1 Supplementary Discussion

1.1 Overview

First, we report how well the trained classifiers could detect the presence of each image category (1.2), for both target and distractor stimuli. We then present *feature importance maps* illustrating which features played the largest role in detecting these categories (1.2.2). After that, we present the results of our quartile analysis in greater detail (1.3): We report pair-wise comparison statistics showing that the reaction-time priming effect varied in size as a function of distractor-processing quartile (1.3.1). In the following section, we present pair-wise comparison statistics measuring how RTs varied across quartiles for the control and ignored-repetition conditions separately (1.3.2). We also present statistics showing that the reaction-time priming effect did not vary significantly as a function of target-processing quartile (1.3.3). Next, we report results from a variant of the analysis where we included error trials (1.3.4) – the classifier analyses in the main paper excluded these trials. Finally, we report priming effects as a function of distractor-processing quartile separately for each of the four categories of images used as stimuli (1.3.5). For all of our statistical comparisons, we used two-tailed paired-samples t-tests to compute the reliability of effects across subjects.

1.2 Detailed classification results

1.2.1 Classifier sensitivity

In the main body of the paper, we described the cross-validation method that we used to evaluate the classifiers' ability to detect the presence of each image category when presented as the target stimulus. Figure 4 in the main paper showed the mean level of cross-validation performance for target stimuli across subjects for each time bin, combining results from all four categories. Supplementary Figure 1 re-plots these results, this time splitting the results by target stimulus category. Supplementary Table 1 reports the mean accuracy results for each category when, for

Prime target category

	Shoe	Chair	House	Face
Mean AUC (SEM)	0.55 (0.01)	0.57 (0.01)	0.60 (0.01)	0.71 (0.03)
Statistics	t(15) = 4.40 (p =0.0005)	t(15) = 5.83 (p <1e-4)	t(15) = 9.18 (p <1e-6)	t(15) = 8.01 (p <1e-6)

Supplementary Table 1: **Sensitivity of each of the category classifiers to the presence of the target image on the prime display.**

each trial, we averaged classifier output across the time bins from 20ms to 960ms post-stimulus-onset (see the *Combining classifier estimates across time bins* part of the Methods section in the main paper). The figure and table show that, for all four categories, the classifiers were significantly above chance at separating trials in which the image category was on-screen-as-the-target from trials in which the image category was not on-screen.

Supplementary Table 2 breaks down the target classification results even further: It shows mean target classification accuracy for each category, conditionalized on the category of the distractor image (e.g., it shows classification accuracy for shoe targets, as a function of whether the distractor image was a face, chair, or house). Face and shoe target classification were both unaffected by distractor category. However, for house and chair targets, there was an effect of distractor category on accuracy as evidenced by one-way repeated-measures ANOVAs (for house targets, $F(2, 30) = 10.36, p < .001$; for chair targets, $F(2, 30) = 6.32, p < .01$). Further inspection shows that sensitivity to house and chair targets was worst when faces were onscreen as distractors.

Next, as described in the main body of the paper, the trained category-specific classifiers were applied to trials in which that category appeared as the distractor, in order to estimate the extent to which the subject processed the distractor on each trial. Figure 5 in the main paper showed the mean level of cross-validation performance for distractor stimuli across subjects for each time bin, combining results from all four categories. Supplementary Figure 2 re-plots these results, this time splitting the results by distractor stimulus category. Supplementary Table 3

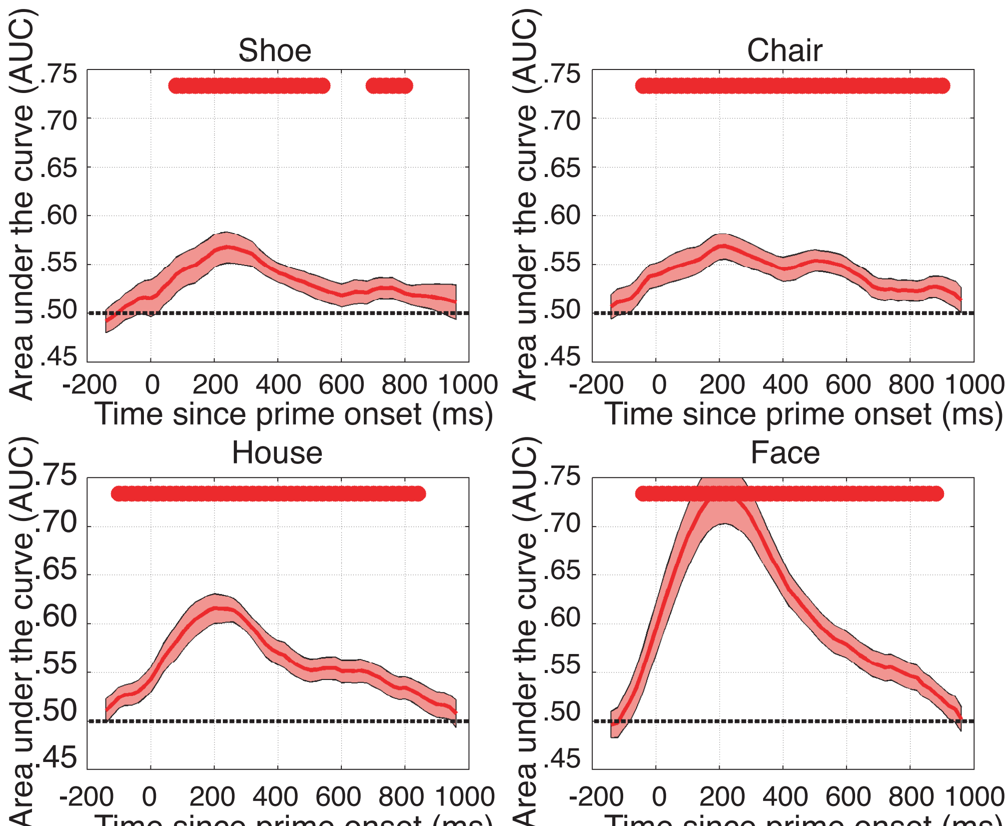

Supplementary Figure 1: Target classification (cross-validation) performance as a function of target stimulus category. Beginning in the upper-left corner and rotating around clockwise, the graphs show the performance of the shoe, chair, face, and house classifiers. The red lines plot the AUC computed at each time bin (see text for description of how this was computed). The shaded regions around each line indicate the standard error across subjects. The dashed black lines at 0.5 mark chance performance. The red dots along the top of each plot indicate which of the time bin classifiers performed significantly above chance at the $p < .05$ level.

		Prime target category			
		Shoe	Chair	House	Face
Prime distractor category	Shoe	-	0.59 (0.02)	0.61 (0.02)	0.71 (0.03)
	Chair	0.55 (0.02)	-	0.64 (0.01)	0.71 (0.03)
	House	0.56 (0.02)	0.57 (0.02)	-	0.70 (0.03)
	Face	0.55 (0.02)	0.55 (0.01)	0.55 (0.02)	-

Supplementary Table 2: **Sensitivity of each of the category classifiers to the presence of the target image on the prime display, conditionalized on the category of the prime distractor image.** Each cell reports the mean (standard error) over subjects of the area under the curve (AUC) indicating the sensitivity to the presence of the target image.

reports the mean accuracy results for each category when, for each trial, we averaged classifier output across the time bins from 20ms to 960ms post-stimulus-onset. The figure and table show that, for all four categories, the classifiers were significantly above chance at separating trials in which the image category was on-screen-as-the-distractor from trials in which the image category was not on-screen. As we discussed in the main paper, we expected that the classifier's sensitivity to distractors would be much lower than its sensitivity to targets, insofar as distractors (on average) are processed much less strongly than targets. The main point of the distractor sensitivity analysis was to show that classification performance was above floor for these items – this is a prerequisite for using the classifier to measure trial-by-trial fluctuations in distractor processing.

Supplementary Table 4 breaks down the distractor classification results even further: It shows mean distractor classification accuracy for each category, conditionalized on the category of the target image (e.g., it shows classification accuracy for shoe distractors, as a function of whether the target image was a face, chair, or house). Classification of face distractors was unaffected by the category of the target image. However, shoe, chair, and house distractor classification were all affected by the category of the target image (for shoe distractors, $F(2,30) = 7.54, p < .01$; for chair distractors, $F(2,30) = 10.37, p < .001$; for house distractors, $F(2,30) = 7.28, p < .01$); in all three cases, distractor classification was worst when faces

	Prime distractor category			
	Shoe	Chair	House	Face
Mean AUC (SEM)	0.52 (0.004)	0.52 (0.007)	0.52 (0.01)	0.58 (0.02)
Statistics	$t(15)$ = 3.98 (p =0.001)	$t(15)$ = 3.13 (p =0.007)	$t(15)$ = 2.67 (p =0.02)	$t(15)$ = 4.29 (p =0.0006)

Supplementary Table 3: **Sensitivity of each of the category classifiers to the presence of the distractor image on the prime display.**

		Prime target category			
		Shoe	Chair	House	Face
Prime distractor category	Shoe	-	0.53 (0.01)	0.55 (0.01)	0.47 (0.01)
	Chair	0.54 (0.01)	-	0.55 (0.01)	0.48 (0.01)
	House	0.55 (0.01)	0.55 (0.01)	-	0.48 (0.02)
	Face	0.58 (0.02)	0.59 (0.02)	0.58 (0.02)	-

Supplementary Table 4: **Sensitivity of each of the category classifiers to the presence of the distractor image on the prime display, conditionalized on the category of the prime target image.** Each cell reports the mean (standard error) over subjects of the area under the curve (AUC) indicating the sensitivity to the presence of the distractor image.

were onscreen as targets. These results all fit with the idea that faces capture subjects' attention more strongly than the other categories: When a face is on screen as the target, the distractor category is processed less strongly (see Supplementary Table 4); also, when a face is present as a distractor, it can impede processing of the target category (see Supplementary Table 2).

As noted in the main paper, the sensitivity of the trained classifiers was greater for detecting processing of the target stimuli than for detecting processing of the distractor stimuli. Figures 4 and 5 in the main paper plot the results of our sensitivity analysis for these two types of stimuli separately. Supplementary Figure 3 plots the mean difference in classifier sensitivity for target vs. distractor stimuli, and indicates which time-bin-specific classifiers were significantly more sensitive to processing of targets vs. distractors. This figure shows that, at almost all time points, the classifiers were significantly more sensitive to processing of the target stimuli.

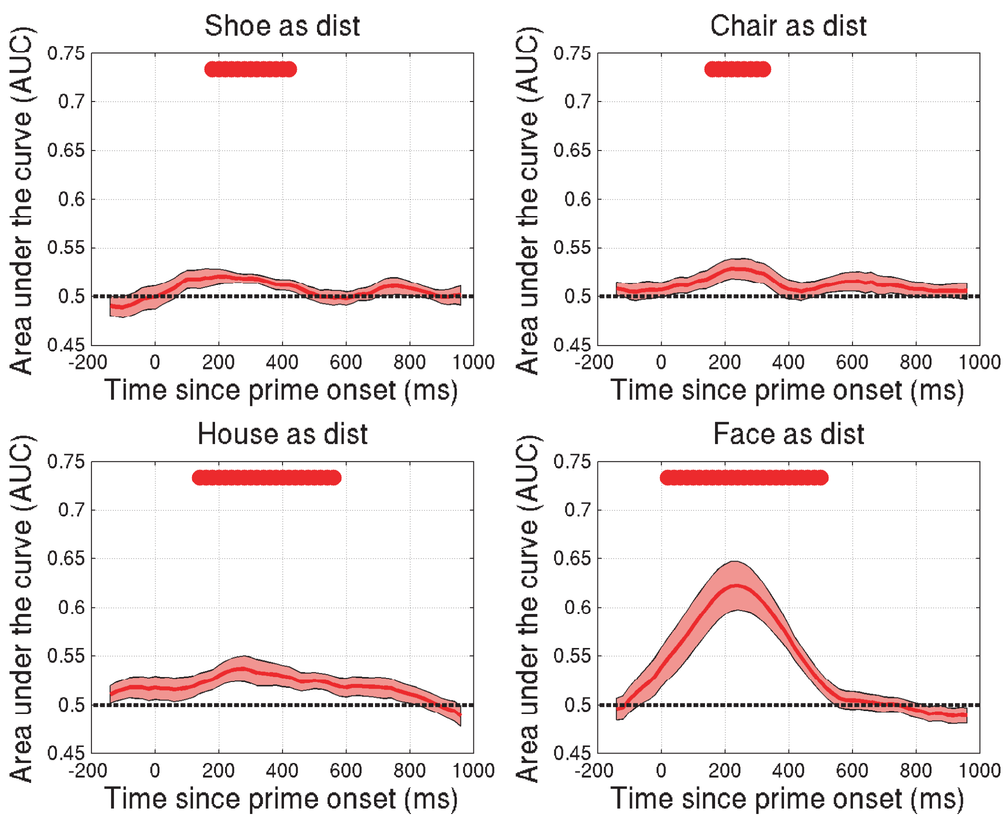

Supplementary Figure 2: Distractor classification (cross-validation) performance as a function of distractor stimulus category. Beginning in the upper-left corner and rotating around clockwise, the graphs show the performance of the shoe, chair, face, and house classifiers. The red lines plot the AUC computed at each time bin (see text for description of how this was computed). The shaded regions around each line indicate the standard error across subjects. The dashed black lines at 0.5 mark chance performance. The red dots along the top of each plot indicate which of the time bin classifiers performed significantly above chance at the $p < .05$ level.

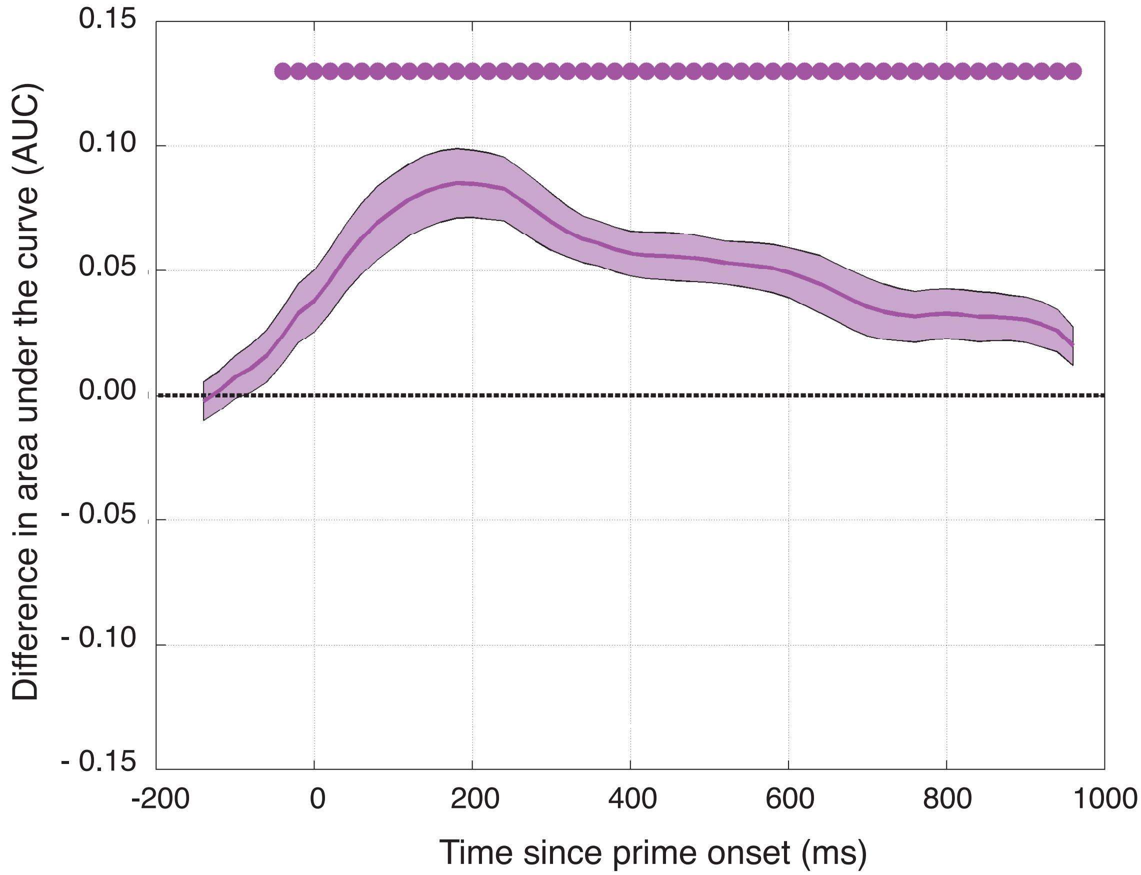

Supplementary Figure 3: Difference in average sensitivity of the classifier to the category of the prime target stimulus and prime distractor stimulus, combining across the four categories. The purple line plots the mean difference in the area under the ROC (AUC) computed for target and distractor stimuli. The shaded region around this line indicates the standard error across subjects. The dashed black line along 0 indicates no difference. The dots along the top of the figure indicate which of the time-bin-specific classifiers performed significantly better in detecting the target stimuli than in detecting the distractor stimuli at the $p < .05$ level.

1.2.2 Features used for classification

To evaluate which features were most important in detecting the presence of each category, we created *importance maps* for each condition (McDuff et al., 2009). These importance maps show *which features were most responsible for driving each classifier's response to the category that it was built to detect*. For example, which features were most responsible for driving the response of the face-as-target classifier, on trials where faces were presented as targets?

In ridge regression classifiers, the net contribution of a feature to detecting a category is a function of the feature's value, multiplied by the weight assigned to the feature. Because of the z-transformation that we applied to feature values prior to classification, the (transformed) features that we fed into the classifier could take on either positive or negative values. As such, there were two ways for a feature to make a net positive contribution to detecting a particular category:

- The feature could have a positive z-transformed value for the category (indicating that its value was above-average when the category was present) and a positive weight. Features meeting this criterion were assigned a positive importance value $imp_{ij} = w_{ij} * avg_{ij} * 100$ where w_{ij} is the weight between the input feature i and the output unit for category j, and avg_{ij} is the mean value of input feature i over all trials from the category j.

- The feature could have a negative z-transformed value for the category (indicating that its value was below-average when the category was present) and a negative weight. In this case, the "double negative" combination of negative feature value and negative weight results in a net positive contribution. Features meeting this criterion were assigned a negative importance value $imp_{ij} = -w_{ij} * avg_{ij} * 100$.

Features where the sign of w_{ij} differed from the sign of avg_{ij} (indicating that the feature made a net negative contribution to detecting the presence of the category) were assigned an

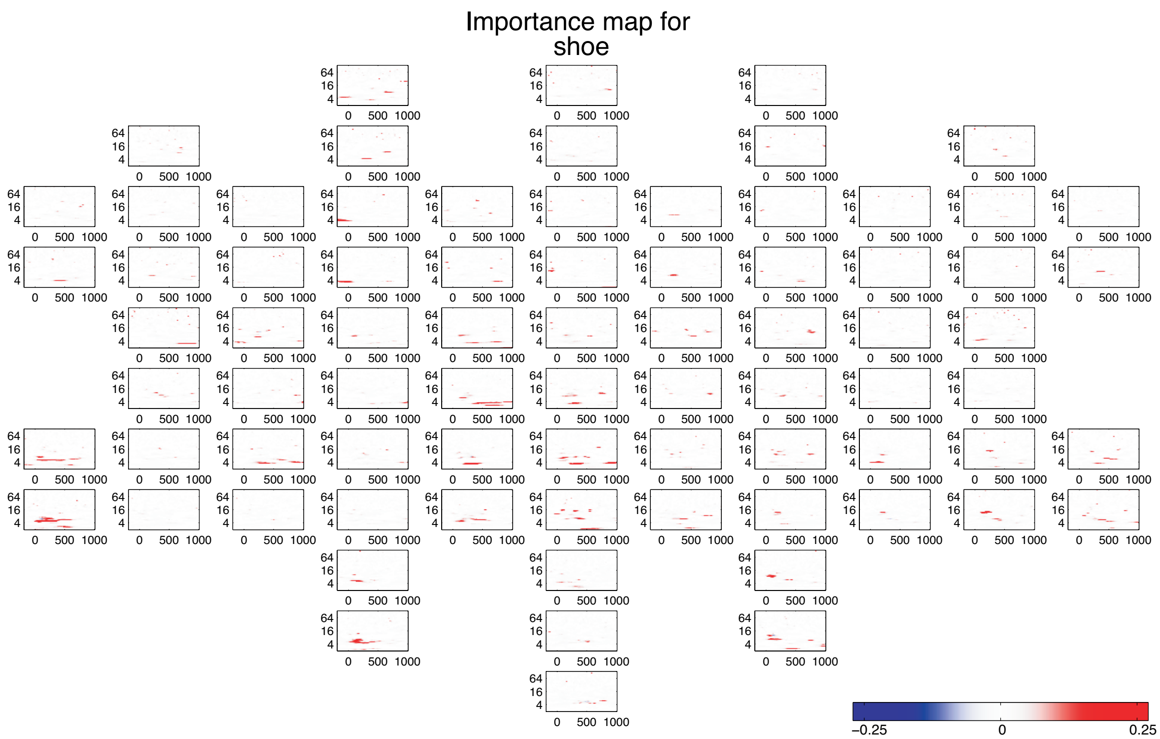

Supplementary Figure 4: Feature importance map for classification of shoes. This figure indicates which features were most important (across subjects) for detecting processing of the shoe category. Each graph plots the importance of frequency-time bin pairings from a particular electrode (see text for description of how importance values were computed). The graphs are laid out as though the viewer were looking down onto the electrode array from behind the subject's head. The top corresponds to the electrodes on the forehead, the right-most graphs correspond to electrodes surrounding the subject's right ear, and vice versa for the bottom and left side of the figure. The x-axis of each graph corresponds to time since the prime stimulus onset, and the y-axis indicates the frequency band of each feature.

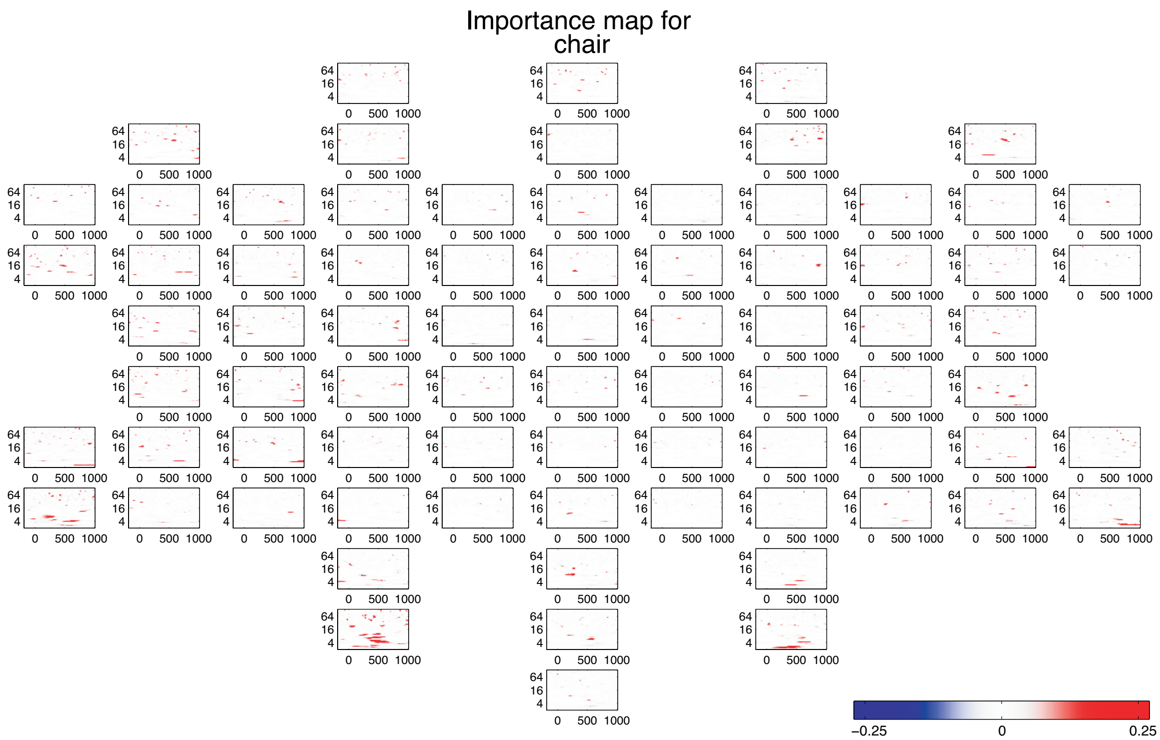

Supplementary Figure 5: Feature importance map for classification of chairs. This figure indicates which features were most important (across subjects) for detecting processing of the chair category. Each graph plots the importance of frequency-time bin pairings from a particular electrode (see text for description of how importance values were computed). The graphs are laid out as though the viewer were looking down onto the electrode array from behind the subject's head. The top corresponds to the electrodes on the forehead, the right-most graphs correspond to electrodes surrounding the subject's right ear, and vice versa for the bottom and left side of the figure. The x-axis of each graph corresponds to time since the prime stimulus onset, and the y-axis indicates the frequency band of each feature.

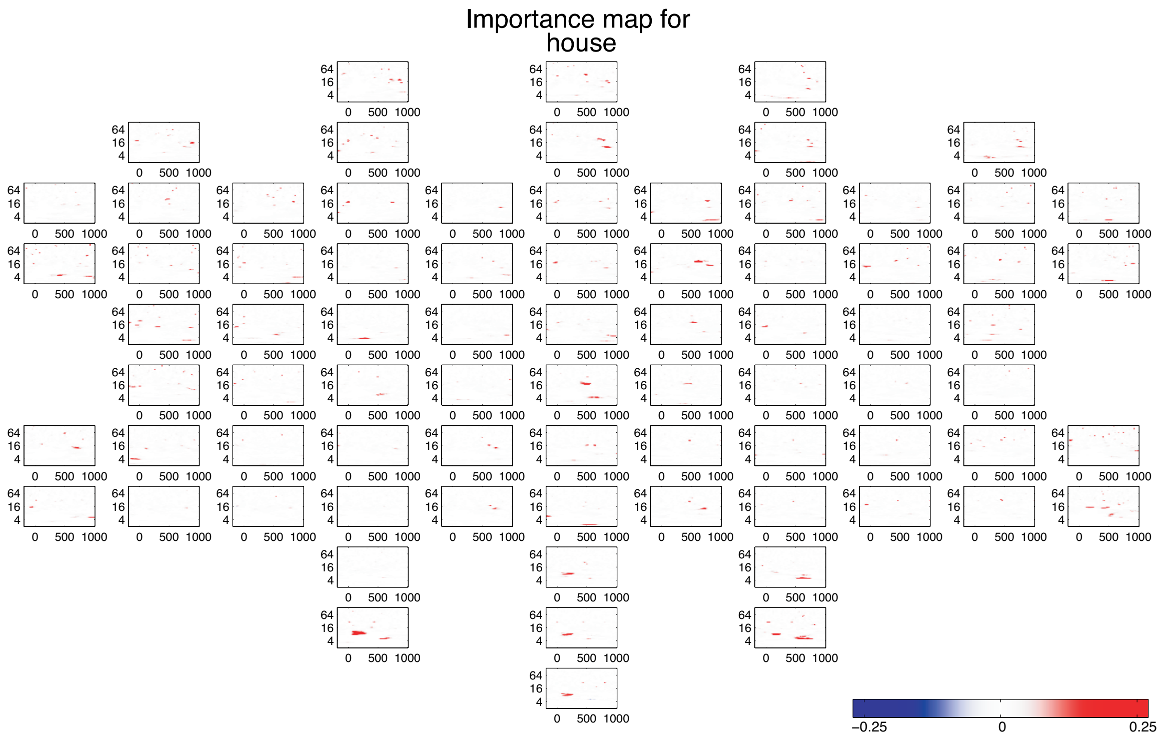

Supplementary Figure 6: Feature importance map for classification of houses. This figure indicates which features were most important (across subjects) for detecting processing of the house category. Each graph plots the importance of frequency-time bin pairings from a particular electrode (see text for description of how importance values were computed). The graphs are laid out as though the viewer were looking down onto the electrode array from behind the subject's head. The top corresponds to the electrodes on the forehead, the right-most graphs correspond to electrodes surrounding the subject's right ear, and vice versa for the bottom and left side of the figure. The x-axis of each graph corresponds to time since the prime stimulus onset, and the y-axis indicates the frequency band of each feature.

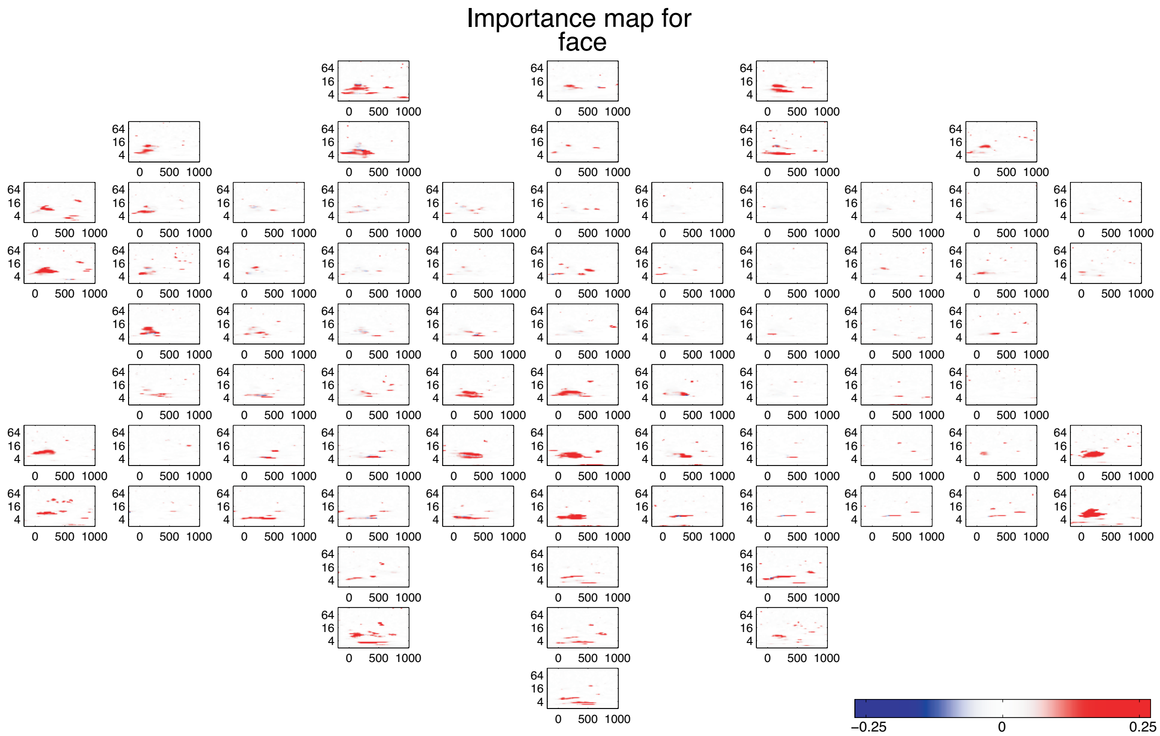

Supplementary Figure 7: Feature importance map for classification of faces. This figure indicates which features were most important (across subjects) for detecting processing of the face category. Each graph plots the importance of frequency-time bin pairings from a particular electrode (see text for description of how importance values were computed). The graphs are laid out as though the viewer were looking down onto the electrode array from behind the subject's head. The top corresponds to the electrodes on the forehead, the right-most graphs correspond to electrodes surrounding the subject's right ear, and vice versa for the bottom and left side of the figure. The x-axis of each graph corresponds to time since the prime stimulus onset, and the y-axis indicates the frequency band of each feature.

importance value of zero.

Importance values were computed using the above equations for each individual subject. We then computed the mean importance value (across subjects) for each feature. The importance map for shoes is shown in Supplementary Figure 4; the importance map for chairs is shown in Supplementary Figure 5; the importance map for houses is shown in Supplementary Figure 6; and the importance map for faces is shown in Supplementary Figure 7. In these figures, positive importance values (indicating that the presence of the category was associated with above-average values of the feature) are plotted in red, and negative importance values (indicating that the presence of the category was associated with below-average values of the feature) are plotted in blue.

1.3 Quartile analysis details

In the main body of the paper, we showed the effect of distractor processing on subjects' subsequent RTs by plotting the priming effect for each distractor-processing quartile (Figure 6); In this section, we present the results of our quartile analyses in greater detail.

1.3.1 Statistics for pair-wise comparisons of distractor-processing quartiles

Figure 6 in the main paper indicates that the priming effect in the medium-low distractor-processing quartile was significantly different from the priming effect in each of the other quartiles. The full statistics for these and all other pair-wise comparisons are reported in Supplementary Table 5. These results show that the only statistically reliable differences were between the medium-low distractor-processing quartile and the other quartiles.

1.3.2 Relationship between RT and distractor processing by condition

Figure 7 in the main paper plots the relationship between RTs and distractor processing within the ignored repetition and control conditions. This figure indicates that, in the control condition,

Quartile

	Low	Med. low	Med. high	High
Low	-	t(15)=2.52 (p =.02)	t(15)=0.37 (p =.72)	t(15)=1.27 (p =.22)
Med. low	36 ms	-	t(15)=2.76 (p =.01)	t(15)=2.94 (p =.01)
Med. high	-6 ms	-41 ms	-	t(15)=1.08 (p =.30)
High	-26 ms	-62 ms	-20 ms	-

Supplementary Table 5: **Pair-wise comparisons of quartile priming effects.** The values shown below the diagonal indicate the priming effect computed by subtracting the quartiles named in the left column from those named across the top row. The values shown above the diagonal indicate the statistical reliability of the difference across subjects.

Quartile

	Low	Med. low	Med. high	High
Low	-	t(15)=0.97 (p =.34)	t(15)=1.07 (p =.30)	t(15)=0.98 (p =.34)
Med. low	36 ms	-	t(15)=1.81 (p =.09)	t(15)=1.57 (p =.14)
Med. high	-6 ms	-41 ms	-	t(15)=0.20 (p =.85)
High	-26 ms	-62 ms	-20 ms	-

Supplementary Table 6: **Pair-wise comparisons of quartile response times in the ignored repetition condition.** The values shown below the diagonal indicate the priming effect computed by subtracting the quartiles named in the left column from those named across the top row. The values shown above the diagonal indicate the statistical reliability of the difference across subjects.

the RTs in the medium-low distractor-processing quartile were significantly different from the RTs in the low and high distractor-processing quartiles. The full statistics for these and all other pair-wise comparisons are reported in Supplementary Tables 6 and 7. These tables show that the only statistically reliable differences were between the medium-low distractor-processing quartile and the low and high distractor-processing quartiles in the control condition.

1.3.3 Statistics for pair-wise comparisons of target-processing quartiles

Figure 8 in the main paper indicates that priming effects did not vary as a function of target-processing quartiles. The full statistics for all of the pair-wise comparisons between target-processing quartiles are reported in Supplementary Table 8.

Quartile

	Low	Med. low	Med. high	High
Low	-	$t(15)$=2.52 (p =.02)	$t(15)$=0.30 (p =.76)	$t(15)$=0.86 (p =.40)
Med. low	36 ms	-	$t(15)$=2.03 (p =.06)	$t(15)$=2.52 (p =.02)
Med. high	-6 ms	-41 ms	-	$t(15)$=1.22 (p =.24)
High	-26 ms	-62 ms	-20 ms	-

Supplementary Table 7: **Pair-wise comparisons of quartile response times in the control condition.** The values shown below the diagonal indicate the priming effect computed by subtracting the quartiles named in the left column from those named across the top row. The values shown above the diagonal indicate the statistical reliability of the difference across subjects.

Quartile

	Low	Med. low	Med. high	High
Low	-	$t(15)$=0.73 (p =.48)	$t(15)$=1.34 (p =.20)	$t(15)$=1.57 (p =.14)
Med. low	-16 ms	-	$t(15)$=0.64 (p =.53)	$t(15)$=0.49 (p =.63)
Med. high	-31 ms	-15 ms	-	$t(15)$=0.28 (p =.78)
High	-24 ms	-8 ms	7 ms	-

Supplementary Table 8: **Pair-wise comparisons of quartile priming effects when the trials were split by the level of prime target processing.** The values shown below the diagonal indicate the priming effect computed by subtracting the quartiles named in the left column from those named across the top row. The values shown above the diagonal indicate the statistical reliability of the difference across subjects.

Quartile

	Low	Med. low	Med. high	High
Low	-	t(15)=2.41 (p =.03)	t(15)=0.53 (p =.60)	t(15)=1.03 (p =.32)
Med. low	36 ms	-	t(15)=2.59 (p =.02)	t(15)=2.81 (p =.01)
Med. high	-9 ms	-44 ms	-	t(15)=0.69 (p =.50)
High	-23 ms	-58 ms	-14 ms	-

Supplementary Table 9: **Pair-wise comparisons of quartile priming effects including trials with errors.** The values shown below the diagonal indicate the priming effect computed by subtracting the quartiles named in the left column from those named across the top row. The values shown above the diagonal indicate the statistical reliability of the difference across subjects. This table differs from Supplementary Table 5 in that the values were computed over all trials, including those with inaccurate responses.

1.3.4 Quartile analysis with error trials included

The classifier analyses presented in the main paper excluded error trials (which were very rare: percent correct accuracy was close to 98% in both the ignored repetition and control conditions). Here, we present the results of quartile analyses that included error trials. First, we wanted to assess whether the main distractor-quartile analysis showed the same results when we included error trials. Supplementary Table 9, like Supplementary Table 5 above, compares the NP effect for the four distractor processing quartiles – the only difference between Supplementary Table 9 and Supplementary Table 5 is that the analysis shown in Supplementary Table 9 includes error trials. The pattern of results shown in Supplementary Table 9 is essentially the same as the pattern with error trials excluded: The priming effect in the medium-low quartile is significantly more negative than the priming effect in the other three quartiles.

Including error trials also makes it possible for us to examine how accuracy varied as a function of distractor processing quartile and condition (control vs. ignored repetition). As shown in Supplementary Table 10, none of the within-quartile effects of condition were significant. Furthermore, as shown in Supplementary Table 11, and the effect of condition on accuracy did not differ significantly across quartiles.

	Quartile			
	Low	Med. Low	Med. High	High
Control accuracy (SEM)	97.9 (1.0)	97.7 (1.3)	98.1 (1.0)	97.8 (1.0)
Ignored repetition accuracy (SEM)	98.7 (0.6)	98.6 (0.7)	98.2 (1.1)	98.1 (1.0)
Statistics	t(15) = 0.73 (p =0.48)	t(15) = 0.75 (p =0.46)	t(15) = 0.10 (p =0.92)	t(15) = 0.25 (p =0.81)

Supplementary Table 10: **Percent accuracy across quartiles for each condition.** The statistics indicate the reliability of the difference in accuracy between the two conditions across subjects.

	Quartile			
	Low	Med. low	Med. high	High
Low	-	t(15)=0.02 (p =.99)	t(15)=0.44 (p =.66)	t(15)=0.41 (p =.68)
Med. low	-0.02	-	t(15)=0.52 (p =.61)	t(15)=0.41 (p =.69)
Med. high	-0.76	-0.73	-	t(15)=0.08 (p =.93)
High	-0.60	-0.58	-0.16	-

Supplementary Table 11: **Pair-wise comparisons of quartile accuracy effects.** The values shown below the diagonal indicate the difference in accuracy computed by subtracting the quartiles named in the left column from those named across the top row. The values shown above the diagonal indicate the statistical reliability of the difference across subjects.

	Quartile			
Shoes	Low	Med. low	Med. high	High
Low	-	t(15)=0.06 (p =.95)	t(15)=0.76 (p =.46)	t(15)=0.28 (p =.78)
Med. low	-2 ms	-	t(15)=0.71 (p =.49)	t(15)=0.26 (p =.79)
Med. high	-44 ms	-43 ms	-	t(15)=0.92 (p =.37)
High	-13 ms	-11 ms	31 ms	-

Supplementary Table 12: **Pair-wise comparisons of quartile priming effects within the shoe-as-distractor trials.** The values shown below the diagonal indicate the priming effect computed by subtracting the quartiles named in the left column from those named across the top row. The values shown above the diagonal indicate the statistical reliability of the difference across subjects.

1.3.5 Quartile analysis by image category

Figure 6 in the main paper shows the quartile priming effects when trials from all four image categories were pooled together. The same qualitative pattern (i.e., maximal suppression after medium-low distractor processing and maximal facilitation after high distractor processing) persisted within the individual stimulus categories. In three of the four categories (all of the categories except shoes), the medium-low distractor processing quartile showed the strongest negative priming effect, and the high distractor processing quartile showed a numerically positive priming effect. The four panels of Supplementary Figure 8 show the priming effect for each quartile, computed separately for each stimulus category. All significant pair-wise comparisons are indicated in the figure with the corresponding p value. Most of the pair-wise tests failed to reach significance because the within-category RTs were computed over fewer trials (compared to the main analysis) and were noisier as a result. Pair-wise comparison statistics for the four categories are reported in Supplementary Tables 12 - 15.

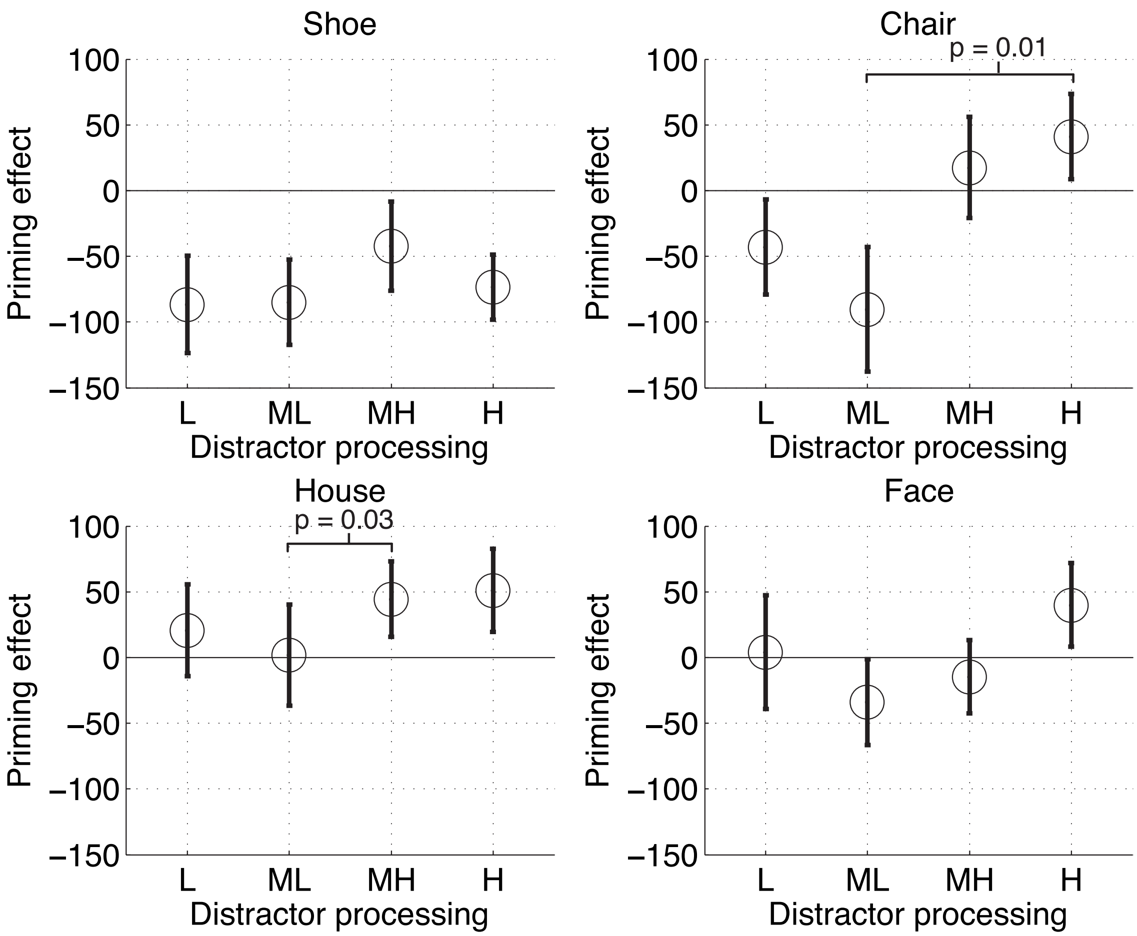

Supplementary Figure 8: Comparison of the priming effects within each quartile, done separately for each stimulus category. (Distractor processing abbreviations: L = low; ML = medium-low; MH = medium high; H = high) Significance values reflect the reliability of the difference across subjects, calculated using a two-tailed paired-samples t-test. Error bars indicate standard errors on the mean priming effect within each quartile across subjects.

	Quartile			
Chairs	Low	Med. low	Med. high	High
Low	-	$t(15)$=0.90 (p =.38)	$t(15)$=1.10 (p =.29)	$t(15)$=1.77 (p =.10)
Med. low	47 ms	-	$t(15)$=1.74 (p =.10)	$t(15)$=2.81 (p =.01)
Med. high	-60 ms	-108 ms	-	$t(15)$=0.65 (p =.53)
High	-84 ms	-131 ms	-23 ms	-

Supplementary Table 13: **Pair-wise comparisons of quartile priming effects within the chair-as-distractor trials.** The values shown below the diagonal indicate the priming effect computed by subtracting the quartiles named in the left column from those named across the top row. The values shown above the diagonal indicate the statistical reliability of the difference across subjects.

	Quartile			
Houses	Low	Med. low	Med. high	High
Low	-	$t(15)$=0.43 (p =.67)	$t(15)$=0.81 (p =.43)	$t(15)$=0.98 (p =.34)
Med. low	19 ms	-	$t(15)$=2.35 (p =.03)	$t(15)$=1.12 (p =.28)
Med. high	-33 ms	-69 ms	-	$t(15)$=0.03 (p =.98)
High	-30 ms	-49 ms	-1 ms	-

Supplementary Table 14: **Pair-wise comparisons of quartile priming effects within the house-as-distractor trials.** The values shown below the diagonal indicate the priming effect computed by subtracting the quartiles named in the left column from those named across the top row. The values shown above the diagonal indicate the statistical reliability of the difference across subjects.

	Quartile			
Faces	Low	Med. low	Med. high	High
Low	-	$t(15)$=1.00 (p =.33)	$t(15)$=0.31 (p =.76)	$t(15)$=0.67 (p =.51)
Med. low	48 ms	-	$t(15)$=0.59 (p =.57)	$t(15)$=1.72 (p =.11)
Med. high	17 ms	-19 ms	-	$t(15)$=1.20 (p =.25)
High	-26 ms	-74 ms	-55 ms	-

Supplementary Table 15: **Pair-wise comparisons of quartile priming effects within the face-as-distractor trials.** The values shown below the diagonal indicate the priming effect computed by subtracting the quartiles named in the left column from those named across the top row. The values shown above the diagonal indicate the statistical reliability of the difference across subjects.

References

McDuff S, Frankel H, Norman K (2009) Multivoxel pattern analysis reveals increased memory targeting and reduced use of retrieved details during single-agenda source monitoring. *Journal of Neuroscience* 29:508–516.