Web Appendix

From “Where” to “What”:
Distributed Representations of Brand Associations in the Human Brain

1. Methods

1.1 General Methods

1.1.1 Subjects

Seventeen adults (6 females) from the San Francisco Bay Area were recruited from Craigslist to participate in the functional magnetic resonance imaging (fMRI) study. Their ages ranged from 26 to 45, with an average of 34 and the standard deviation of 6.5. Informed consent was obtained using a consent form approved by the Internal Review Board at University of California, Berkeley. The total time for the whole experiment was approximately 3 hours, including the instruction, the scanning session, and the post-experiment questionnaires. Each participant was paid $70 in cash upon completion of the experiment.

1.1.2 Experimental Protocol

We imaged participants' brains with functional magnetic resonance imaging (fMRI) while they underwent a passive viewing task involving 44 well-known brands. The stimuli were 44 brand logos selected from Best Global Brands by Interbrand (Figure 1A), with significant diversity in brand meaning and industry. Each of the 44 stimulus items was presented four times in a pseudo-random sequence on the gray background (Figure 1B), and each presentation lasted for 4-8s. There were twelve rest periods distributed across the session lasting for 4-8 seconds, during which participants were instructed to fixate on a cross at the center of the screen. In addition, there were six self-paced catch questions, in which participants had to press a particular button using an MRI-compatible button box to continue to the next trial (Figure S1).

When a brand logo was presented, the participants' task was to think about the characteristics or traits associated with the brand. Each participant was free to think about any characteristic or trait they associated with the brands, and there was no attempt to obtain consistency of the associations neither across participants nor across repetition times.

After the scanning session, participants were asked to complete a survey about the brands they saw in the scanner. The survey included the 42-item brand personality scale (Aaker 1997) (Table S1), familiarity, and preference for each brand. The brand personality scale involved judgment of the descriptiveness of 42 traits to each brand, with a five-point scale from not at all descriptive (rating=1) to extremely descriptive (rating=5). Participants reported their familiarity and preference toward the brands from a four-point scale, ranging from dislike/unfamiliar (rating=1), somewhat dislike/unfamiliar (rating=2), somewhat like/familiar (rating=3) to like/familiar (rating=4). We obtained 1,936 ratings in total per participant in the survey.

1.1.3 fMRI Data Acquisitions

Functional images were acquired on a Siemens 3T TIM/Trio scanner at Henry H. Wheeler Jr. Brain Imaging Center at University of California, Berkeley. An EPI sequence was used to acquire the functional data (repetition time (TR) = 2,000 ms; echo time (TE) = 30 ms; voxel resolution = 3mm×3mm×3mm; FOV read = 192 mm; FOV phase = 100%; interleaved series order). The scan sequences were axial slices approximately flipped 30 degrees to the AC-PC axis. High-resolution structural T1-weighted scans (1mm×1mm×1mm) were acquired by using an MPRage sequence.

1.1.4 Survey of an Additional Sample

An additional sample of 25 undergraduate students completed the survey online on the same set of brands and traits of the brand personality scale for course credits. The average responses of the trait scores were highly correlated among our neuroimaging subjects and the follow-up participants (Figure 1E).

1.2 Behavioral Data analysis

To conceptualize the brands, we first characterized the set of psychological features associated with the brands using participant responses in the survey outlined in Aaker's framework (Aaker 1997). Specifically, we used a factor analytic approach to summarize variation in trait ratings and reduce collinearity issues. First, we averaged responses from all of the participants to calculate the average descriptiveness of each trait to each brand. Then, the average scores were factor-analyzed with SPSS ("IBM SPSS Statistics for Windows, Version 20.0" 2011) using principal components analysis and a varimax rotation. We selected the factors with eigenvalues greater than one, and the solution explained a high level of variance. Finally, each brand can be re-expressed in terms of its feature vector, defined as the strength of association between the brand and the factors (the personality features).

1.3 fMRI Data Preprocessing

Prior to analysis, the images were corrected for slice time artifacts, realigned, coregistered to the subject's T1 image, and normalized to Montreal Neurological Institute coordinates, using SPM8 (Statistical Parametric Mapping, Wellcome Trust Centre for Neuroimaging). Consistent with previous decoding studies, we did not smooth the images.

To identify the representative fMRI image of a brand for the analysis, we followed the methods proposed by Mumford et al. (2012) for event-related designs (Mumford et al. 2012). We first used a general linear model in SPM8 to estimate a single fMRI image for each of the 176 brand presentations, specifically method LS-S in Mumford et al. (2012). The model included the brand of interest as an individual regressor and another regressor consisting of all the other brands and the catch questions. The duration of all of the events was set to be zero. The beta values estimated for the first regressor of the brand of interest was used as the brain activation patterns associated with a brand at a particular repetition time. Alternative procedures to estimate the representative fMRI images were used (section 2.5), but consistent with Mumford et al. (2012), the LS-S method with each event modeled as an impulse function yields the best performance.

Using brain images for each brand at each repetition time, we standardized the activation levels for each voxel by z-scoring over the 176 files. Then, for each brand, we averaged the four brain images of the four repetition times to obtain the averaged fMRI image associated with thinking about the brand. The fMRI images were then exported to Matlab using Princeton MVPA toolbox. Finally, we applied the individual grey matter mask to include voxels within the grey matter. For each participant, the grey matter mask was created by segmenting the individual's normalized average EPI image using SPM8.

1.4 Decoding Analysis

1.4.1 Overview

We used model-based decoding analysis to predict the fMRI image of a brand j for each of the participant using his/her neural responses to other brands and the personality features of this brand j, following the methods proposed by Mitchell et al. (2008). We assumed that for each participant, the neural response y_j^v in voxel v to brand j is given by $y_j^v = c_1^v f_{1,j} + c_2^v f_{2,j} + \cdots + c_n^v f_{n,j}$, where $f_{n,j}$ is the value of the nth personality feature for brand j, n is the number of personality features, and c_n^v is a scalar parameter that specifies the degree to which the nth personality feature activates voxel v. Notice that the model was estimated independently for each of the participants.

1.4.2 Training the Model

The personality features associated with a brands $f_{n,j}$ were specified with the factors (section 1.2) that quantitatively capture the characteristics or traits associated with the brands. Then, the parameters c_n^v that define the neural signature contributed by the nth personality feature to the vth voxel were estimated. This is accomplished by training the model using a set of observed fMRI images associated with known stimulus brands. Each training stimulus, brand j, was first re-expressed in terms of its personality feature vector $< f_{1,j} \ldots f_{n,j} >$, and multiple regression is then used to obtain maximum likelihood estimates of the c_n^v values; that is, the set of c_n^v values that minimize the sum of squared errors in reconstructing the training fMRI images. Since the number of personality features is less than the number of training examples in our model, this multiple regression problem is well posed and a unique solution is obtained.

Once trained, the resulting computational model can be used to predict the full fMRI activation image for any other brand. Given an arbitrary new brand k, we first expressed the brand with the personality features $< f_{1,k} \ldots f_{n,k} >$ in Aaker's framework. Then, we applied the above formula using the previously estimated values for the parameters c_n^v. The computational model and corresponding theory can be directly evaluated by comparing their predictions for brands outside the training set to observed fMRI images associated with those brands.

1.4.3 Training and Evaluating the Model

The computational model was trained and evaluated using a cross validation approach, in which the model was repeatedly trained using only 42 of the 44 available stimulus items, then tested using the two stimulus items that had been left out. On each iteration, the trained model was tested by giving it the two stimulus brands it had not yet seen (b_1 and b_2), plus their observed fMRI images (i_1 and i_2). We used two ways to evaluate the performance of the model. First, we compared the similarity between the predicted and the observed brain images (section 1.4.4). Second, we required the model to predict which of the two brain images was associated with which of the two brands using a matching procedure described in section 1.4.5. This leave-two-out train-test procedure was iterated 946 times, leaving out each of the possible brand pairs.

1.4.4 Similarity between Predicted and Actual Images

Given a trained computational model, two new brands (b_1 and b_2) and two new images (i_1 and i_2), the trained model was first used to create predicted image p_1 for brand b_1 and predicted image p_2 for brand b_2. The model was evaluated by comparing these predicted fMRI images to the observed fMRI data. We first compared two possibilities: matched pairs (p_1=i_1 and p_2=i_2, dark arrows in Figure 2D) and mismatched pairs (p_1=i_2 and p_2=i_1, light arrows in Figure 2D). Under the null hypothesis of no association, the predicted fMRI image for a brand will be equally predictive of the matched brand as with the mismatched brand.

Because we do not expect every voxel in the brain to be involved in representing the personality features of the brands, only a subset of voxels was used for assessing the similarity between images. This subset of voxels was selected automatically during training, using only the data for the 42 training brands, and excluding the data from the two testing brands. The voxel selection method is described in section 1.4.6.

Let sel(i) be the vector of values of the selected subset of voxels for image i. The similarity score between a predicted image, p, and observed image, i, was calculated as the Pearson correlation coefficient of the vectors sel(p) and sel(i). In Figure S5, we compared the average similarity score of matched pairs ($p_1=i_1$ and $p_2=i_2$) to the average similarity score of mismatched pairs ($p_1=i_2$ and $p_2=i_1$). In Figure 3B, we compared, for each brand pair, the similarity score of predicted and observed brand images (p_A and i_B) against psychological similarity of brands as measured using Pearson correlation coefficient of trait ratings of the two brands.

1.4.5 Matching Predicted to Actual Images

Given the two testing brands (b_1 and b_2) and two observed images (i_1 and i_2), we then required the model to predict which of the two brain images was associated with which of the two brands. The trained model was first used to create predicted image p_1 for brand b_1 and predicted image p_2 for brand b_2. It then decided which was a better match: ($p_1=i_1$ and $p_2=i_2$) or ($p_1=i_2$ and $p_2=i_1$), by choosing the image pairing with the larger sum of similarity scores. Similarly, we selected a subset of voxels sel(i) of image i to calculate the similarity score. More specifically, the match of the pair of brands is calculated by:

$$\text{Match } (p_1=i_2 \text{ and } p_2=i_1) = \text{Correlation } (\text{sel}(p_1), \text{sel}(i_2)) + \text{Correlation } (\text{sel}(p_2), \text{sel}(i_1)).$$

The expected accuracy in matching the two left-out brands to their left-out fMRI images is 0.50 if the matching is performed at chance levels.

Pearson correlation coefficient was the first similarity measure we considered, but we subsequently also considered the Euclidean distance of the two vectors and found that the two yielded similar results (Figure S8). All results reported in the current paper use Pearson correlation coefficient.

1.4.6 Voxel Selection

As described above, similarity between two images was calculated using only a subset of the image voxels, following methods proposed in Mitchell et al. (2008). Voxels were selected automatically during training, using only the 42 training brands on each of the leave-two-out cross validation folds. To select voxels, all voxels were first assigned a "stability score" using the data from the 4 presentations of each of the 42 training stimuli. Given these 4*42 = 168 presentations (168 fMRI images), each voxel was assigned a 4×42 matrix, where the entry at row i, column j, is the value of this voxel during the i^{th} presentation of the j^{th} brand. The stability score for this voxel was then computed as the average pairwise correlation over all pairs of rows in this matrix. In essence, this assigns highest scores to voxels that exhibit a consistent (across different presentations) variation in activity across the 42 training stimuli. The 500 voxels ranked highest by this stability score were used in the similarity test in Figure 3. However, our result is robust when including more voxels in the analysis (section 2.7, Figure S9 & S10) or excluding voxels in the occipital cortex (section 2.8, Figure S11).

Selecting voxels based on the similarity score was the first voxel selection method we considered, but we subsequently also considered selecting voxels based on the significance in the multiple regression equation when training the model. The result is robust for different ways of voxel selection methods. All results reported in the current paper use the voxels selected by the stability scores.

1.4.7 Empirical Distribution to Determine Statistical Significance

The expected chance accuracy of an uninformed model correctly matching two stimuli outside the training set to their two fMRI images is 0.5. The observed accuracies of our trained models, based on 946 iterations of a leave-two-out cross validation train/test regime, are higher than 0.5. Here we used a permutation test to determine the p value based on observed accuracies, in order to reject the null hypothesis that the trained model has true accuracy of 0.5. Given our leave-two-out train/test regime, no closed-form formula is available to assign such a p value. Therefore, we computed the p value based on an empirical distribution of observed accuracies obtained from 10,200 independently trained single-participant models that we expect will have true accuracy very close to 0.5. The empirical distribution of accuracies for these null models was 0.50, with standard deviation 0.06, indicating that observed average accuracy above 0.55 for 17 participants is statistically significant at $p<0.0001$. Below we describe our approach in more detail.

We created this empirical distribution of accuracies by training multiple models using the observed fMRI images for the 44 stimulus brands, but using different brand labels. More specifically, it is a form of permutation test, permuting the 44 stimulus labels. For example, in one model, we used Google's personality features to describe Gucci, IBM's features to describe Campbell's, and so on. In another model, we used an independent scrambled set of the feature scores to describe the brands. Models were trained and tested using the leave-two-out test regime, exactly as elsewhere in this paper, with one minor exception: in these models the 500 most stable voxels were selected using data from all 44 brands, whereas elsewhere this selection of stable voxels was based only on the 42 training brands. This exception was made because it dramatically improves the computational speed.

For each of the 17 participants, we trained and tested 600 such randomly generated models, resulting in 10,200 models in total. The mean accuracy over these models was 0.50, with standard deviation 0.06.

1.4.8 Accuracy Map

The accuracy map in Figure 4 of the main text shows voxel clusters with the highest correlation between predicted and actual voxel values for an average subject. To obtain the accuracy map, we first averaged the fMRI patterns for each brand at each repetition time across 17 subjects. More specifically, we performed the second-level analysis for the subjects' estimated beta files (section 1.3) associated with each brand at each repetition time in SPM8. Using the 176 average brain images, we standardized the activation levels for each voxel by z-scoring over the 176 files. Then, for each brand, we averaged the four brain images of the four repetition times to obtain the averaged fMRI image associated with thinking about the brand. Finally, we calculated 44 predicted images for the 44 brands, by training a model on the other 43 brands to predict the remaining brand. For each voxel, this produced a set of 44 predicted values. The accuracy score of each voxel was calculated as the Pearson correlation coefficient between this vector of its predicted values and the corresponding vector of its observed values. An image map containing these voxel scores was created.

To determine the significance of the correlation between the predicted vector and the observed vector for a voxel, we took a permutation approach. We computed p values based on an empirical distribution of correlation coefficients obtained from 100,000 independently permuted values. For each voxel, we calculated the Pearson correlation coefficient of the permuted vector and the observed vector. The permuted vectors were created by scrambling the 44 predicted values.

The clusters shown in Figure 4 were then produced using SPM8, to identify clusters containing at least 10 contiguous voxels whose score was greater than the permuted threshold value $p<0.05$.

2. Supplementary Results

2.1 Factor Analysis Results

Using the fMRI subjects' average responses of the descriptiveness of the traits to the brands, our factor analysis and the criteria yielded five factors. We found that a substantial proportion of the variance (86%) was captured by these 5 factors (Table S2), and they were labeled as excitement, competence, sincerity, ruggedness, and sophistication as shown in the factor loadings of traits (Figure S2). Further inspection of the factor loadings showed that our results largely replicated those of previous studies. Using this factor analytic framework, we characterized each brand as a vector of personality features consisting of these five factor scores that summarizes the set of characteristics participants associate with these brands. Each brand thus was re-expressed in terms of its feature vector (Figure S3; Table S3).

2.2 Familiarity and Preference

Our subjects were highly familiar with the brands used in the experiment. The average familiarity score was 3.58 out of 4 across all participants and all brands.

Behaviorally, different people had different preference toward the brands. On average, Google (average preference=3.83) and Amazon.com (average preference=3.72) were the most preferred brands, while Goldman Sachs (average preference=1.89) and Marlboro (average preference=1.83) were the least preferred ones. Neurally, we found that activation in striatum was positively correlated with the subject's reported brand preference. Striatum is a region of the brain known to respond to primary and secondary rewards (Fliessbach et al. 2007; Izuma, Saito, and Sadato 2008), and is consistent with the idea that our brains respond to preferences of abstract objects such as brands.

2.3 Individual Results

Given the two testing brands (b_1 and b_2) and two observed images (i_1 and i_2), we required the computational model to predict which of the two brain images was associated with which of the two brands. The average performance of the model for iterating over all of the possible combination of hold-out data across 17 subjects is 58%, compared with 50% if the model performs at chance. For individual subjects, the average hit rate across all of the possible combinations of brand pairs is plotted in Figure S4. Subjects were sorted by the average hit rate.

2.4 Accuracy Map

The accuracy map (Figure 4) shows voxel clusters with the highest correlation between predicted and actual voxel values for an average subject. The clusters contained at least 10 contiguous voxels whose correlation value was greater than the permuted threshold value $p<0.05$. The correlation values, locations, and regions of these voxels are listed in Table 1. The surface rendering and the glass brain of locations of the most accurately predicted voxels are shown in Figure S6.

2.5 Robustness to Extracting Representative fMRI Response to the Brands

To identify the representative fMRI image of a brand, we used the procedure outlined in Mumford et al. (2012) to account for the fact that in rapid event-related designs the evoked BOLD signal for adjacent trials will overlap in time. We first used a general linear model in SPM8 to estimate a single fMRI image for each of the 176 brand presentations using method LS-S in Mumford et al. 2012. More specifically, for each subject, 176 general linear models were estimated, with each model estimating 2 regressors: one regressor for the event of interest (corresponding to a presentation of one particular brand) and one regressor for all other events that are combined into a single nuisance regressor, where each event was modeled as an impulse function (the duration of each event is set to be zero) convolved with a double gamma hemodynamic function.

We run further robustness checks using alternative methods of estimating representative fMRI images. First, we estimate the fMRI images using the LS-S model with the full duration of the events instead of setting the duration to be zero. Second, we estimate the fMRI images using the "standard" general models, where all events of interest are modeled in one general linear model per subject. Consistent with Mumford et al. 2012, the prediction rates are somewhat worse than using the duration 0 LS-S model, but remain quite significant ($p<0.005$, Figure S7A). Other features, such as modulation by psychological similarity, remain qualitatively unchanged (Figure S7B).

2.6 Robustness to the Measure of Psychological Similarity of Brands

We examined the relationship between the psychological organization of brands and the discriminability of the associated brain images in Figure 3, where the psychological similarity of brands were measured by the Pearson correlation coefficient of the averaged trait ratings of the 42 items in the survey. It is not the only way to measure the similarity of brands' psychological features. For example, instead of correlations of traits, we calculated the similarity as the Euclidean distance between brands' feature vectors of the five factor scores. Our results are robust when using the Euclidean distance between brands in association space (Figure S8). More specifically, the difference in correlation between the matched brain images and the mismatched brain images is larger when there is a larger Euclidean distance between the two brands' feature vectors (Figure S8A). Also, the strength of neural correlation between predicted and observed brain images is robustly modulated by the similarity of brands' psychological features (Figure S8B). Finally, separating the brands based on subjective similarity into quartiles as assessed based on Euclidean distance of factor scores, we find a significant relationship between hit rate and subjective similarity (Figure S8C).

2.7 Robustness to Varying Number of Voxels Used in Decoding

The main result comes from comparing the predicted and observed brain images using 500 most stable voxels (section 1.4.6) selecting from 40,000-50,000 voxels in the whole brain. As a robustness check, we ran the analysis including different numbers of voxels. More specifically, we compared our results of discriminative accuracies and the relationship between the psychological organization of brands and the correlation between predicted and observed brain images, using 500, 1000, 5000, and 25,000 voxels with the highest stability scores. Our result is robust when more voxels were included. First, the overall hit rate for hold-out classification did not drop significantly when more voxels were included in the comparison between the predicted and the observed brain images (Figure S9). Second, the significant relationship between hit rate and subjective similarity of brands was also robust when different numbers of voxels were included (Figure S9). Finally, we found that strength of neural correlation is robustly modulated by the similarity of brands' psychological features when different voxels were included (Figure S10).

2.8 Robustness to Excluding Visual Cortex Voxels

The most predictive voxels were distributed in the cortex, including occipital lobes (Figure 4 and Figure S6). As a robustness check, we ran the analysis excluding voxels in the occipital cortex using an ROI mask (Figure S11A), which was estimated in a general linear model with a regressor including all of the brand-viewing tasks. The mask was created using the group-average beta values across 17 subjects with the criteria of $p<0.002$ and at least 1000 contiguous voxels. We then ran the computational model for each participant excluding voxels within the mask. Specifically, we performed the same analysis but selected the 500 most stable voxels outside the mask. Our result is robust to masking for the overall hit rate and the significant relationship between hit rate and subjective similarity (Figure S11B), for the strength of neural correlation modulated by the similarity of brands' psychological features (Figure S11C), and for the significant relationship between the subjective similarity and the difference between the correlation of the predicted brain images and the actual brain images for correctly matched pair and the incorrectly matched pair (Figure S11D).

2.9 Robustness to Controlling for Physical Properties of Brand Logos

As a robustness check, we account for brain activities associated with the visual activation when viewing the brand logos, by comparing the result in three models with different sets of explaining variables. The first model is our main result, using the five factors of psychological features. In the second one, we used the other five variables obtained from the ratings of an independent population regarding to the physical properties of the brand logos, such as whether the logo is red, blue, round, whether it has hard edges, and whether it contains words (Figure S12A). In the third model, we included both sets of variables: five psychological features and five physical properties. We found that although the model of physical properties yielded a higher prediction rate compared with the psychological-feature model (68% versus 58%), the combination of psychological features and physical properties does the best (73%) (Figure S12B). Also, the hit rate is modulated by the similarity in traits only in the psychological-feature model (Figure S12C).

2.10 Regression Coefficients

To visualize the regression coefficients (c_n^v) of the five dimensions of psychological features, we regress the activation levels on the five dimensions for each voxel of the average subject. Notice that we did not hold out any data. The average regression coefficients within each of the clusters shown in Figure 4 are plotted in Figure S13-S17. The brain regions were ordered by anterior/posterior axis in clockwise fashion.

2.11 Univariate Analyses of Brand Personality Factors

We use a univariate approach to identify the brain regions significantly correlated with each of the five dimensions of brand personality, with parametric modulations in a general linear model using SPM8. The brain regions associated with each of the five factors are shown in Figure S18. Univariate analyses are typically less sensitive than multivariate analyses because the former does not consider information that is distributed among activity patterns between voxels. Consistent with this, we find only a few patches of the visual cortex that respond to brand personality factors at the $p<0.001$ level (Figure S18).

3. Figures

Figure S1: Experimental Protocol

Subjects engaged in a passive viewing task, and were instructed to think about the characteristics or traits associated with each brand. For each trial, a brand logo was presented for 4-8 seconds on a gray background. In addition, there were twelve fixations lasting for 4-8 seconds and six self-paced catch questions.

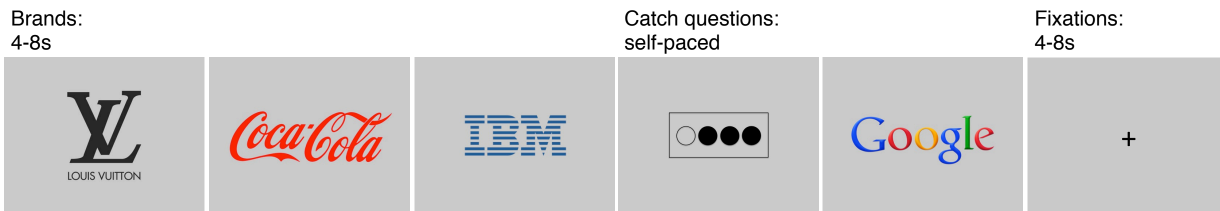

Figure S2: Factor Loadings of Traits

The factor analysis and the criteria yielded five factors, labeled as excitement, competence, sincerity, ruggedness, and sophistication. The factor loadings of traits showed that our results largely replicated those of previous studies. For example, the first factor loaded highly on the traits "trendy", "unique", and "original"—commonly referred to as the Excitement factor. The third factor, referred as Sincerity, loaded highly on traits such as "friendly", "family-oriented", and "wholesome".

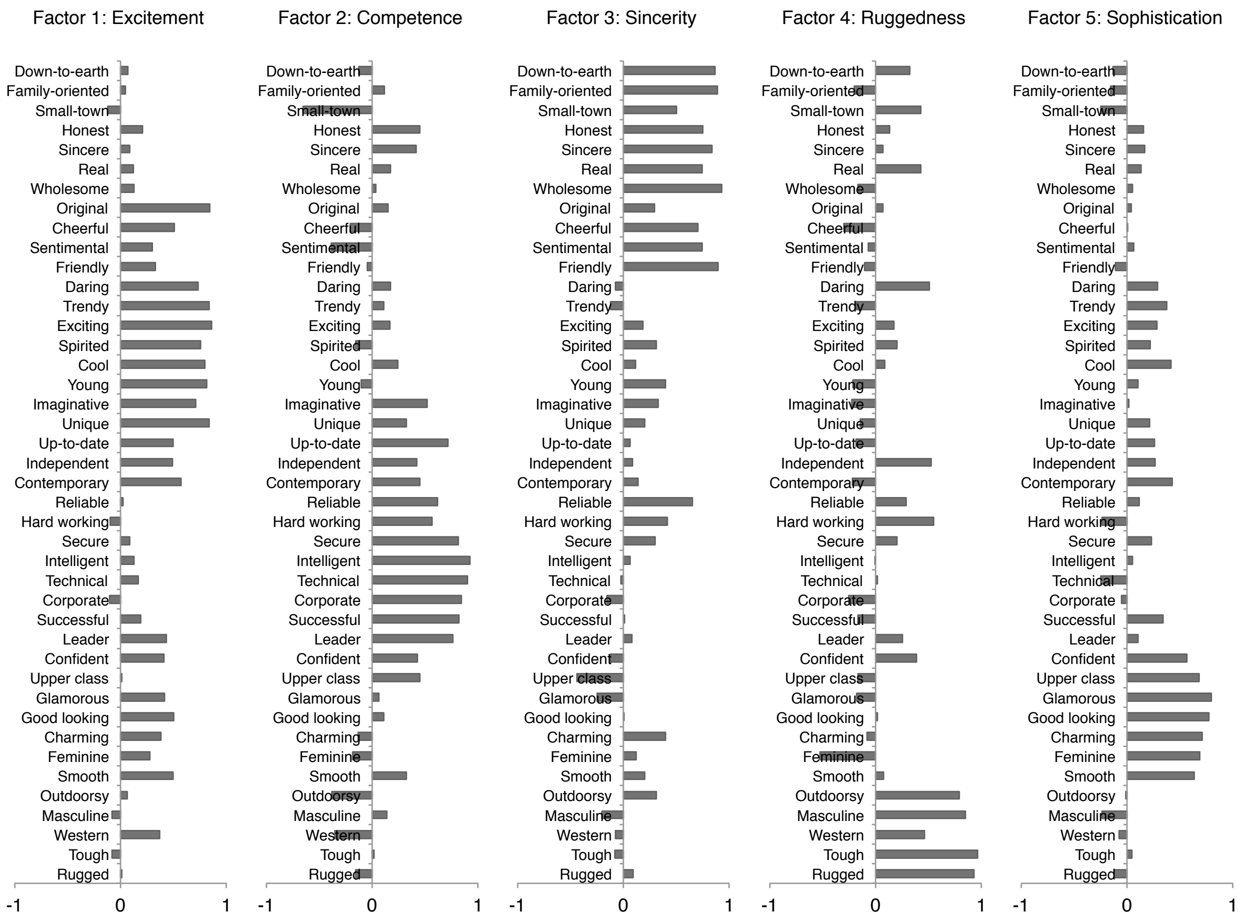

Figure S3: Radar Charts of the Factor Scores of Brands

Each brand was re-expressed in terms of its feature vector, defined as the strength of association between the brand and the personality factors. These factor scores for each brand are shown in the radar charts (Ex: excitement, Com: competence, Sin: sincerity, Rug: ruggedness, and So: sophistication). Green (Red) regions indicate positive (negative) factor scores.

Figure S4: Model Performance on the Individual Level

Average hit rate over all of the possible combination of hold-out data for each subject. Error bars represent 95% confidence intervals. Subjects were sorted by the performance of the model.

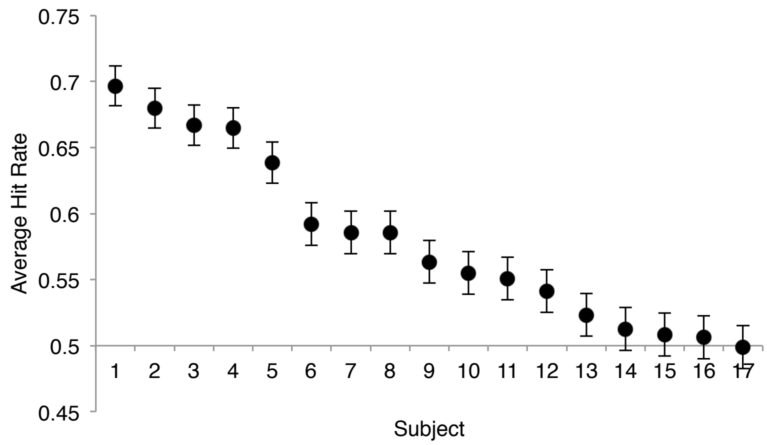

Figure S5: Correlation between Neural Similarity and Psychological Similarity.
Separating the brands based on subjective similarity into quartiles, we find a significant relationship between the subjective similarity and the difference between the correlation of the predicted brain images and the actual brain images for correctly matched pair (dark arrow in Figure 2D) and the incorrectly matched pair (light arrow in Figure 2D). That is, the difference in correlation between the matched brain images and the mismatched brain images is larger when brands are more dissimilar. When brands are highly similar (mean Pearson r=0.75), there is no significant difference between the correlation of matched images and the correlation of mismatched images. Errorbars indicate SEM.

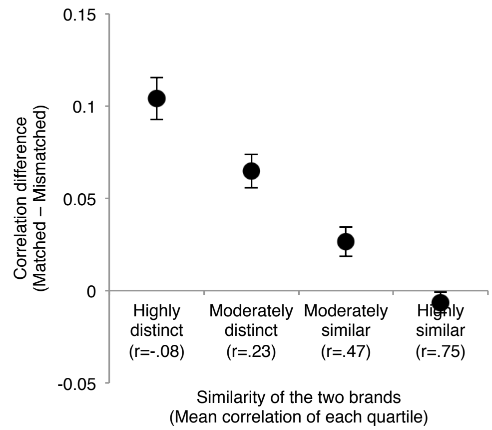

Figure S6: Accuracy Map

Surface rendering (top) and the glass brain (bottom) of locations of the most accurately predicted voxels, i.e., voxels with highest correlation between predicted and actual activations for the average participant. Each panel shows clusters containing at least 10 contiguous voxels where predicted-actual correlation is significantly greater than zero, with $p<0.05$ from the permutation test. These voxel clusters are distributed throughout the cortex and located in the left and right occipital and frontal lobes (Table 1). Note that decoding results were robust to exclusion of visual cortices.

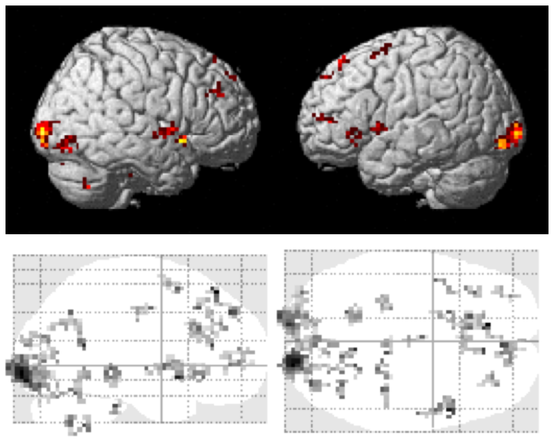

Figure S7: Robustness to Extracting Representative fMRI Response to the Brands

Our result is robust to different ways of extracting the representative of neural responses to brands. We are able to predict significantly better than chance which brand the participant was thinking about using the brain activities estimated (1) with impulse function in the LS-S model, (2) with the full duration in the LS-S model, and (3) with the standard one GLM model. **(A)** Consistent with Mumford et al. 2012, the prediction rates of the alternative models are somewhat worse than using the duration 0 LS-S model, but remain quite significant (p<0.005, Figure S7A). **(B)** The prediction rates modulated by psychological similarity remain qualitatively unchanged (Figure S7B).

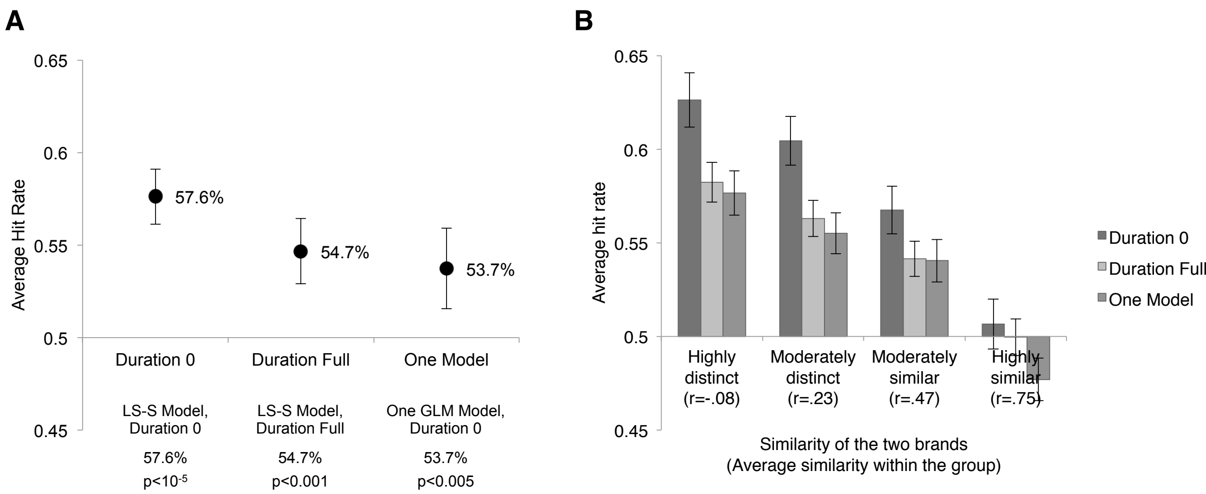

Figure S8: Robustness to the Measure of Psychological Similarity of Brands

(A) The difference in correlation between the matched brain images and the mismatched brain images is larger when brands are more dissimilar, measured as the larger Euclidean distance between the two brands' feature vectors. **(B)** We plotted, for each brand pair, the correlation between predicted and observed brand image (y-axis) against similarity in brand meaning as measured using Euclidean distance of factor scores (x-axis). We found that strength of neural correlation is robustly modulated by the similarity of brands' latent properties ($r = -0.56, p < 10^{-7}$). **(C)** Separating the brands based on subjective similarity into quartiles as assessed based on Euclidean distance of factor scores, we find a significant relationship between hit rate and subjective similarity.

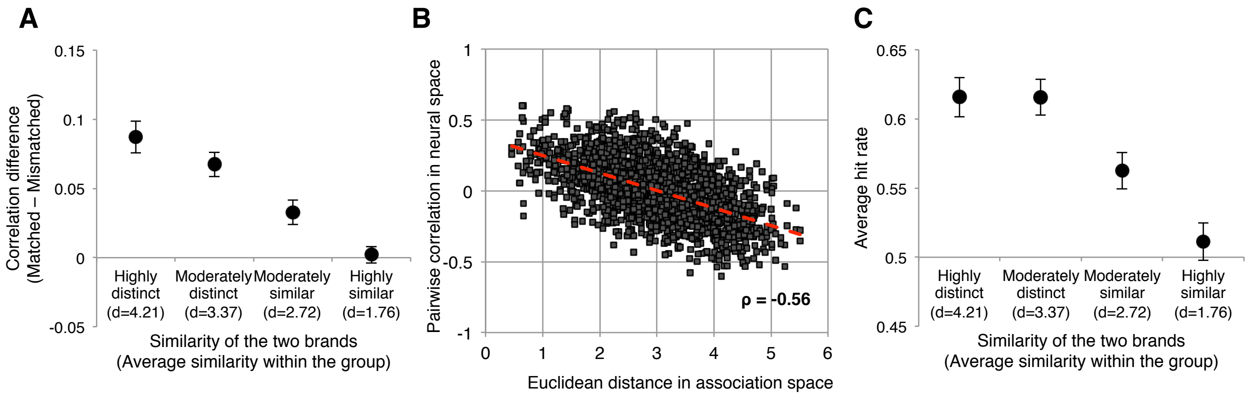

Figure S9: Robustness to Number of Voxels (Hit Rate)

The overall hit rate for hold-out classification was 58% when comparing the predicted and observed brain images using the 500 most stable voxels. When more voxels were included in the comparison, our result was still robust. Separating the brands based on subjective similarity into quartiles as assessed based on correlation of trait ratings, we find a significant relationship between hit rate and subjective similarity when different number of voxels were included.

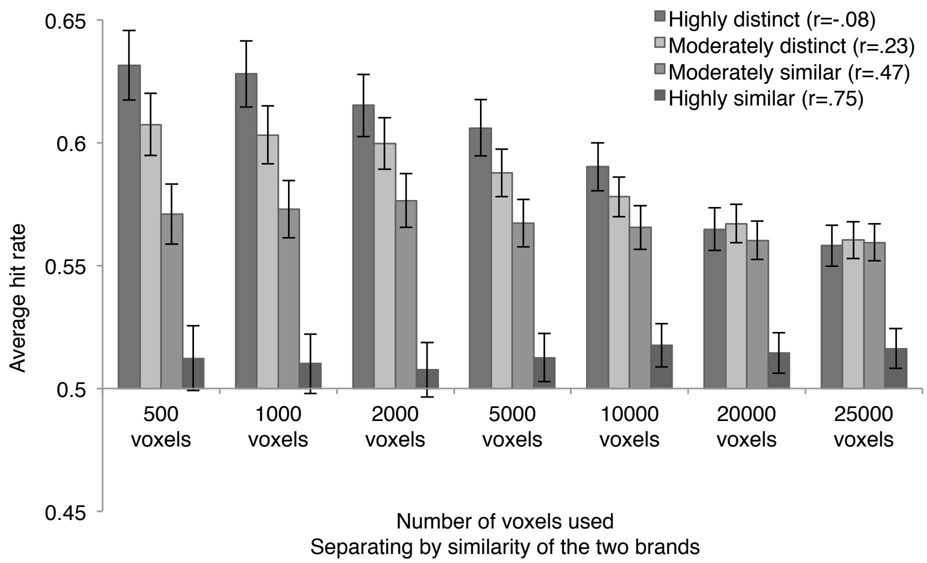

Figure S10: Robustness to Number of Voxels (Correlation)

To compare similarity between neural and psychological measures of brand associations, we plotted, for each brand pair, the correlation between predicted and observed brain images (y-axis) against similarity in brand meaning as measured using correlation of trait ratings (x-axis). The correlation between the predicted and the observed brain images was calculated using **(A)** 500 **(B)** 1000 **(C)** 5000, and **(D)** 25000 most stable voxels. We found that strength of neural correlation is robustly modulated by the similarity of brands' latent properties when different voxels were included.

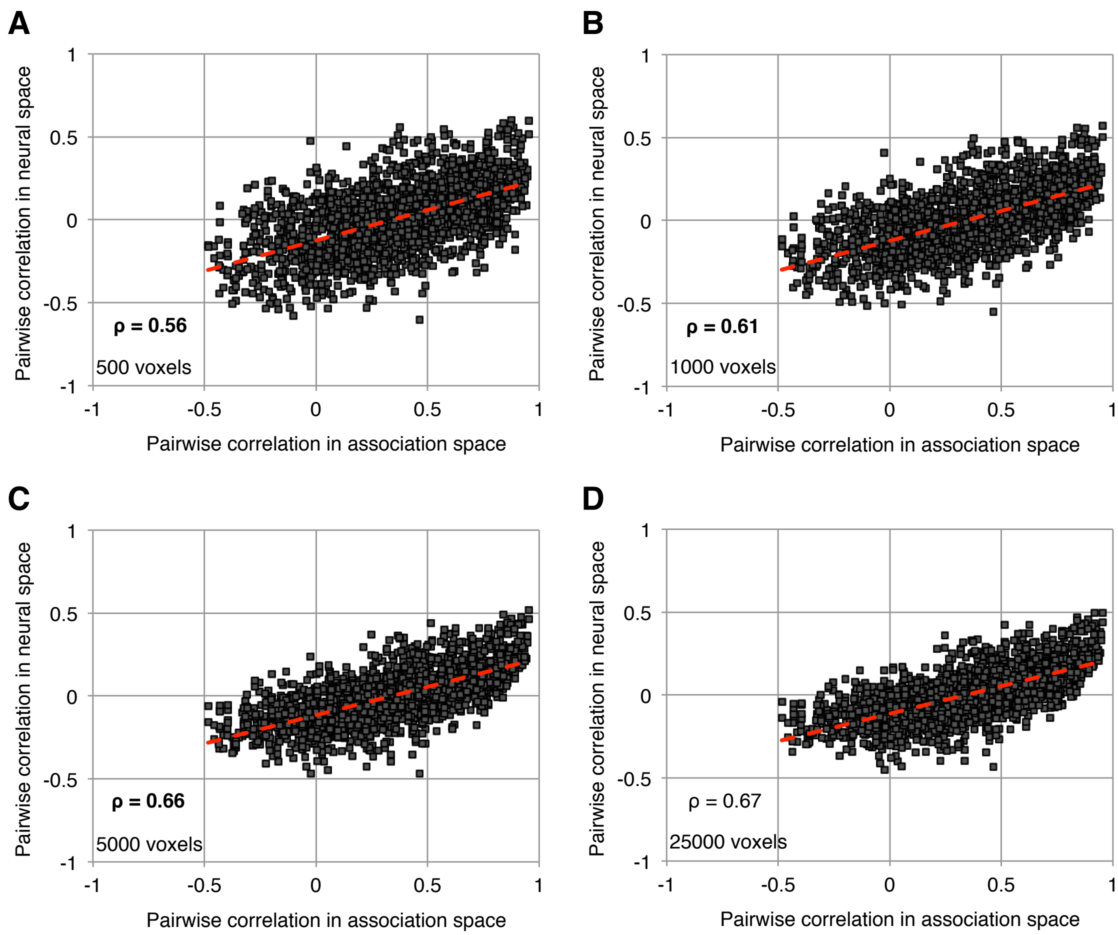

Figure S11: Robustness to Excluding Visual Cortex Voxels

We ran the analysis excluding voxels in the occipital cortex as a robustness check of the result. **(A)** The ROI mask used to exclude voxels within occipital lobes. **(B)** The overall hit rate was significantly better than chance, and the significant relationship between hit rate and subjective similarity was robust to masking. **(C)** The strength of neural correlation modulated by the similarity of brands' personality properties was robust to masking. **(D)** The significant relationship between the subjective similarity and the difference between the correlation of the predicted brain images and the actual brain images for correctly matched pair and the incorrectly matched pair was robust to masking.

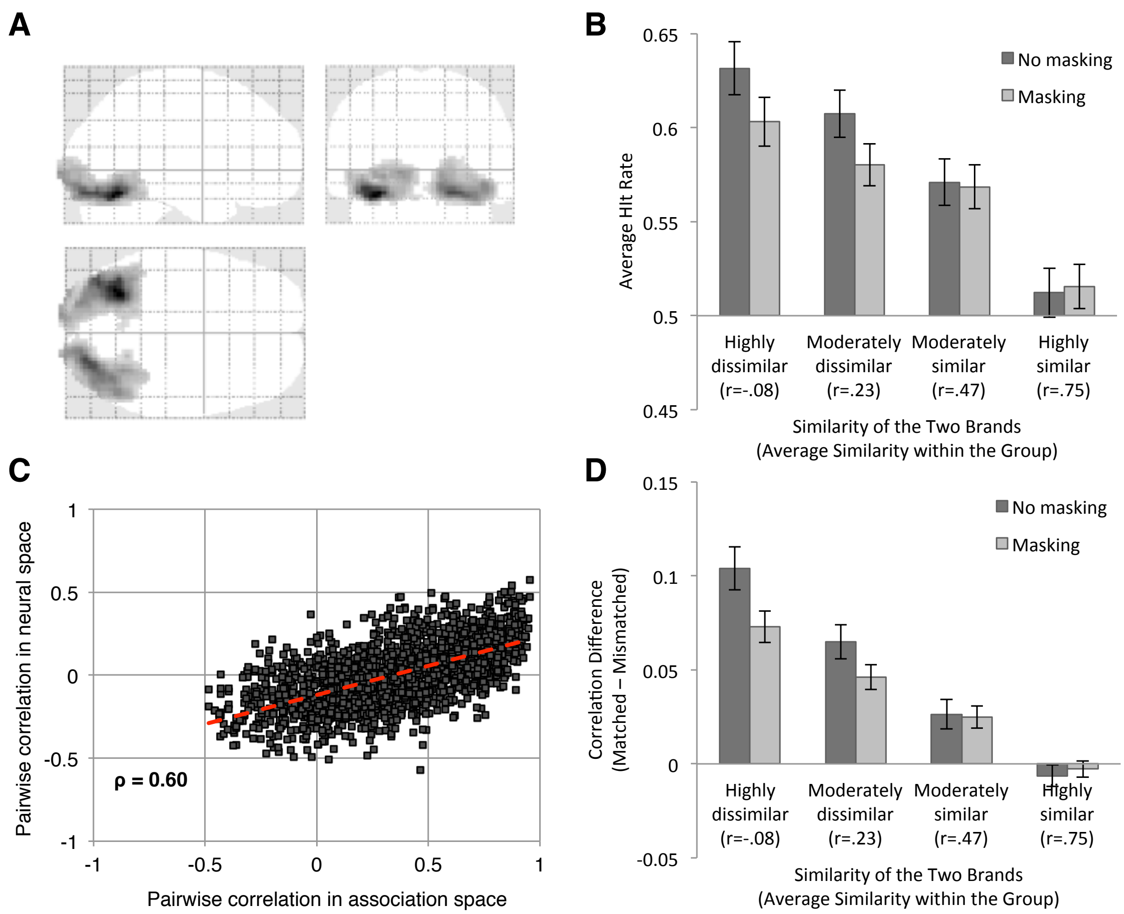

Figure S12: Robustness to Controlling for Physical Properties of Brand Logos

We compared the result in three models with different sets of explaining variables to account for brain activities associated with the visual activation. **(A)** Variables used in the models. **(B)** Although the model of physical properties yields a higher prediction rate compared with the psychological association model, the combination of factors and physical properties does the best. **(C)** The hit rate is modulated by the similarity in traits only in the psychological association model.

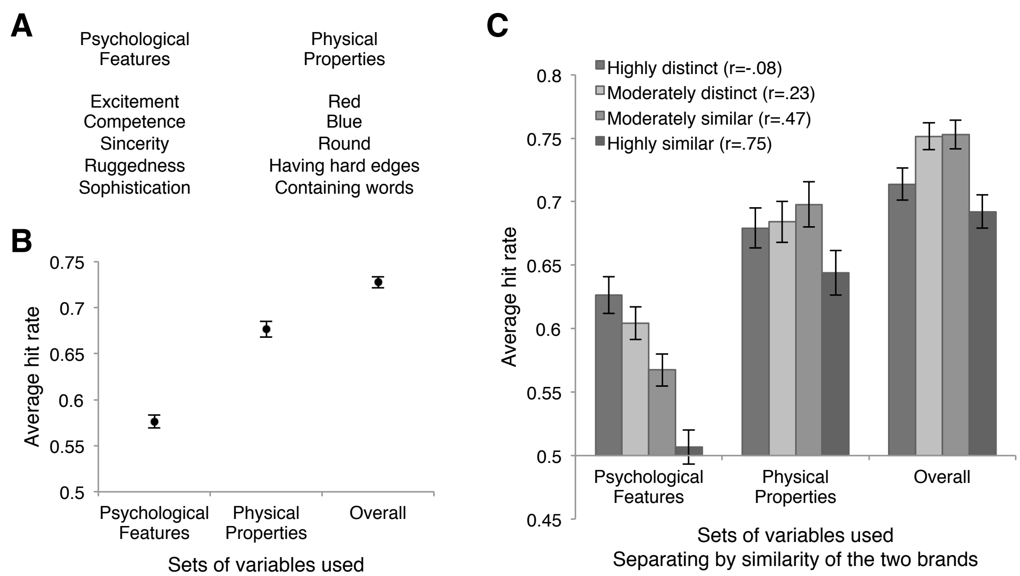

Figure S13: Average Regression Coefficients of Excitement

The average regression coefficients of excitement within each cluster in Figure 4 are shown in the figure, with the coordinate of the peak voxel in the parentheses.

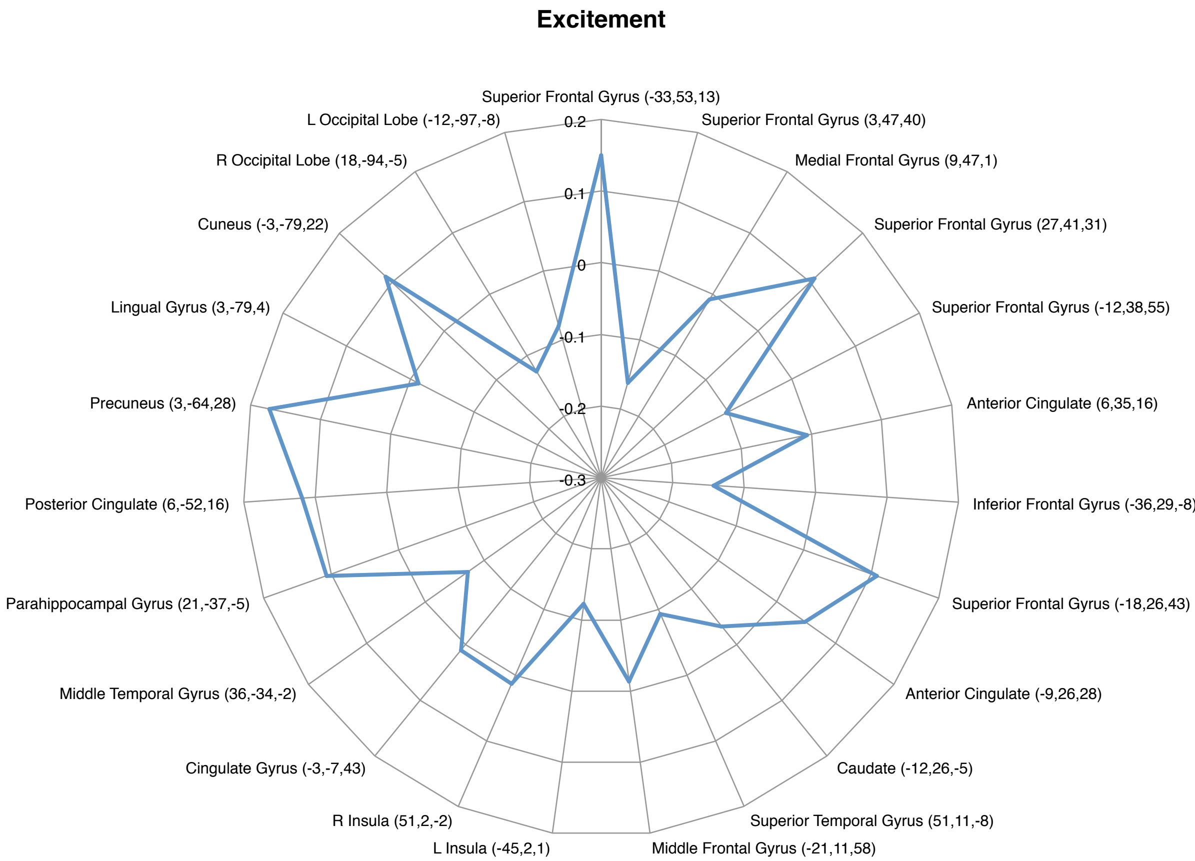

Figure S14: Average Regression Coefficients of Competence

The average regression coefficients of competence within each cluster in Figure 4 are shown in the figure, with the coordinate of the peak voxel in the parentheses.

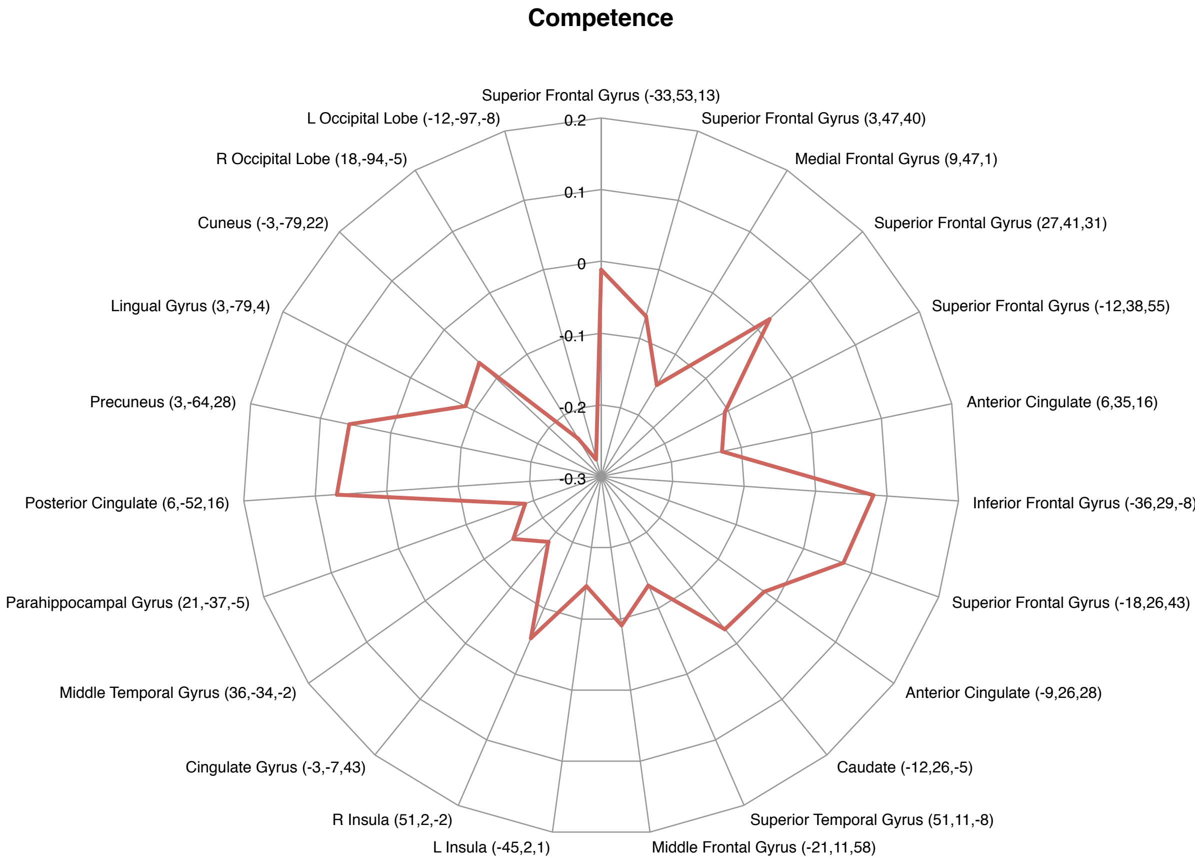

Figure S15: Average Regression Coefficients of Sincerity

The average regression coefficients of sincerity within each cluster in Figure 4 are shown in the figure, with the coordinate of the peak voxel in the parentheses.

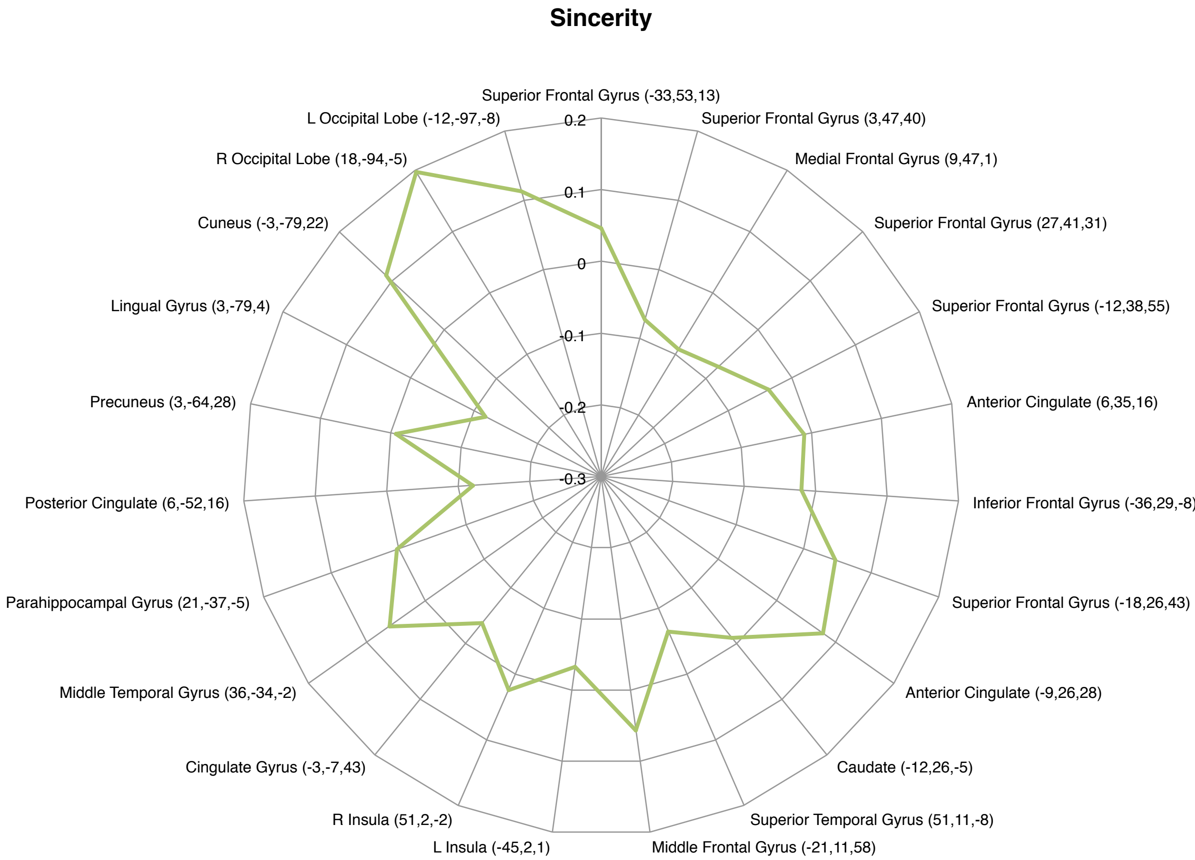

Figure S16: Average Regression Coefficients of Ruggedness

The average regression coefficients of ruggedness within each cluster in Figure 4 are shown in the figure, with the coordinate of the peak voxel in the parentheses.

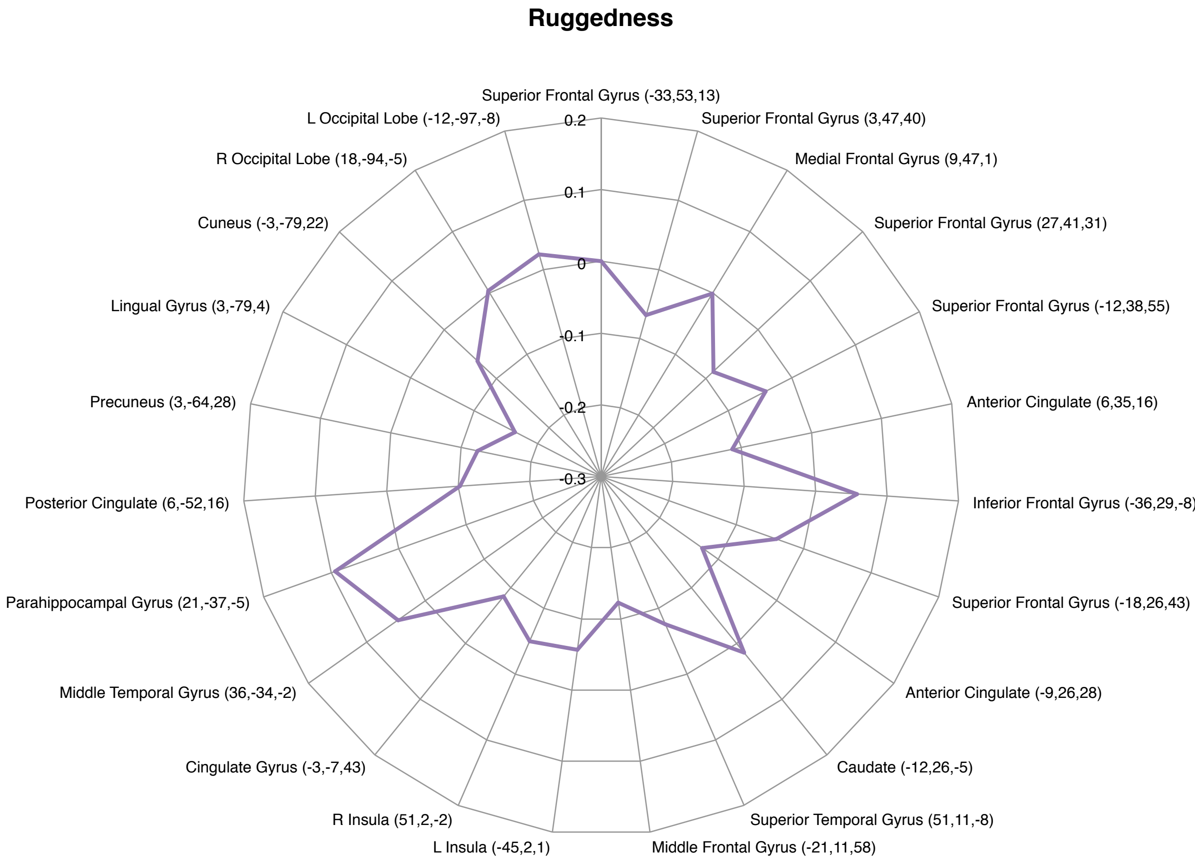

Figure S17: Average Regression Coefficients of Sophistication

The average regression coefficients of sophistication within each cluster in Figure 4 are shown in the figure, with the coordinate of the peak voxel in the parentheses.

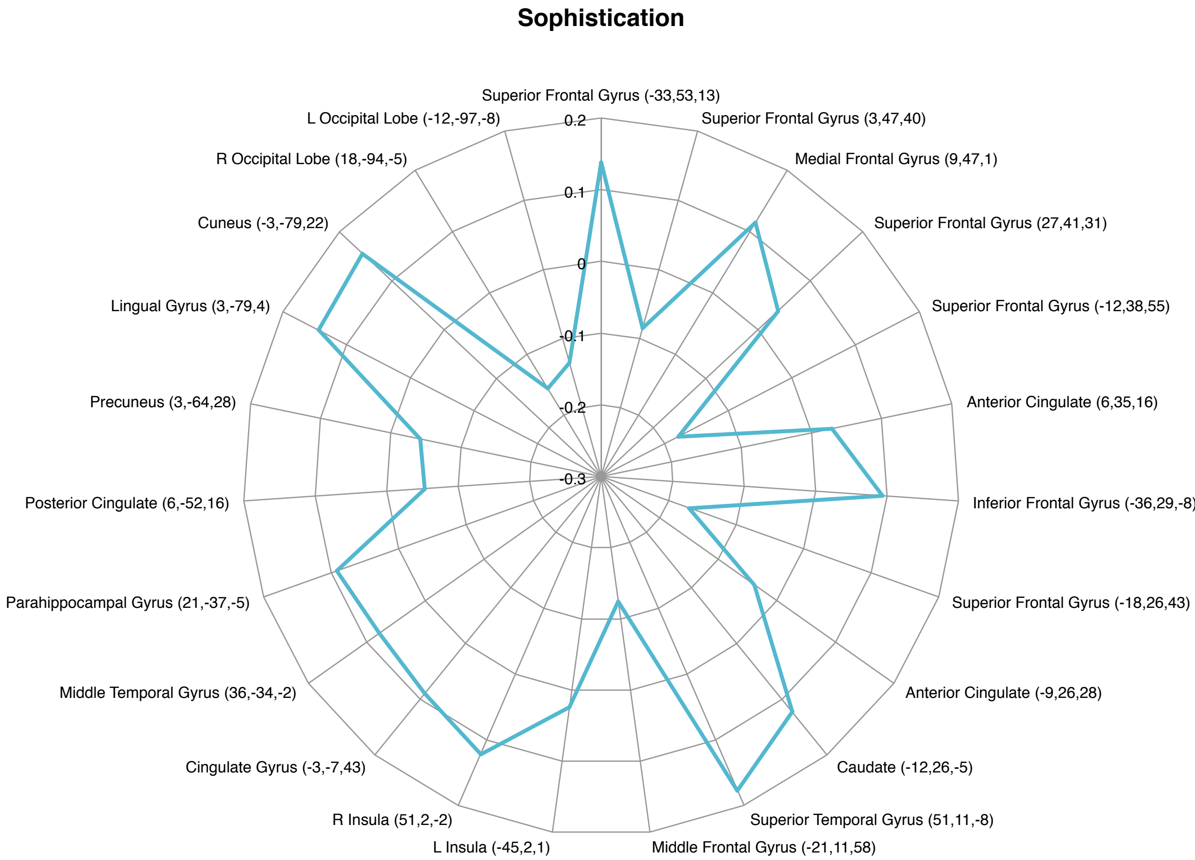

Figure S18: Univariate Analyses of Brand Personality Factors

Unlike our decoding analysis, univariate analyses are typically less sensitive than multivariate analyses because the former does not consider information that is distributed among activity patterns between voxels. Consistent with this, we find only a few patches of the visual cortex that respond to brand personality factors at the p<0.001 level.

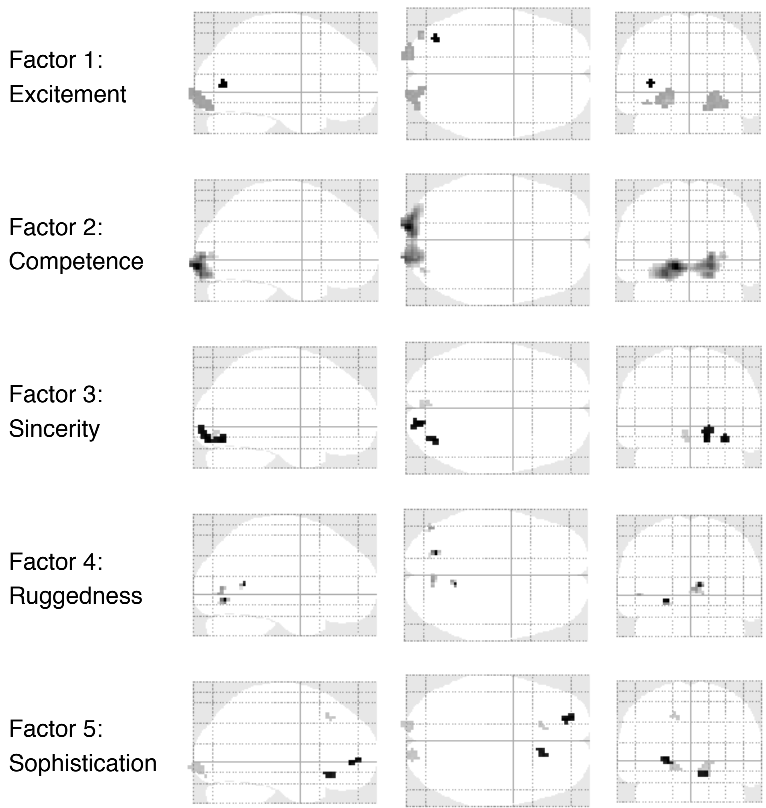

4. Tables

Table S1: Traits Used in the Survey

After the scanning session, participants were asked to complete a survey about the brands the saw in the scanner. For each brand, participant rated the descriptiveness of 42 traits, with a five-point scale from not at all descriptive to extremely descriptive.

Dimension	Sincerity	Excitement	Competence	Sophistication	Ruggedness
Traits	Down-to-earth Family oriented Small-town Honest Sincere Real Wholesome Original Cheerful Sentimental Friendly	Daring Trendy Exciting Spirited Cool Young Imaginative Unique Up-to-date Independent Contemporary	Reliable Hard-working Secure Intelligent Technical Corporate Successful Leader Confident	Upper-class Glamorous Good-looking Charming Feminine Smooth	Outdoorsy Masculine Western Tough Rugged

Table S2: Dimensions of Brand Characteristics
Summary of the five factors obtained from the factor analysis.

Name	Dimension	Variance Explained	Eigenvalue	Traits with highest item-to-total correlations
Excitement	1	35.0%	14.69	Exciting, Original, Unique, Trendy, Young.
Competence	2	19.1%	8.03	Intelligent, Technical, Corporate, Successful, Secure.
Sincerity	3	14.8%	6.22	Wholesome, Friendly, Family-oriented, Down-to-earth, Sincere.
Ruggedness	4	12.3%	5.19	Tough, Rugged, Masculine.
Sophistication	5	4.7%	1.96	Glamorous, Good-looking, Charming.

Table S3: Factor Scores of the Brands Used in the Experiment

Brand	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5
Amazon.com	0.924	1.282	0.638	0.083	0.101
American Express	-1.177	0.452	-0.898	-0.480	0.444
Apple	1.695	1.208	0.126	0.073	0.435
BMW	0.257	0.937	-0.696	0.683	1.467
Budweiser	-0.597	-1.998	-0.382	0.907	-1.168
Campbell’s	-1.175	-1.082	2.061	-0.139	0.011
Canon	0.195	1.082	0.532	0.050	-0.121
Cisco	-0.787	0.975	-0.397	-0.050	-0.780
Coca-Cola	0.339	-0.719	0.757	-0.265	0.170
Dell	-0.880	0.902	0.072	-0.140	-0.458
Disney	1.427	-0.277	1.634	-1.060	-0.330
Ford	-0.370	-0.739	0.432	1.678	-0.441
GE	-1.315	0.919	0.494	0.123	-0.200
Gillette	-0.849	0.033	0.131	1.377	0.887
Goldman Sachs	-1.343	0.025	-2.272	-0.247	-0.172
Google	1.980	1.505	0.339	0.141	-0.604
Gucci	0.662	-0.465	-1.383	-0.953	2.175
H&M	0.732	-1.124	-0.377	-1.069	0.329
Harley-Davidson	1.529	-0.835	-0.878	2.703	-0.455
HP	-0.564	0.678	-0.328	-0.715	-0.884
Honda	-0.104	0.384	0.918	0.434	0.467
IBM	-0.526	1.761	-0.515	-0.038	-1.089
IKEA	0.740	-0.112	1.114	-0.741	0.048
Intel	-0.356	1.699	-0.201	-0.143	-0.829
J.P. Morgan	-1.775	0.436	-1.609	-0.186	0.137
Jack Daniel’s	-0.077	-1.363	-0.533	1.604	0.194
Kellogg's	-0.914	-0.701	1.954	-0.574	0.170
L’Oréal	0.199	-0.718	-0.388	-1.229	1.455
Lancôme	-0.565	-1.090	-0.363	-1.534	2.325
Levi's	-0.150	-0.967	1.544	2.064	0.876
Louis Vuitton	0.646	-0.099	-1.517	-0.815	1.681
Marlboro	-0.587	-1.749	-1.655	1.350	-1.509
McDonald's	-0.576	-1.414	-0.837	-1.526	-2.114
Mercedes-Benz	-0.496	0.881	-0.339	0.585	2.072
Microsoft	-0.474	1.070	-0.263	-0.429	-1.169
MTV	2.534	-1.146	-1.187	-0.638	-0.879
Nestlé	-0.724	-1.217	1.757	-1.428	0.193
Nike	1.554	0.268	0.139	1.772	0.635
Nintendo	1.347	0.023	0.104	-1.207	-1.629
Pepsi	0.277	-0.781	0.395	-0.298	-0.292
Sony	0.517	1.160	-0.330	0.164	-0.136
Toyota	-0.527	0.288	0.884	-0.024	0.282
UPS	-1.233	0.769	1.126	0.851	-0.271
Yahoo!	0.586	-0.142	0.195	-0.715	-1.026

5. Reference

Aaker, Jennifer L (1997), “Dimensions of Brand Personality,” *Journal of Marketing Research*, 34(3), 347–56.

Fliessbach, K, B Weber, P Trautner, T Dohmen, U Sunde, C E Elger, and a Falk (2007), “Social Comparison Affects Reward-Related Brain Activity in the Human Ventral Striatum.,” *Science (New York, N.Y.)*, 318(5854), 1305–8.

“IBM SPSS Statistics for Windows, Version 20.0” (2011).

Izuma, Keise, Daisuke N Saito, and Norihiro Sadato (2008), “Processing of Social and Monetary Rewards in the Human Striatum.,” *Neuron*, 58(2), 284–94.

Mumford, Jeanette a, Benjamin O Turner, F Gregory Ashby, and Russell a Poldrack (2012), “Deconvolving BOLD Activation in Event-Related Designs for Multivoxel Pattern Classification Analyses.,” *NeuroImage*, 59(3), 2636–43.