these images fell into a variety of categories, did not depict consumer products, and varied highly across normative valence ratings. As such, they were more representative of objects encountered in everyday life than the stimuli used in previous studies. We then asked whether the general desirability of these objects could still be predicted from brain activity shortly after stimulus presentation, despite having waived the

characteristic of high valence. Moreover, following previous studies (Bode et al., 2014a; Tusche et al., 2010) we made the stimuli task-irrelevant by presenting them in the background while participants focussed on a foreground task. This corresponded closely to everyday scenarios in which objects are perceived incidentally, and are not the focus of direct attention or deliberation.

There are many ways to characterise the general desirability of, or tendency to be drawn towards, an object, including: preferences (Lebreton et al., 2009), approach motivation (Harmon-Jones & Allen, 1998), utility (Fishburn, 1970), and valence (Mogg et al., 2003). However, most of these concepts are not easily applied to everyday stimuli, which are not in fact owned (or can be owned) by an individual. They also do not adequately capture the intuitive concept of a general desire to 'want' or have/own/interact with an object that could be the result of a fast and automatic aggregated semantic analysis of decision-relevant stimulus features. Hence, for the sake of this study, this decision-relevant aggregate concept was simply termed 'Wanting', which is an intuitive terminology which our participants indicated they understood. The concept of Wanting has been used in addiction- and food-related research to refer to the implicit desire of a rewarding stimulus that is distinct from the hedonic impact, or 'liking', of the stimulus (Berridge et al., 2009). In the current study, participants made explicit Wanting judgements; thus, the definition of Wanting used here differs slightly from that which is typically used in addiction research. However, it is noted that the concept of explicit Wanting has also been used in a number of previous studies (e.g., Finlayson et al., 2008). Due to practical constraints, the current study did not attempt to fully unpack which single stimulus features were integrated when forming aggregate Wanting judgements. We did, however, include four additional stimulus attributes, which might plausibly influence Wanting judgements and have been suggested to be processed pre-attentively and automatically: Relevance (Sander et al., 2003), Familiarity (Wang, et al., 1994), Aesthetic Pleasantness (Kühn & Gallinat, 2012; O'Doherty et al., 2003), and Time Reference (Bode et al., 2014a). The latter provided a test of whether the findings by Bode and colleagues (2014a) would replicate for everyday stimuli, which were not optimised for this dimension. Crucially, participants were only asked to explicitly rate the stimuli with regards to these attributes following initial stimulus exposure. Thus, they were not primed in any way to evaluate the stimuli with regards to these attributes.

Following Bode and colleagues (2014a), we used multivariate pattern analysis (MVPA; Blankertz et al., 2011; Bode et al., 2012; King & Dehaene, 2014; Philiastides & Sajda, 2006), optimised for event-related potential data (Bode et al., 2012; Bode & Stahl, 2014), to directly predict the participants' ratings from spatially distributed activity patterns within the first few hundred milliseconds following stimulus presentation. This approach has been shown to be sensitive to even subtle decision-related information, which cannot be detected using conventional methods (e.g., Bode et al., 2012; 2014a). The second advantage of MVPA is that it allows for a (spatially and temporally) unbiased search for this information that does not rely on *a priori* knowledge of specific neural generators or components of the event-related potential linked to these dimensions. MVPA involves training a classifier (a pattern classification algorithm) to distinguish between patterns of brain activity associated with different experimental variables of interest (here the ratings for each dimension). Successful prediction provides evidence that the respective information of interest is represented in brain activity patterns during distinct analysis time windows. Finally, in addition to attempting to predict the ratings, we also related them to the NAPS normative ratings for valence, arousal and approach/avoidance.

It was hypothesised that it would be possible to predict participants' subsequent subjective Wanting ratings from patterns of brain activity recorded during the presentation of non-optimised, task-irrelevant everyday objects. Additionally, we investigated whether information regarding stimulus Relevance, Familiarity, Aesthetic Pleasantness and Time Reference was also represented, and how these representations might have contributed to an integrated neural Wanting-signal.

2. Materials and methods

2.1 Participants

Twenty-five right-handed students from the University of Melbourne, Australia, fluent in English and with normal or corrected to normal vision, gave written informed consent and were compensated with AUD 20 for their time. Three data sets had to be discarded due to excessive skin potential artifacts in the EEG. The final sample consisted of 22 participants (mean age = 21.7 years, range = 18-34 years; 16 females). The experimental procedures were approved by the ethics committee of the University of Melbourne (ID 1443258).

2.2 Stimuli and apparatus

Fifty images of objects were selected from the Nencki Affective Picture System (NAPS; Marchewka et al., 2014) based on category membership and normative NAPS ratings. First, for standardisation purposes, all vertically oriented images were rejected from the 329 available object images. The remaining 253 images were sorted into 17 semantic categories (vegetables, meat, drinks, other foods, cars/trucks, bicycles/motorcycles, boats/planes/trains, toys, weapons/tools, rubbish, electronics, bottles/glass, bathroom items, kitchen items, shoes, sculptures/art, and miscellaneous). Then, 50 items were selected based on an algorithm that sampled ~3 images from each category such that one item each was of low, medium and high valence, respectively (see Supplement A for list of chosen items and their normative ratings). Note that unlike in previous studies, images were not all highly valanced, were not optimised for a specific rating dimension (Bode et al., 2014a), nor were they explicit consumer products (e.g., Tusche et al., 2010). Each image was resized to 600 x 450 pixels for the experiment.

The experimental task was presented on a Dell P2210 LCD monitor, with a resolution of 1680 x 1050 pixels and a frame rate of 60 Hz, using Psychtoolbox (Matlab R2012b; Brainard, 1997). Participants were seated with their chin on a chinrest approximately 50 cm from the computer monitor such that the images subtended 18.07° x 13.57° of visual angle.

2.3 Experimental procedures

On each trial, a central fixation box (22 x 22 pixels; 0.69° x 0.69° of visual angle) was shown on a grey background for 2-3 s (drawn from a uniform distribution). Following this, an image was presented in the background for 3.2 s while the fixation box constantly remained on the screen. Image presentation was again followed by another jittered fixation-only period (2-3 s) before the next image was shown. Each image was presented three times (a total of 150 trials) in an individually randomised order. Participants were given a break at the halfway point.

Participants were asked to monitor the fixation box, which would fill into a solid black square (after a 1-2 s delay following image presentation; delays were normally distributed) on randomly selected trials in 30% of all trials (see Figure 1). Participants were instructed to press the space bar (on a standard computer keyboard) immediately after the box filled. This task was designed to render the images task-irrelevant and prevent uneven allocation of attention, in line with previous studies, which used similar but slightly more demanding foreground tasks (e.g., Bode et al., 2014a; Tusche et al., 2010).

After the experiment, participants completed a computer-based image-rating task in which the same 50 images were rated on 9-point scale (1 = not at all, 9 = very much / 1 = past, 9 = future) with respect to Wanting ("How much do you want this object?"), Relevance ("How relevant is this image to you personally?"), Familiarity ("How familiar is this object for you?"), Aesthetic Pleasantness ("How aesthetically pleasant is the image for you?"), and Time Reference ("Is this item for you related to the past or the future?"). The order in which images and questions were presented was again randomised. Crucially, whilst viewing the images during the experiment, participants were neither aware of this subsequent rating task, nor of the attributes which they would be asked about.

2.4 EEG data recording and processing

During the experiment, EEG activity was recorded using 64 Ag/AgCl electrodes (Fp1, Fpz, Fp2, AF7, AF3, AFz, AF4, AF8, F7, F5, F3, F1, Fz, F2, F4, F6, F8, FT7, FC5, FC3, FC1, FCz, FC2, FC4, FC6, FT8, T7, C5, C3, C1, Cz, C2, C4, C6, T8, TP7, CP5, CP3, CP1, CPz, CP2, CP4, CP6, TP8, P9, P7, P5, P3, P1, Pz, P2, P4, P6, P8, P10, PO7, PO3, POz, PO4, PO8, O1, Oz, O2, Iz) interfacing a Biosemi Active-II system running ActiView acquisition software. The electrodes were attached to a Biosemi fabric cap according to the international 10-20 system. The EEG was continuously recorded at a sampling rate of 512 Hz. An implicit reference was used during recording, and all channels were re-referenced off-line against the average of two electrodes placed on the mastoid bones. Further external electrodes were placed below and next to the left eye to record the vertical and horizontal electrooculogram.

Data were pre-processed in Matlab R2012b using the EEGLAB Toolbox (Delorme & Makeig, 2004). A standard 0.1 to 70 Hz band-pass filter was applied, and a 50 Hz notch filter was used offline to remove electrical line noise. The data were divided into epochs ranging from 100 ms before to 600 ms after presentation of each stimulus. The 100 ms period before stimulus presentation was used as the baseline. Epochs containing muscle movements and skin potentials were removed after visual screening. Channels with poor quality data were visually identified and corrected using spline interpolation. An independent-components analysis (ICA) was performed to identify and remove components related to eye-blinks and eye-movements. A stricter visual rejection of artifacts followed in which all epochs with amplitudes exceeding ± 500 µV were excluded. Finally, a standard current source density (CSD) analysis was conducted at each of the electrode sites using the CSD toolbox (Kayser & Tenke, 2006). This analysis involves calculating the second derivative of the distribution of the voltage over the scalp (Perrin et al., 1987; 1989), and has been shown to benefit support vector machine (SVM) classification of ERPs (Bai et al., 2007; Bode et al., 2012; 2014a; Bode & Stahl, 2014).

2.5 Multivariate Pattern Analysis (MVPA) of ERP data

Due to the non-normal distributions of participants' ratings for single dimensions, there was not sufficient variance to allow for a regression approach of specific rating values using multivariate support vector regression (SVR) analysis (as conducted in Bode et al.,

2014a; see Supplement B for individual rating distributions). Consequently, ratings were converted to a binary score ('high', 'low') relative to the group median for the particular attribute being analysed. A linear SVM classification approach (Bode et al., 2012; Bode & Stahl, 2014) was then used to predict these scores from distributed patterns of ERP data (Figure 2). To avoid potential sample size biases due to different numbers of trials, all analyses were based on a balanced number of trials associated with 'low' and 'high' ratings, randomly chosen from all available trials for each participant (Bode et al., 2014a; Bode & Stahl, 2014). Participants who had less than 20 trials relating to either 'low' or 'high' ratings for a particular attribute were excluded from the respective analysis. This means, separated by dimension, we included on average: Wanting 60.95 trials (range 20-100), Relevance 72.63 trials (range 20-120), Aesthetic Pleasantness 77.0 trials (range 40-120), Familiarity 71.0 trials (range 20-120), and Time Reference 87.62 trials (range 60-120).

We analysed the 100 ms of data pre-stimulus (as a neutral baseline) and the 600 ms following stimulus onset using 20 ms analysis time windows, which were moved through the data using a sliding-window approach with 10 ms step-size (Figure 2). For each step, the data from all channels within the 20 ms analysis time window were transformed into vectors representing the spatio-temporal activity patterns and sorted into trials corresponding to images with 'high' or 'low' post-experimental ratings. Data from both trial types was then randomly sorted into ten sets of equal size. A linear SVM classifier (as implemented by the LIBSVM Toolbox, Chang & Lin, 2011; standard regulation parameter C = 1) was trained on 90% of the data (9 sets) and estimated a decision hyperplane/boundary for the classification of 'high' and 'low' rated objects. The estimated decision boundary was then used to classify the vectors of the remaining 10% of the data (1 set) as either 'high' or 'low'. This process was repeated using a 10-fold cross-validation procedure, using each set as the test data set once while the classifier was independently re-trained on data from the remaining 9 sets. To rule out any potential drawing biases, a conservative accuracy estimation approach was taken, and the entire cross-validated classification procedure was itself repeated ten times (Bode et al., 2014a; Bode & Stahl, 2014). The classification accuracy for each analysis time window (i.e. the accuracy with which the recoded binary ratings could be predicted) was obtained by averaging the results from these 100 analyses. Independent analyses were performed for each participant, and again separately for each rating dimension. Thus, each participant's individual CSD-EEG data were separately used for training the SVM classifier and for testing for each analysis.

Statistical testing (*T*-tests) at group level (independently conducted for each rating dimension) was performed by comparing the classification accuracy to an empirical chance distribution instead of the theoretical chance level, which has been criticised recently as too lenient (Combrisson & Jerbi, 2015). The empirical chance distribution was obtained by repeating the classification procedure exactly as outlined above, but with the labels ('high' or 'low') randomly shuffled before classification. Any systematic biases in the data would therefore also affect the chance distribution, further decreasing the probability of false positive results. Statistical testing was carried out separately for each time window and a cluster-based correction for multiple comparisons was applied (1000 permutations, alpha level of 0.05; Bullmore et al., 1999; Maris & Oostenveld, 2007). This correction tests for above chance classification across clusters of time-points and involves summing the t-statistic across consecutive time windows, which showed statistically significant decoding. These summed values (cluster masses) are then compared against a distribution of cluster masses generated from permutation samples of decoding accuracies (see Maris & Oostenveld, 2007). Critically, we chose this correction method as it has been shown to be sensitive to small, temporally sustained effects (Grootswagers et al., 2016), thus, it enabled us to test for

sustained information in the EEG data that was predictive of participants' ratings. In addition to testing significance at single time-points, we tested for significant differences across 50 ms time windows as another measure for sustained predictive information within the EEG data. Finally, the absolute and z-standardised feature weights were extracted for each successful analysis time window, averaged across all single time points within each analysis window (note that single feature weights for all 2 ms time points within analysis windows would be impossible to interpret), and assigned to each channel. This resulted in one value for each of the 64 channels, representing a rough estimate of their importance for the classification. Note that these maps were not interpreted as the 'source' (or origin) of this information (for a critique of this approach see e.g., Haufe et al., 2014), but merely obtained to compare similarities of successful classification between rating dimensions.

2.6 Control Analyses

2.6.1 Low-level visual image features

To ensure that participants' ratings (and potentially their successful prediction) did not simply reflect low-level visual properties of the images, the relationships between seven low-level features of the images, provided by the NAPS, and the participants' ratings were assessed (c.f. Bode et al., 2014a). The low-level features were: a) luminance, defined as the average pixel value for the images, b) contrast, defined as the standard deviation across pixels for the images, c) JPEG size and entropy, both used as indices of image complexity, and e) the three dimensions of the CIE L*a*b* colour space referred to as LABL, LABA, and LABB (Marchewka et al., 2014), which indexed the colour composition of the images. The relationships between these low-level parameters and the participants' ratings were investigated by correlating the parameter values and the mean rating values for each image for each dimension. Since participants' ratings were not normally distributed, non-parametric Spearman's correlations were used in all analyses.

2.6.2 Normative ratings

Additionally, the relationships between participants' ratings and the valence, approach/avoidance and arousal normative ratings provided by the NAPS were assessed. Note that given the large number of images used in this study, it was not possible to obtain individual ratings on these dimensions, but the use of the normative ratings allowed for an approximated assessment of these relationships. Spearman's correlations were again used in all analyses.

3. Results

3.1 Performance on foreground task

Participants responded with high accuracy to the box filling task (M = 98.76%, SE = .67; two trials excluded in total due to technical problems with recording the button press). The average response time across participants was 435 ms (SD = 139). The high average accuracy suggests that participants performed the task correctly, and that attention was maintained across the task and was independent of the category of the background stimuli.

3.2 Image ratings

Ratings for all dimensions were relatively skewed and not normally distributed for many participants (see Supplement B). Descriptive statistics for all dimensions are provided in Table 1. For MVPA, ratings were converted into binary scores according to whether or not they fell above or below the (group level) median for the attribute being analysed (see section 2.5).

Table 1

Means, Medians and Standard Deviations (SD) for the participants' ratings

Stimulus Attribute	Mean	Median	SD
Wanting	3.12	2	2.45
Relevance	3.46	3	2.29
Familiarity	4.82	5	2.73
Aesthetic pleasantness	4.03	4	2.63

Note. A 9-point scale (1 = not at all, 9 = very much) was used for all dimensions.

3.3 Multivariate pattern analysis of ERP data

A linear SVM classifier was used to predict participants' post-experimental ratings (following conversion into binary 'high' and 'low' values) for the images from the single-trial spatio-temporal patterns of their EEG data.

Wanting ratings (N = 16 participants) could be predicted significantly above chance between 100 – 150 ms following stimulus presentation (Figure 3). Similarly, Relevance ratings (N = 19) could be predicted significantly above chance during the same time period (Figure 4). When we analysed single time windows, wanting ratings could not be predicted significantly above chance following cluster-based correction for multiple comparisons. However, Relevance ratings could be predicted significantly above chance during 90 – 130 ms following stimulus presentation. Finally, the classification accuracies for Familiarity (N = 18), Aesthetic Pleasantness (N = 20) and Time Reference (N = 21) were again not significantly different from the empirical chance distribution across any of the time windows (Figure 5).

Additionally, we repeated the classification analyses using different analysis window widths (10 ms, 40 ms), different step sizes (20 ms, 40 ms), and alternative approaches to converting ratings into binary variables (e.g., above and below scale midpoint, or excluding mid-range ratings) but these analyses did not improve classification accuracy and are not reported here.

3.4 Relationship between Wanting and Relevance ratings

One possible interpretation of the high similarity of classification time-courses for Wanting and Relevance was that these two concepts were strongly inter-related for participants. To test this, we analysed the Spearman's correlations between average Wanting ratings and Relevance ratings. For 15 of the 50 images, these correlations were significant after Bonferroni-correction (median r = .53, range: .01 to. 88; see Supplement C for full results). When calculated separately for each participant, the correlations for Wanting and Relevance ratings across images for 20 of the 25 participants showed significant positive correlations after Bonferroni-correction (median r = .54, range: .15 to .88; see Supplement D

for full results). These results provided evidence that participants' assessment of the Relevance and general desirability of an object were inter-related.

Given these results, we additionally visually assessed the similarity between the z-standardised absolute feature weight maps for Wanting and Relevance classification, averaged across 100-150 ms following stimulus onset (see Figure 6). While these maps cannot be used to determine with certainty the sources of predictive information (Haufe et al., 2014), their high similarly nevertheless suggests at least similar neural generators for the predictive signals for both dimensions.

3.5 Control analyses

3.5.1 Low-level visual image features

No significant correlations were found between any of the low-level image features and the average Wanting and Relevance ratings across images (Table 2). Furthermore, no significant correlations were found between individual participant ratings and any of the low-level image features (see Supplement E for full results). These results suggest that ratings (and the classification thereof) were not driven by simple visual properties of the images.

Table 2

Correlations between mean Wanting and Relevance ratings and the seven low-level stimulus features

Low-level feature	Wanting		Relevance	
	r_s	p	r_s	p
Luminance	.05	.75	.11	.46
Contrast	-.15	.30	-.23	.11
JPEG size	-.09	.55	-.13	.37
LABL	.06	.69	.13	.37
LABA	.10	.51	.12	.42
LABB	-.04	.77	.02	.91
Entropy	.05	.72	.13	.37

Note. r_s = Spearman's rho; p = uncorrected significance level

3.5.2 Normative ratings

Next, we assessed the relationship between the rating dimensions, which could successfully be predicted from patterns of ERPs, and the normative image ratings provided in the NAPS. Using Spearman correlation analyses, it was found that average Wanting ratings were significantly positively correlated with valence ($r = .84, p < .001$) and approach/avoidance ($r = .83, p < .001$), and negatively correlated with arousal ($r = -.59, p < .001$). Similarly, Relevance ratings were significantly positively correlated with valence ($r = .67, p < .001$) and approach/avoidance ($r = .58, p < .001$) and negatively correlated with arousal ($r = -.52, p < .001$).

Correlations between individual participants' ratings and valence, approach/avoidance and arousal normative ratings were also assessed (all Bonferroni-corrected for multiple comparisons; see Supplement F). For the Wanting ratings, 20 participants showed significant

positive correlations with valence ratings, 18 participants showed significant positive correlations with approach/avoidance ratings, and 9 participants showed significant negative correlations with arousal normative ratings. Similarly, for the Relevance ratings, 10 participants showed significant positive correlations with valence normative ratings, 7 participants showed significant positive correlations with approach/avoidance, and 3 participants showed significant negative correlations with arousal normative ratings.

These results provide evidence of substantial associations between Wanting and valence, approach/avoidance and arousal. The same was found for Relevance, although it is noted that these associations were less pronounced.

4. Discussion

Using MVPA for ERPs, we investigated whether information relating to the general desirability of task-irrelevant, 'everyday' object stimuli was automatically processed in the hundreds of milliseconds following their background presentation. It was found that 'Wanting' and 'Relevance' ratings could be predicted significantly above chance from distributed patterns of brain activity between 100-150 ms following stimulus presentation. 'Familiarity', 'Aesthetic Pleasantness', and 'Time Reference' ratings could not be predicted. Following cluster-based correction for multiple comparisons 'Relevance' ratings could be predicted significantly above chance in consecutive single time windows, yet 'Wanting' ratings could not. Further, Wanting and Relevance ratings were highly correlated and displayed highly similar feature weight maps. Ratings for these two dimensions also showed substantial associations with all three normative dimensions of the images: valence, approach/avoidance and arousal.

4.1 Fast and automatic processing of subjective stimulus Relevance

The fact that participants' Relevance ratings could robustly be predicted from distributed patterns of their brain activity during an early stage of stimulus processing suggests that information relating to the subjective relevance of stimuli was rapidly and automatically processed, even more so than Wanting. While we have no direct evidence for the claim that this semantic analysis was truly automatic and pre-conscious, several aspects of the study design and results support such an interpretation. First, participants were unaware that they would have to rate the stimuli with regards to Relevance and Wanting until after the experimental task (i.e. when their brain activity had already been recorded). Hence, they were not primed in any way to appraise the stimuli with regards to this dimension. Second, prediction of Relevance (and to some extent Wanting) ratings was possible around 100 ms after stimulus onset, supporting the idea that this reflected an early processing stage, preceding potentially more in-depth semantic analyses dependent on feedback-loops. It is therefore unlikely that the processing of such information resulted from conscious deliberation, but could instead reflect the fast extraction of the semantic 'gist' of the visual environment (Oppermann et al., 2011). Finally, participants responded with high accuracy to the foreground attention task. Whilst this cannot unequivocally establish that attention remained centrally fixated throughout the task, especially since the attention task only required responses on 30% of trials and was less demanding than in previous studies (e.g., Bode et al., 2014a; Tusche et al., 2010), it nevertheless suggests that the processing of predictive information did not result from conscious deliberation on the behalf of the participants. This is further supported by the fact that the images were truly task-irrelevant.

The claim that participants appear to have automatically processed object information is consistent with previous research suggesting that high-level stimulus attributes were encoded in brain activity even when participants did not consciously attend to the stimuli or were unaware of the attribute of interest (Bode et al., 2014a; Lebreton et al., 2009; Peelen et al., 2009; Simanova et al., 2010). In particular, our results extend a number of past EEG studies, which found that specific semantic stimulus attributes (Bode et al., 2014a) and category membership (Simanova et al., 2010; Taghizadeh-Sarabi et al., 2015; Wang et al., 2012) were processed in the hundreds of milliseconds following stimulus presentation. Our results are also consistent with previous studies showing that information relating to individual preferences for specific stimuli (consumer products) was automatically processed during passive stimulus exposure (e.g., Telpaz et al., 2015; Tusche et al., 2010). Importantly however, our results suggest that the processing of such information may not occur for all stimuli. Our strong prediction of Relevance ratings in particular may indicate that stimuli/objects initially undergo automatic processing with regards to their subjective relevance to the decision maker and, if judged to be relevant, may then undergo more in-depth processing. This view is consistent with previous studies which used stimuli of high relevance to participants and found that information predictive of individual preferences could reliably be decoded (e.g., Tusche et al., 2010). Crucially, if the processing of stimulus relevance indeed 'gates' in-depth semantic processing, the majority of positive findings (e.g., Bode et al., 2014; Tusche et al., 2010) showing that various semantic features of stimuli could be predicted from brain activity during passive exposure may be misleading. These studies seem to imply that a great variety of semantic aspects of stimuli are processed during passive exposure. However, this might only be true for highly relevant stimuli, or some dimensions (e.g., purchase decisions) might even simply reflect potentially correlated stimulus relevance judgements. Ultimately, the current study was not explicitly designed to test whether the processing of stimulus relevance gates the processing of more in-depth information. The processing of Wanting, which operationalised stimulus desirability in our study, may indeed be gated by subjective relevance processing. However, given the great similarity of both dimensions in terms of correlations between ratings and feature weight maps, these may both tap into the same underlying, rapid cognitive assessment.

It has been suggested that the Time Reference of stimuli is automatically appraised (Bode et al., 2014a), since, for most potentially rewarding objects, it is important to know whether an object is currently available, or will not be available until sometime in the future. Our current results seem to challenge this assumption as we could not replicate the prediction of post-experimental Time Reference ratings. However, there were marked differences between the studies that could explain this discrepancy. First, previous research used only positively valenced stimuli (such as desirable foods, positive social interactions, status symbols, romantic scenes), which were further optimized for Time Reference ratings in pretests (Bode et al., 2014a). In contrast, the current study deliberately did not optimise stimuli but used an equal amount of high, medium and low valenced stimuli. Given that Relevance ratings were found to be highly correlated with normative valence ratings in the current study, the fact that Time Reference ratings could not be predicted may also indicate that decision-makers only automatically processed time-related information if an object was judged to be generally relevant to them. This would potentially suggest a hierarchy of processing stages, which we are not able to resolve with our (rather long) 20 ms analysis time windows. Subdividing our image sample by high and low relevance to assess this question was also precluded due to the low resulting trial numbers. Another difference is that the current study used a scale ranging from 'past' to 'future', with 'present' being the implicit mid-point, while the previous study only asked participant to rate objects on a scale from

'present' to 'future', and might therefore have captured a slightly different (and more decision-relevant) Time Reference concept.

4.2 Interpreting the prediction of ratings

One question which arises when considering the results of this study and previous decoding studies is: what is the exact nature of the predictive information which is being decoded? It is interesting to note that, with pattern classification studies in general, interpretation of the predictive information relies heavily on the variable being predicted. For example, in the current study, when decoding Wanting ratings, it is tempting to think the predictive information must directly relate to the processing of stimulus desirability, yet when decoding Relevance ratings, it is tempting to think the predictive information must directly relate to the processing of stimulus relevance. However, it is also possible that in both analyses the ratings are being predicted from the same source of predictive information which ultimately relates to the processing of some alternative stimulus dimension(s). In the current study, the close relationships between Relevance and Wanting ratings and the normative arousal, valence and approach/avoidance ratings suggests that judgments regarding stimulus relevance and desirability may have been driven by the aggregation of information regarding dimensions such as arousal, valence and/or approach/avoidance. This is supported by the fact that such features (arousal and valence) are rapidly reflected in brain activity following stimulus presentation (e.g. Feng et al., 2012; Harmon-Jones & Allen, 1998; see Olofsson et al., 2008 for review) and are also relatively time-locked in their processing so could easily be decoded across a broad range of tasks. Thus, one interpretation of the current pattern of results is that participants rapidly and automatically processed affective stimulus attributes, which later informed their relevance and desirability judgements. Ultimately, the current results alone do not allow us to determine which factors contribute to a general relevance signal in the brain and how these are integrated. Indeed, this might depend on both the nature of the stimuli and the exact questions used to assess the constructs. Exploring these questions poses an exciting avenue for future research.

4.3 Semantic stimulus analysis and decision-making

Critically, it could be argued that the participants' Relevance and Wanting ratings only reflected low-level visual image features. If this were true, there would be no evidence that stimuli were automatically processed at a semantic level. However, our control analyses used a variety of low-level features, such as luminance, colour space and visual complexity, which are known to be processed during the same time period from which participants' ratings could be predicted (e.g., Johannes et al., 1995; Tarkiainen et al., 2002), and we found no evidence for any associations. This replicated previous control analyses (Bode et al., 2014a), which also found no such associations.

Note, however, that our results did not directly show that the either the subjective relevance or desirability judgments were used in decision-making. Notably, our participants were neither asked to choose between objects (e.g., Tusche et al., 2010), nor did we test for any spill-over effects of semantic attribute processing to related decision processes, as in previous research (e.g., Wilson & Daly, 2004; Murawski et al., 2012). Future studies are therefore needed to establish whether, given the possibility of actually acquiring an object, the Wanting statements predict real decision behaviour. Given the strong correlation of Wanting with other dimensions – in particular, valence and approach/avoidance – it is further possible that these related concepts might capture decision-relevant general desirability better than the definition of Wanting used here. Future studies might also aim to assess more

directly whether more detailed constructs such as preferences (Lebreton et al., 2009), approach motivation (Harmon-Jones & Allen, 1998), or utility (Fishburn, 1970) might describe the early information encoded in brain activity patterns better than our rating dimension. Given that incorporating more ratings into the current study would have excessively increased the length of the study, as well as demands on the participant, our results cannot answer these conceptual questions.

4.4 Limits to the prediction of semantic information

It is tempting to assume that because Familiarity and Aesthetic Pleasantness ratings could not be predicted, this implies that information relating to these stimulus features was not automatically processed. This conclusion is, however, inconsistent with previous research (Jacobsen & Höfel, 2003; Kühn & Gallinat, 2012; Peissig et al., 2007; Scott et al., 2006; Wang et al., 1994). It should be noted that the absence of significant results in MVPA analyses does not provide evidence for the absence of a neural representation. Another explanation for this null-effect could be that the coding of these attributes might involve deeper brain structures (e.g., Brown et al., 2011; Kühn & Gallinat, 2012), which are less likely to generate reliably detectable patterns of activity in the EEG (e.g., Koike et al., 2015; Srinivasan, 1999). Alternatively, these dimensions might have been less important for participants and therefore received less attention, or they might not have been consistently analysed for all objects, or for all participants, or semantic analyses might have occurred with less temporal consistency, attenuating the temporal patterns. Finally, as noted above, it may be the case that the processing of stimulus Relevance 'gates' the processing of more specific stimulus information, such as Familiarity, Aesthetic Pleasantness and Time Reference. All these possibilities would lead to weaker or less time-locked and more variable neural signatures, which cannot currently be detected by MVPA, despite its high sensitivity. Future studies could test these questions by again designing optimised stimulus sets, and investigating systematically under which circumstances neural signatures of these dimensions occur. Our 'everyday' object stimuli might simply have been too varied to warrant a semantic assessment of these dimensions on each trial.

This study was also somewhat limited in its ability to successfully predict ratings in general as many participants did not display strong variance on several dimensions. This prohibited the use of SVR as an approach to regress ratings directly from brain activity patterns (Bode et al., 2014a). This lack of variance in responses again might have been due to the fact that the current stimuli were not optimised for any of the attributes of interest. Alternatively, it might have been due to the use of ordinal nine-point scales. In previous food-related research investigating explicit Wanting, measures such as these have been criticised as being restricted in accuracy and limited by methodological problems such as 'end avoidance' (Finlayson & Dalton, 2012). Future research could investigate whether continuous scales, such as the bipolar semantic slider scales used to record the NAPS normative ratings (Marchewka et al., 2014), are better suited at capturing fine-grained differences in ratings. This might also allow the use of individual median splits to define categories for classification, which was precluded here due to the limited use of the scale in some participants. Further, while we could not explicitly separate between participants with high and low biases to respond "not at all", as this led to insufficient statistical power, future experiments with larger sample sizes could address this question and investigate whether such groups differ in their neural signatures for these and other dimensions.

4.5 Conclusion

To our knowledge, the current study is the first to provide evidence that the subjective relevance of task-irrelevant, everyday object stimuli can be predicted from brain activity patterns recorded with EEG, around one hundred milliseconds following stimulus presentation. These results suggest that a rapid, automatic semantic analysis occurs immediately following stimulus presentation during which decision-related information is extracted. The current results further suggest that subjective relevance and desirability judgments may result from the fast and automatic aggregation of information regarding stimulus valence, approach/avoidance and arousal. Disentangling which stimulus dimensions are automatically extracted for which kind of stimuli poses an interesting challenge for future research.

5. Acknowledgements

6. References

Aarts, H., Custers, R., & Holland, R. W. (2007). The nonconscious cessation of goal pursuit: when goals and negative affect are coactivated. *Journal of personality and social psychology*, *92*(2), 165.

Bai, O., Lin, P., Vorbach, S., Li, J., Furlani, S., & Hallett, M. (2007). Exploration of computational methods for classification of movement intention during human voluntary movement from single trial EEG. *Clinical Neurophysiology*, *118*(12), 2637-2655.

Benoit, R. G., Gilbert, S. J., & Burgess, P. W. (2011). A neural mechanism mediating the impact of episodic prospection on farsighted decisions. *The Journal of Neuroscience*, *31*(18), 6771-6779.

Berridge, K. C., Robinson, T. E., & Aldridge, J. W. (2009). Dissecting components of reward: 'liking', 'wanting', and learning. *Current opinion in pharmacology*, *9*(1), 65-73.

Blankertz, B., Lemm, S., Treder, M., Haufe, S., & Müller, K. R. (2011). Single-trial analysis and classification of ERP components—a tutorial. *NeuroImage*, *56*(2), 814-825.

Bode, S., Bennett, D., Stahl, J., & Murawski, C. (2014a). Distributed patterns of event-related potentials predict subsequent ratings of abstract stimulus attributes. *PloS one*, *9*(10), e109070.

Bode, S., Murawski, C., Soon, C.S., Bode, P., Stahl, J., & Smith, P.L. (2014b). Demystifying "free will": The role of contextual information and evidence accumulation for predictive brain activity. *Neuroscience & Biobehavioral Reviews*, *47*, 636-645.

Bode, S., Sewell, D. K., Lilburn, S., Forte, J. D., Smith, P. L., & Stahl, J. (2012). Predicting perceptual decision biases from early brain activity. *The Journal of Neuroscience*, *32*(36), 12488-12498.

Bode, S., & Stahl, J. (2014). Predicting errors from patterns of event-related potentials preceding an overt response. *Biological psychology*, *103*, 357-369.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial vision*, *10*, 433-436.

Brown, S., Gao, X., Tisdelle, L., Eickhoff, S. B., & Liotti, M. (2011). Naturalizing aesthetics: brain areas for aesthetic appraisal across sensory modalities. *Neuroimage*, *58*(1), 250-258.

Bullmore, E. T., Suckling, J., Overmeyer, S., Rabe-Hesketh, S., Taylor, E., & Brammer, M. J. (1999). Global, voxel, and cluster tests, by theory and permutation, for a difference between two groups of structural MR images of the brain. *IEEE transactions on medical imaging*, *18*(1), 32-42.

Chang, C. C., & Lin, C. J. (2011). LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, *2*(3), 27.

Cheng, Y.Y., Shein, P.P., & Chiou, W.B. (2012). Escaping the impulse to immediate gratification: the prospect concept promotes a future-oriented mindset, prompting an inclination towards delayed gratification. *British Journal of Psychology*, *103*, 129-141.

Combrisson, E., & Jerbi, K. (2015). Exceeding chance level by chance: The caveat of theoretical chance levels in brain signal classification and statistical assessment of decoding accuracy. *Journal of neuroscience methods*, *250*, 126-136.

Creswell, J. D., Bursley, J. K., & Satpute, A. B. (2013). Neural reactivation links unconscious thought to decision making performance. *Social Cognitive and Affective Neuroscience*, *8*(8), 863-869.

Custers, R., & Aaarts, H. (2010). The unconscious will: How the pursuit of goals operates outside of conscious awareness. *Science*, *329*, 47-50.

Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods, 134*(1), 9-21.

Dijksterhuis, A., & Aarts, H. (2010). Goals, attention, and (un) consciousness. *Annual Review of Psychology*, *61*, 467-490.

Dijksterhuis, A., Bos, M. W., Nordgren, L.F., & Van Baaren, R.B. (2006a). On making the right choice: The deliberation-without-attention effect. *Science*, *311*(5763), 1005-1007.

Dijksterhuis, A., & Nordgren, L.F. (2006b). A theory of unconscious thought. *Perspectives on Psychological Science*, *1*(2), 95-109.

Felsen, G., & Dan, Y. (2005). A natural approach to studying vision. *Nature Neuroscience*, *8*(12), 1643-1646.

Feng, C., Wang, L., Liu, C., Zhu, X., Dai, R., Mai, X., & Luo, Y. J. (2012). The time course of the influence of valence and arousal on the implicit processing of affective pictures. *PloS one*, *7*(1), e29668.

Finlayson, G., & Dalton, M. (2012). Hedonics of food consumption: are food 'liking' and 'wanting' viable targets for appetite control in the obese?. *Current Obesity Reports*, *1*(1), 42-49.

Finlayson, G., King, N., & Blundell, J. (2008). The role of implicit wanting in relation to explicit liking and wanting for food: implications for appetite control. *Appetite*, *50*(1), 120-127.

Fishburn, P. C. (1970). Utility theory for decision making (No. RAC-R-105). Research analysis corp McLean VA.

Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2016). Decoding dynamic brain patterns from evoked responses: a tutorial on multivariate pattern analysis applied to time series neuroimaging data. *Journal of Cognitive Neuroscience*, *29*(4), 677-697.

Harmon-Jones, E., & Allen, J. J. (1998). Anger and frontal brain activity: EEG asymmetry consistent with approach motivation despite negative affective valence. *Journal of personality and social psychology*, *74*(5), 1310.

Haufe, S., Meinecke, F., Görgen, K., Dähne, S., Haynes, J. D., Blankertz, B., & Bießmann, F. (2014). On the interpretation of weight vectors of linear models in multivariate neuroimaging. *Neuroimage*, *87*, 96-11.

Jacobsen, T., & Höfel, L. (2003). Descriptive and evaluative judgment processes: behavioral and electrophysiological indices of processing symmetry and aesthetics. *Cognitive, Affective, & Behavioral Neuroscience*, *3*(4), 289-299.

Johannes, S., Münte, T. F., Heinze, H. J., & Mangun, G. R. (1995). Luminance and spatial attention effects on early visual processing. *Cognitive Brain Research*, *2*(3), 189-205.

Kayser, J., & Tenke, C. E. (2006). Principal components analysis of Laplacian waveforms as a generic method for identifying ERP generator patterns: I. Evaluation with auditory oddball tasks. *Clinical Neurophysiology*, *117*(2), 348-368.

Kim, B. K., & Zauberman, G. (2013). Can Victoria's Secret change the future? A subjective time perception account of sexual-cue effects on impatience. *Journal of Experimental Psychology: General*, *142*(2), 328.

King, J. R., & Dehaene, S. (2014). Characterizing the dynamics of mental representations: the temporal generalization method. *Trends in Cognitive Sciences*, *18*(4), 203-210.

Koike, T., Tanabe, H. C., & Sadato, N. (2015). Hyperscanning neuroimaging technique to reveal the "two-in-one" system in social interactions. *Neuroscience Research*, *90*, 25-32.

Kühn, S., & Gallinat, J. (2012). The neural correlates of subjective pleasantness. *Neuroimage*, *61*(1), 289-294.

Lebreton, M., Jorge, S., Michel, V., Thirion, B., & Pessiglione, M. (2009). An automatic valuation system in the human brain: evidence from functional neuroimaging. *Neuron*, *64*(3), 431-439.

Levy, I., Lazzaro, S. C., Rutledge, R. B., & Glimcher, P. W. (2011). Choice from non-choice: predicting consumer preferences from blood oxygenation level-dependent signals obtained during passive viewing. *The Journal of Neuroscience*, *31*(1), 118-125.

Marchewka, A., Żurawski, Ł., Jednoróg, K., & Grabowska, A. (2014). The Nencki Affective Picture System (NAPS): Introduction to a novel, standardized, wide-range, high-quality, realistic picture database. *Behavior Research Methods*, *46*(2), 596-610.

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG-and MEG-data. *Journal of Neuroscience Methods*, *164*(1), 177-190.

Mogg, K., Bradley, B. P., Field, M., & De Houwer, J. (2003). Eye movements to smoking-related pictures in smokers: relationship between attentional biases and implicit and explicit measures of stimulus valence. *Addiction*, *98*(6), 825-836.

Murawski, C., Harris, P. G., Bode, S., & Egan, G. F. (2012). Led into temptation? Rewarding brand logos bias the neural encoding of incidental economic decisions. *PloS one*, *7*(3), e34155.

O'Doherty, J., Winston, J., Critchley, H., Perrett, D., Burt, D. M., & Dolan, R. J. (2003). Beauty in a smile: the role of medial orbitofrontal cortex in facial attractiveness. *Neuropsychologia*, *41*(2), 147-155.

Olofsson, J. K., Nordin, S., Sequeira, H., & Polich, J. (2008). Affective picture processing: an integrative review of ERP findings. *Biological Psychology*, *77*(3), 247-265.

Oppermann F, Hassler U, Jeschniak JD, Gruber T (2011). The rapid extraction of gist–Early neural correlates of high-level visual processing. *Journal of Cognitive Neuroscience*, *24*(2), 521–529.

Peelen, M. V., Fei-Fei, L., & Kastner, S. (2009). Neural mechanisms of rapid natural scene categorization in human visual cortex. *Nature, 460*(7251), 94-97.

Peissig, J. J., Singer, J., Kawasaki, K., & Sheinberg, D. L. (2007). Effects of long-term object familiarity on event-related potentials in the monkey. *Cerebral Cortex*, *17*(6), 1323-1334.

Perrin, F., Pernier, J., Bertrand, O., & Echallier, J. F. (1989). Spherical splines for scalp potential and current density mapping. *Electroencephalography and clinical neurophysiology*, *72*(2), 184-187.

Perrin, F., Pernier, J., Bertnard, O., Giard, M. H., & Echallier, J. F. (1987). Mapping of scalp potentials by surface spline interpolation. *Electroencephalography and Clinical Neurophysiology*, *66*(1), 75-81.

Peters, J., & Büchel, C. (2010). Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal-mediotemporal interactions. *Neuron*, *66*(1), 138-148.

Pettersen, K. H., Devor, A., Ulbert, I., Dale, A. M., & Einevoll, G. T. (2006). Current-source density estimation based on inversion of electrostatic forward solution: effects of finite extent of neuronal activity and conductivity discontinuities. *Journal of Neuroscience Methods*, *154*(1), 116-133.

Philiastides, M. G., & Sajda, P. (2006). Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cerebral Cortex*, *16*(4), 509-518.

Sander, D., Grafman, J., & Zalla, T. (2003). The human amygdala: an evolved system for relevance detection. *Reviews in the Neurosciences*, *14*(4), 303-316.

Scott, L. S., Tanaka, J. W., Sheinberg, D. L., & Curran, T. (2006). A reevaluation of the electrophysiological correlates of expert object processing. *Journal of Cognitive Neuroscience*, *18*(9), 1453-1465.

Simanova, I., Van Gerven, M., Oostenveld, R., & Hagoort, P. (2010). Identifying object categories from event-related EEG: toward decoding of conceptual representations. *PloS one*, *5*(12), e14465.

Simmank, J., Murawski, C., Bode, S., & Horstmann, A. (2015). Incidental rewarding cues influence economic decisions in people with obesity. *Frontiers in behavioral neuroscience*, *9*:278.

Srinivasan, R. (1999). Methods to improve the spatial resolution of EEG. International *Journal of Bioelectromagnetism*, *1*(1), 102-111.

Taghizadeh-Sarabi, M., Daliri, M. R., & Niksirat, K. S. (2015). Decoding objects of basic categories from electroencephalographic signals using wavelet transform and support vector machines. *Brain topography*, *28*(1), 33-46.

Tarkiainen, A., Cornelissen, P. L., & Salmelin, R. (2002). Dynamics of visual feature analysis and object-level processing in face versus letter-string perception. *Brain*, *125*(5), 1125-1136.

Telpaz, A., Webb, R., & Levy, D. J. (2015). Using EEG to Predict Consumers' Future Choices. *Journal of Marketing Research*, *52*(4), 511-529.

Tusche, A., Bode, S., & Haynes, J. D. (2010). Neural responses to unattended products predict later consumer choices. *The Journal of Neuroscience*, *30*(23), 8024-8031.

Tusche, A., Kahnt, T., Wisniewski, D., & Haynes, J. D. (2013). Automatic processing of political preferences in the human brain. *Neuroimage*, *72*, 174-182.

Van den Bergh, B., Dewitte, S., & Warlop, L. (2008). Bikinis instigate generalized impatience in intertemporal choice. *Journal of Consumer Research*, *35*(1), 85-97.

Vanrullen, R. (2011). Four common conceptual fallacies in mapping the time course of recognition. *Perception Science*, *2*, 365.

Wang, Q., Cavanagh, P., & Green, M. (1994). Familiarity and pop-out in visual search. *Perception & Psychophysics*, *56*(5), 495-500.

Wang, C., Xiong, S., Hu, X., Yao, L., & Zhang, J. (2012). Combining features from ERP components in single-trial EEG for discriminating four-category visual objects. *Journal of Neural Engineering*, *9*(5), 056013.

Wilson, M., & Daly, M. (2004). Do pretty women inspire men to discount the future?. *Proceedings of the Royal Society of London B: Biological Sciences*, *271*(Suppl 4), S177-S179.

Zaki, J., & Ochsner, K. (2009). The need for a cognitive neuroscience of naturalistic social cognition. *Annals of the New York Academy of Sciences*, *1167*(1), 16-30.

Zedelius, C. M., Veling, H., Custers, R., Bijleveld, E., Chiew, K. S., & Aarts, H. (2014). A new perspective on human reward research: How consciously and unconsciously perceived reward information influences performance. *Cognitive, Affective, & Behavioral Neuroscience*, *14*(2), 493-508.

Zhong, C. B., & DeVoe, S. E. (2010). You Are How You Eat Fast Food and Impatience. *Psychological Science, 21*(5), 619-622.

7. Figure captions

Figure 1. Illustration of the paradigm. A jittered fixation period of 2-3 s preceded image presentation. During this period a grey screen with a central fixation box was presented. Following the fixation period each image was presented for 3.2 s. On 30% of trials the central fixation box would 'fill' into a solid black square, between 1-2 s following stimulus onset. Participants were tasked with monitoring the fixation box and pressing the space bar whenever it 'filled'.

Figure 2. Multivariate SVM classification. Spatio-temporal vectors were created from 20 ms time windows of data. The vectors were labelled 'high' or 'low' depending on whether the image presented had been given a 'high' or 'low' Wanting rating (same for Relevance, Familiarity, Aesthetic Pleasantness and Time Reference in the case of those analyses). A linear SVM classifier was trained to distinguish between 'high' and 'low' vectors. The ability of the classifier to predict whether novel data were associated with 'high' or 'low' ratings was then tested. After the full classification procedure the 20 ms analysis window would step forward 10 ms in time and the procedure would be repeated.

Figure 3. Spatio-temporal decoding of Wanting ratings. Linear SVM classification was used to predict Wanting ratings from distributed patterns of CSD ERPs. The black line represents the classification accuracy. The blue line represents the empirical chance distribution. The red bar denotes a statistically significant difference between classification accuracy and the empirical chance distribution in the 50 ms time window spanning 100-150 ms.

Figure 4. Spatio-temporal decoding of Relevance ratings. Linear SVM classification was used to predict Relevance ratings from distributed patterns of CSD ERPs. The black line represents the classification accuracy. The blue line represents the empirical chance distribution. The red bar denotes a statistically significant difference between classification accuracy and the empirical chance distribution in the 50 ms time window spanning 100-150 ms. The red data points denote single time points where classification accuracy was significantly different from the empirical chance distribution following cluster-based correction for multiple comparisons.

Figure 5. Spatio-temporal decoding of A) Aesthetic Pleasantness, B) Familiarity, and C) Time Reference ratings. Linear SVM classification was used to predict these ratings from distributed patterns of CSD ERPs. The black line represents the classification accuracy. The blue line represents the empirical chance distribution. No single time points showed significant decoding following cluster-based correction for multiple comparisons. Note: the final 50 ms time window for Time Reference approached significance.

Figure 6. Scalp maps of the z – standardised absolute feature weights, averaged across 100-150ms for A) Relevance ratings and B) Wanting ratings.