

# The Neural Basis of Visual Word Form Processing: A Multivariate Investigation

Adrian Nestor<sup>1,2</sup>, Marlene Behrmann<sup>1,2</sup> and David C. Plaut<sup>1,2</sup>

<sup>1</sup>Center for the Neural Basis of Cognition, Carnegie Mellon University, 4400 Fifth Avenue, Pittsburgh, PA 15213, USA

<sup>2</sup>Department of Psychology, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA, 15213, USA

Address correspondence to Adrian Nestor. Email: anestor@andrew.cmu.edu

**Current research on the neurobiological bases of reading points to the privileged role of a ventral cortical network in visual word processing. However, the properties of this network and, in particular, its selectivity for orthographic stimuli such as words and pseudowords remain topics of significant debate. Here, we approached this issue from a novel perspective by applying pattern-based analyses to functional magnetic resonance imaging data. Specifically, we examined whether, where and how, orthographic stimuli elicit distinct patterns of activation in the human cortex. First, at the category level, multivariate mapping found extensive sensitivity throughout the ventral cortex for words relative to false-font strings. Secondly, at the identity level, the multi-voxel pattern classification provided direct evidence that different pseudowords are encoded by distinct neural patterns. Thirdly, a comparison of pseudoword and face identification revealed that both stimulus types exploit common neural resources within the ventral cortical network. These results provide novel evidence regarding the involvement of the left ventral cortex in orthographic stimulus processing and shed light on its selectivity and discriminability profile. In particular, our findings support the existence of sublexical orthographic representations within the left ventral cortex while arguing for the continuity of reading with other visual recognition skills.**

**Keywords:** face recognition, fMRI, multivariate mapping, visual Word Form Area

Extensive research has established the involvement of the left human ventral cortex in reading, specifically with regard to orthographic processing (for meta-analyses, see Fiez and Petersen 1998; Turkeltaub et al. 2002; Jobard et al. 2003; Mechelli et al. 2003). However, the existence of dedicated neural circuitry subserving the encoding and processing of high-level orthographic representations continues to provoke vigorous debate (Dehaene and Cohen 2011; Price and Devlin 2011).

First and foremost, the existence of a left ventral region specialized in processing orthographic stimuli has been the source of significant disagreement—by “orthographic stimuli” we refer here to alphabetic strings such as words and pseudowords that obey orthographic rules, unlike consonant strings. This region, known as the “visual Word Form Area” (vWFA; Cohen et al. 2000), responds robustly and consistently to word stimuli (Turkeltaub et al. 2002; Vigneau et al. 2005) and exhibits invariance to a number of visual characteristics such as case (Dehaene et al. 2001; Polk and Farah 2002), language-specific script (Bolger et al. 2005; Baker et al. 2007) and types of visual stimulation (Rauschecker et al. 2011). Invariance to such characteristics unrelated to stimulus identity (i.e. independent of the specific word displayed) is particularly telling in that it suggests that the vWFA encodes high-level orthographic representations that abstract away from the specifics of the visual display (for short, “visual word forms”). On the other hand, the inconsistent response profile of this region

across different studies (e.g. when contrasted with objects) casts doubt on its domain specificity (Price and Devlin 2003; Wright et al. 2008). For instance, a number of studies failed to find an advantage in the left ventral cortex for words over pictures (Wright et al. 2008; Van Doren et al. 2010) or over orthographic control stimuli such as unfamiliar scripts (Xue and Poldrack 2007; Vogel et al. in press). Such findings prompt a reassessment of the vWFA in terms of more general visual processing that does not imply functional specificity for orthographic stimuli.

Secondly, the involvement of the vWFA in sublexical rather than full-fledged lexical processing has also been disputed (Kronbichler et al. 2004; Glezer et al. 2009; Schurz et al. 2010). The original proposal for a vWFA (Dehaene et al. 2002, 2005) argued for its encoding of sublexical structures, that is, linguistic structures such as pseudowords that obey orthographic and phonotactic rules without necessarily forming actual words. This proposal found ground in the equivalent levels of activation detected for words and pseudowords at the level of the vWFA (Dehaene et al. 2002; Binder et al. 2006; Vinckier et al. 2007). However, more recent results challenged the original proposal after finding marked sensitivity to words over pseudowords in the vWFA through the use of finer-grained analyses and tools such as fMRI adaptation (Glezer et al. 2009). Such findings suggest that the role of the vWFA is more specific than initially thought. More precisely, these data speak to its involvement in the storage and recognition of lexical items (i.e. of actual words).

The debates summarized above are fueled by empirical discrepancies arising from multiple sources, including the type of stimuli tested (Szwed et al. 2011), the spatial resolution of the neural data (Baker et al. 2007) and the statistical rigor underlying the analysis (Wright et al. 2008). More generally, these discrepancies may also reflect the intrinsic limitations of the univariate analyses on which the results are based. Univariate estimates of neural responses aggregate activation across entire regions at the cost of the rich pattern of information encoded within each region. In a different domain, that of face recognition, these limitations have been successfully overcome by using multivariate pattern analyses to evaluate the selectivity and domain specificity of relevant cortical areas (Haxby et al. 2001; Spiridon and Kanwisher 2002; Kriegeskorte et al. 2007; Nestor et al. 2011). Similarly, the present work employed novel multivariate analyses and contrasted their results with those of standard univariate analyses, in order to clarify and potentially resolve ongoing debates concerning the neural basis of visual word form processing.

Our investigation took a 2-pronged approach. First, we used multivariate “searchlight” mapping (Kriegeskorte et al. 2006) to locate ventral regions able to support category-level discrimination (i.e. distinguishing words from false-font strings). Secondly, we applied pattern classification to assess

identity-level discrimination (i.e. distinguishing a particular pseudoword from others despite variations in font). The ability to discriminate among different orthographic stimuli irrespective of extraneous visual properties (e.g. font) represents a critical aspect of reading. Thus, our analysis tested a direct implication of current theories of reading by determining whether and where different orthographic stimuli are encoded by distinct neural patterns. In addition, since orthographic stimuli and faces appear to compete for resources within the left ventral cortex (Dehaene et al. 2010; Plaut and Behrmann 2011), we conducted similar analyses within the reading network to assess its contribution to individual face discrimination.

In summary, the present work investigates coarse category-level as well as fine-grained identity-level encoding of orthographic stimuli within the ventral cortex and assesses the specificity of the neural mechanisms for orthographic stimulus processing by means of multivariate analysis.

## Materials and Methods

### Subjects

Each of 8 subjects (age range 18–22, 5 females) was scanned across 3 different sessions carried out on different days. All subjects were right-handed native English speakers with normal or corrected-to-normal vision and no history of neurological or cognitive disorders. Data were collected for 2 additional subjects who were excluded from the analysis due to large head movements (more than a voxel) on at least 1 out of 3 sessions. All subjects provided written informed consent. The Institutional Review Board of Carnegie Mellon University approved all imaging and behavioral procedures.

### Stimuli

To investigate category-level differences in discriminability and selectivity, subjects viewed stimuli belonging to 5 visual categories: words, false-font strings, faces, houses, and objects.

Word stimuli consisted of 5-letter high-frequency nouns (Kucera–Francis frequency >60) extracted from the MRC psycholinguistic database (Coltheart 1981). As a visual control for these stimuli, we constructed false-font strings by rearranging the strokes of each Roman letter into a new character (1 false font per letter) and assembled them in groups of 5 (see Fig. 1A for example). Thus, the 2 categories are comparable in terms of visual complexity as measured by the type and number of features in their makeup. Also, both categories were presented in high contrast (white characters against a

black background). In addition, subjects were presented with color images of front-view faces, houses, and common objects. These images were normalized with respect to mean luminance, contrast, and size.

Each of the stimulus categories described above contained 104 images that were presented exactly once to each subject. To be clear, false-font strings served as a control category for words, while houses and objects served as control categories for faces. However, false-font strings and words, on the one hand, and the remaining categories, on the other hand, were markedly different with respect to their visual properties and were not optimized for comparison against each other.

To investigate identity-level differences, participants were shown images of 4 different pseudowords designed to be highly and equivalently word-like (with high summed positional bigram and trigram frequencies). Pseudowords are used, instead of actual words, to emphasize the visual rather than the semantic aspect of reading. To address invariance to image changes, each pseudoword was presented in 4 different types of font (Arial Black, Comic Sans MS, Courier, and Lucida Handwriting), as illustrated in Figure 1B. The use of different fonts is relevant here in that it introduces a variety of low-level changes in different magnitudes unlike other factors responsible for visual appearance (e.g. the use of upper/lower case is restricted to 2 classes). With respect to orthographic structure, pseudowords were similar in that they were all composed of 5-letter strings (structured CCVCC) but were dissimilar from each other in that they did not share letters in the same position (or within onsets/codas). In addition, participants were also presented with faces. Specifically, they were shown front-view images of 4 unknown young adult male faces (Supplementary Fig. S1)—each face was presented in 4 different versions (displaying different emotional expressions: happy, sad, disgusted, and neutral).

### Procedures

Each subject was scanned for a total of 21 functional runs collected across 3 different sessions to acquire a sufficient number of observations for the purpose of multivariate analysis.

Four of the scans employed a block design (14 s blocks, 933 ms trials). During these runs, participants performed a 1-back task (same/different image as the previous one). Binary responses were made using the index fingers of the 2 hands (right for “same” and left for “different”). Stimulus blocks were separated by 10 s of fixation. An additional 10 s fixation interval was also introduced at the beginning of each run. Any given run contained a total of 10 blocks, 2 for each of the 5 categories: words, false-font strings, faces, houses, and objects. The total duration of a run was 250 s.

The remaining 17 scans employed a widely spaced event-related design. Each trial had the following structure: a bright fixation cross was presented in the middle of the screen for 100 ms, then a stimulus appeared for 400 ms, and a lower-contrast fixation replaced it for 9.5 s until the end of the trial. Participants were instructed to identify each of 32 stimuli (16 images of pseudowords and 16 of faces) at the individual level across changes in font/expression by pushing a button associated with each identity. More precisely, for each subject, the 4 pseudowords (Fig. 1B) were randomly assigned to the fingers of one hand and the 4 facial identities (Supplementary Fig. S1) were assigned to the fingers of the opposite hand. Similarly, the 2 categories, faces and pseudowords, were randomly assigned to the 2 hands. All stimuli were shown exactly once in each scan. The stimulus order was pseudo-randomized so as to maximize the entropy of the sequence with respect to stimulus identity and category under the constraint that no more than 2 stimuli of the same identity could be presented consecutively (Wager and Nichols 2003). The length of a run was 330 s (including 10 s of fixation at the beginning).

Stimuli were presented in the center of the screen against a dark background and subtended a visual angle of  $3.2^\circ \times 4.1^\circ$ . Stimulus presentation and response recording relied on Matlab (Mathworks, Natick, MA, USA) and Psychtoolbox 3.0.8 (Brainard 1997; Pelli 1997).

All subjects were familiarized with the stimuli and practiced each task until identification accuracy reached ceiling (>95%). Additional tests confirmed that subjects maintained this level of performance



**Figure 1.** Experimental orthographic stimuli. (A) Examples of words and false-font strings. (B) Pseudoword stimuli (4 pseudowords  $\times$  4 types of font).

throughout the experiment. No reliable differences were also found in either accuracy or reaction time across different stimulus identities—see Supplementary Material for more details on behavioral procedures and results. As such, our study was aimed at investigating the neural basis of visual processing underlying effortless and (near-) faultless recognition performance.

### **MRI Data Acquisition**

Subjects were scanned in a Siemens Allegra 3T scanner with a single-channel head coil. Functional images were acquired with an echo-planar imaging (EPI) pulse sequence (TR 2 s, TE 31 ms, flip angle 79°, 2.5-mm isotropic voxels, field of view 240 × 240 mm<sup>2</sup>; 27 slices parallel with the AC-PC line covered the ventral cortex of each subject). A  $T_1$ -weighted anatomical image (1-mm<sup>3</sup> voxels; 192 slices of size 256 × 256 mm<sup>2</sup>) was also acquired on each session.

### **Preprocessing and Conventional Univariate Mapping**

All preprocessing and univariate analyses were carried out using AFNI (Cox 1996). Preprocessing of functional data consisted in the following steps: slice scan-time correction, motion correction, co-registration to the same anatomical image, and normalization to percentage of signal change. For the purpose of univariate analysis, data were also smoothed with a Gaussian kernel of 7.5 mm FWHM. No spatial or temporal smoothing was performed on the data previous to multivariate analyses to preserve the high-frequency information in the activation patterns (Swisher et al. 2010).

Conventional univariate analysis was conducted on the data by fitting each type of stimuli with a boxcar predictor, convolving it with a gamma hemodynamic response function and applying a general linear model to estimate voxelwise coefficients for each stimulus category. Statistical maps were computed by pairwise comparisons between different categories (e.g. words versus false-font strings) and corrected for multiple comparisons using the false discovery rate (FDR).

### **Region of Interest Selection**

Multiple regions were localized both at the individual and at the group level by means of univariate analysis. Specifically, 3 different regions were consistently localized across subjects based on estimates of activation elicited by orthographic stimuli relative to fixation ( $q < 0.001$ ). We note that, while extensive areas of visual cortex are activated by such stimuli, the stringent correction for multiple comparisons ensured that only regions with the highest and most robust response were selected by this contrast. For the purpose of region of interest (ROI) analyses, a spherical mask with a 5-voxel radius was placed at the peak of each region in the native space of each subject (see Supplementary Material for more details regarding the design of the masks). A control ROI also located in the early visual cortex (EVC) was constructed by placing a spherical mask at the center of the anatomically defined calcarine sulcus of each subject.

### **Multivariate Mapping**

Multivariate searchlight mapping (Kriegeskorte et al. 2006) was carried out using a cortical mask similar to that deployed for ROI selection and construction. The mask was exhaustively walked voxel-by-voxel across the volume of each subject (restricted to a cortical mask) and was used to constrain the local information available for pattern classification.

### **Category-level Discrimination**

Observations corresponding to each block were constructed by averaging local patterns of activation across time (4 s through 18 s from block onset in order to accommodate the delay of BOLD responses). Binary classification was then applied across observations for each pair of categories using linear support vector machines (SVM) with a trainable  $c$  term— $c$  was optimized by nested cross-validation using a grid search (Chang and Lin 2011). Unbiased discriminability estimates were computed using a leave-one-pair-out cross-validation scheme:

each time, 1 block of either type was left out for testing while the classifier was trained on the remaining blocks. Finally, classifier performance was encoded as  $d'$  sensitivity (Green and Swets 1966) and aggregated into subject-specific information-based maps (Kriegeskorte et al. 2006; Nestor et al. 2011). Specifically, each voxel in a map was labeled with the discrimination performance computed relative to the mask centered on that voxel.

For the purpose of group analysis, all subject-specific maps were normalized to Talairach space, averaged across subjects, and submitted to voxelwise statistical tests—FDR-corrected  $t$  tests against chance ( $d' = 0$ ).

### **Identity-level Discrimination**

Patterns of activation corresponding to each stimulus were constructed by concatenating voxel responses elicited by the corresponding stimulus at 4, 6, and 8 s after its onset (Mourão-Miranda et al. 2007). In this respect, concatenation (as opposed to averaging) ensures the availability of temporal as well as spatial information for classification purposes, a procedure particularly suitable for widely spaced designs (Nestor et al. 2011). As far as temporal masking is concerned, the 4–8 s window was selected to capture the peak of the hemodynamic response function (Friston et al. 1994).

Next, single observations were constructed for each run by averaging activation patterns across presentations of the same identity (e.g. a given pseudoword displayed with different fonts). Pairwise discriminability maps were then computed following a procedure similar to that used for category-level discrimination—in order to extend binary classification to a multi-class case (i.e. 4 different stimulus identities), we used a “one-against-one” procedure by discriminating each class from every other one (see Supplementary Material). Finally, these maps were averaged across pairs (6 pairs corresponding to 4 individual stimuli) to deliver identity-level information-based maps for each type of within-category discrimination (i.e. of pseudowords across font or of fonts across pseudowords).

The procedure above was applied both for mapping purposes (using a searchlight approach) and for ROI analysis (using regions selected based on univariate mapping).

Of note, the possibility that discrimination estimates are overoptimistic due to autocorrelation (i.e. by affecting the similarity of test and training observations) (Pereira and Botvinick 2011) is not of concern here. This is clear in the case of identity-level discrimination since test and training observations of the same category are constructed from different functional runs. Also, this is unlikely for our category-level results since, in our design, different blocks of the same type are separated by at least 34 s (and an intervening block of a different type).

Multivariate analyses were carried out in Matlab 7.12, using the Parallel Processing Toolbox and the SVM LIB 2.88 library for pattern classification (Chang and Lin 2011) running on a ROCKS+ multiserver environment.

### **Multivariate Ranking of Voxel Diagnosticity**

The contribution of each voxel to a given type of discrimination was ranked by means of recursive feature elimination (RFE). This multivariate technique (Guyon et al. 2002) has been previously applied to fMRI data in the attempt to reduce the dimensionality of activation patterns (De Martino et al. 2008) and to map voxel diagnosticity (Hanson and Halchenko 2008). Here, we use it to assess and compare the contribution of a set of voxels (within a given ROI) to independent types of discrimination.

The method works by repeatedly eliminating the feature that is least diagnostic for a given type of classification—diagnosticity here was measured by a common metric, the square of the classification weights computed across features (Hanson and Halchenko 2008). In detail, the method proceeded as follows: (i) a linear SVM classifier was trained on a given feature set, (ii) a ranking score was computed for all features (based on classification weights), (iii) the feature with the smallest rank was eliminated, and (iv) the procedure was repeated until feature depletion. Thus, for any discrimination, the method produces a ranking of classification features from the least diagnostic

(eliminated first) to the most diagnostic (eliminated last). In our case, it produces a ranking of voxel-time features, that is, the ranking of all voxels for any of the 3 given time points (4, 6, and 8 s after stimulus onset). For this reason, additional averaging across time points was carried out resulting in voxel-specific ranking for each type of classification.

In order to compare the contribution of the same voxels for pseudoword and face discrimination, RFE-based ranking was computed separately for the 2 types of classification and the results were then correlated with each other. This analysis was separately applied to each subject and each ROI (more precisely, to each ROI able to support above-chance pseudoword and face discrimination). Finally, correlation values were converted to  $z$ -scores using Fisher's  $z$  transform and were compared against chance across subjects.

Of note, RFE ensures a more robust and reliable ranking than that based on single-pass classifications (e.g. by ordering the weights of a single classification model). However, it is computationally more demanding since classification is performed  $n$  times (where  $n$  is the number of features in the set) instead of just a single time.

## Results

### Category-level Mapping of Orthographic Processing: Univariate Analysis

First, in order to localize regions responsive to orthographic stimuli in the ventral cortex, we contrasted activation for words versus fixation. This approach is commonly used to target regions that respond robustly albeit not necessarily exclusively to orthographic stimuli (Cohen et al. 2000; Dehaene et al. 2002; Glezer et al. 2009). Three clusters were consistently identified by this approach both at the individual level and at the group level following a stringent correction for multiple comparisons ( $q < 0.001$ ). Two of these clusters, located in the left hemisphere, had peaks in the inferior occipital gyrus (IOG) on the border with the posterior fusiform

gyrus and in the inferior frontal gyrus (IFG) pars triangularis, while a third cluster was found in the right IOG. Table 1 shows the coordinates of individually defined ROIs and Figure 2A displays the corresponding results of the group-based mapping.

To relate the location of these regions to the traditional coordinates of the vWFA, we note that the left IOG peak falls within the boundaries of the vWFA ( $-50 < x < -30$ ,  $-80 < y < -30$ ,  $z < 0$ ) (Cohen et al. 2002) in the posterior proximity of the average vWFA peak coordinates ( $x, y, z = -44, -58, -15$ ) (Cohen et al. 2002; Vigneau et al. 2005). Thus, the left IOG ROI is a plausible vWFA candidate. More generally, consistent with our knowledge about the reading network (Turkeltaub et al. 2002; Mechelli et al. 2003; Bolger et al. 2005), the bilateral IOG regions are presumably involved in orthographic processing. As far as the role of the IFG is concerned, this region is known to be involved in lexical and phonological processing (Fiez and Petersen 1998; Jobard et al. 2003). Its activation in relation to orthographic processing is not surprising though, given the importance of the mapping between orthographic and phonological representations for reading and for efficient processing at the level of the VWFA in particular (Brem et al. 2010).


Secondly, we attempted to localize regions that are selective, rather than merely responsive, to orthographic stimuli, namely regions that respond with higher activation levels to orthographic stimuli relative to control categories exhibiting complex and comparable visual structures (for a recent discussion of selectivity and functional specialization, see Kanwisher 2010). Selectivity thus construed is usually assessed by contrasting words with objects (e.g. Szwed et al. 2011), unfamiliar types of script (e.g. Reinke et al. 2008) or other orthographic-like stimuli (e.g. Ben-Shachar et al. 2007). To this goal, our analysis contrasted words and false-font controls. The results found sensitivity to words in bilateral middle temporal (MTG) regions ( $q < 0.05$ ; Figure 2B). The MTG regions involved in reading are commonly found to mediate lexical processing (Jobard et al. 2003). However, no fusiform or occipital regions typically associated with sublexical processing were detected by this contrast. These findings are illustrative of the difficulty of mapping selectivity for orthographic stimuli by means of univariate analysis (Price and Devlin 2003; Wright et al. 2008; Vogel et al. in press).

Thirdly, to verify the replicability and robustness of our mapping results, we performed a separate analysis based on the data from our event-related scans. Specifically, we attempted to locate regions responsive to orthographic stimuli by contrasting pseudowords and fixation in each

**Table 1.**  
Regions responsive to orthographic stimuli

Region	Coordinates			%SC
	$x$	$y$	$z$	
L IOG (W)	-43 ( $\pm 4$ )	-71 ( $\pm 3$ )	-16 ( $\pm 3$ )	1.03 ( $\pm 0.25$ )
L IOG (PW)	-37 ( $\pm 5$ )	-68 ( $\pm 8$ )	-9 ( $\pm 4$ )	1.60 ( $\pm 0.69$ )
R IOG (W)	26 ( $\pm 2$ )	-86 ( $\pm 3$ )	-10 ( $\pm 3$ )	1.17 ( $\pm 0.24$ )
R IOG (PW)	32 ( $\pm 6$ )	-77 ( $\pm 5$ )	-7 ( $\pm 4$ )	1.43 ( $\pm 0.41$ )
L IFG (W)	-36 ( $\pm 3$ )	23 ( $\pm 3$ )	6 ( $\pm 2$ )	0.31 ( $\pm 0.08$ )
L IFG (PW)	-30 ( $\pm 6$ )	20 ( $\pm 4$ )	12 ( $\pm 5$ )	1.32 ( $\pm 0.74$ )

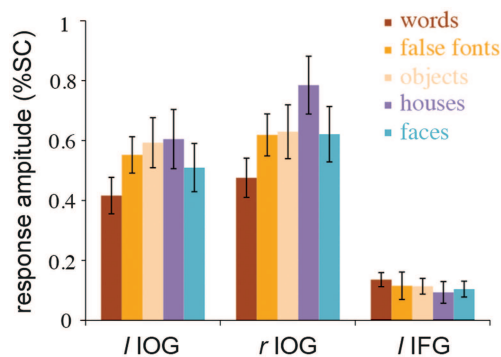
Note: The table shows average Talairach peak coordinates ( $\pm 1$ SD) and response amplitudes in percent signal change (%SC) for individually defined ROIs responsive to words (W) and pseudowords (PW) relative to fixation.



**Figure 2.** Univariate mapping of word processing (group results). (A) Words elicit extensive activation in the bilateral ventral cortex ( $q < 0.05$ ). Activation peaks ( $q < 0.001$ ) are located in the bilateral inferior occipital gyrus (IOG) and the left inferior frontal gyrus (IFG). Crosshairs mark the location of the peaks (left IOG: -41, -69, -16; right IOG: 26, -84, -9; left IFG: -34, 26, 4). (B) Words elicit higher activation than false-font strings in the bilateral middle temporal gyrus (MTG) but not in more ventral occipital/fusiform regions ( $q < 0.05$ ). Crosshairs mark the location of the peaks in the right (51, -31, -1) and left (-49, -36, -1) hemispheres.

subject ( $q < 0.05$ ). This comparison is of interest in that both words and pseudowords have been shown previously to elicit vWFA activation, often of comparable magnitude (Dehaene et al. 2002; Polk and Farah 2002; Vinckier et al. 2007). Critically, our pseudowords were composed of high-frequency letter strings, a factor enhancing the similarity of the activation profiles characteristic of pseudowords and actual words (Vinckier et al. 2007). In agreement with our initial results, we identified bilateral IOG regions in all subjects. A left IFG region was also found in all but 1 subject; however, its location was markedly superior to that of our initial estimate (Table 1).

Fourthly, the response of these individually defined regions to 5 different categories was independently estimated based on the data collected in our block-design runs. Consistent with our initial mapping results, no ROI revealed higher activation to words compared with false fonts (Fig. 3). Also, no ROIs exhibited face selectivity as indicated by higher activation to faces than houses or objects. To further quantify our observations, a 2-way repeated-measures ANOVA (ROI  $\times$  category) was conducted for words and their false-font controls—a separate analysis targeted faces and their control categories (see Supplementary Material). The analysis found



**Figure 3.** Response amplitudes of 3 ROIs involved in orthographic processing. Average BOLD responses are shown for each of 5 different visual categories ( $\pm 1SE$ ).

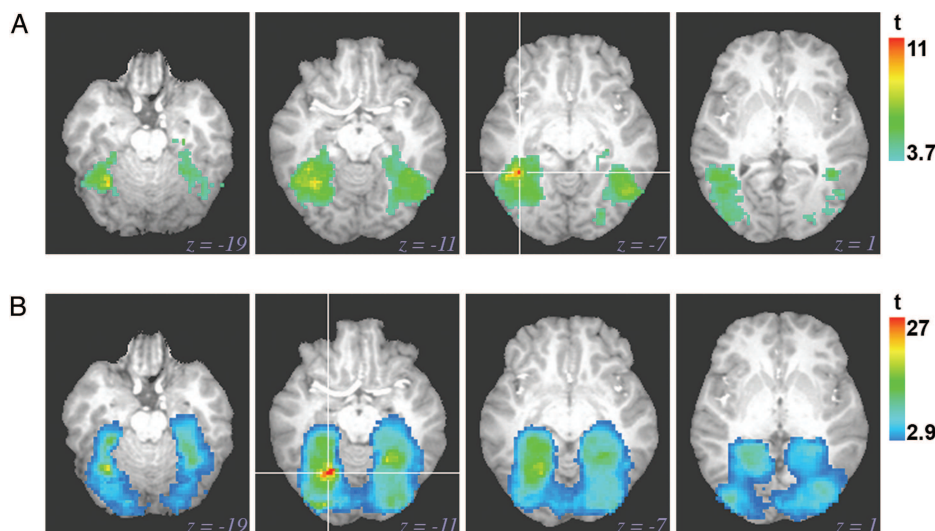
significant main effects for both factors (ROI:  $F_{2,14} = 22.74$ ,  $P < 0.001$ ; category:  $F_{1,7} = 9.09$ ,  $P < 0.05$ ) as well as a significant interaction ( $F_{2,14} = 8.01$ ,  $P < 0.01$ ). Pairwise contrasts (Bonferroni-corrected for multiple comparisons) showed that the left IFG responded less robustly than IOG regions (left IOG:  $t_7 = 4.73$ ,  $P < 0.01$ ; right IOG:  $t_7 = 5.42$ ,  $P < 0.01$ ), but there was no difference between the latter 2 ( $P > 0.10$ ). We also found that words elicited lower instead of higher activation than their visual controls in the left IOG ( $t_7 = 4.16$ ,  $P < 0.05$ ) (for similar findings, see Wang et al. 2011; Vogel et al. in press).

To conclude, univariate analyses were instrumental in locating a network of regions involved in the processing of orthographic stimuli. Overall, the agreement between our 2 mappings (i.e. using words or pseudowords) as well as the ROI analyses associated with them attest to the generality of the present results across different types of orthographic stimuli, different behavioral tasks, and different designs (block and event-related). However, they did not reveal a word-specific advantage at the location of the vWFA. Thus, it appears that the functional specificity of orthographic stimuli, if genuine, is more difficult to capture by univariate analysis than analogous effects in different domains (see Baker et al. 2007, for a similar argument).

#### Category-level Mapping of Orthographic Processing: Multivariate Analysis

Multivariate mapping was performed by estimating the discriminability of words versus false-font strings. Specifically, a spherical searchlight (Kriegeskorte et al. 2006) was walked voxel-by-voxel across the volume of each participant and an unbiased estimate of discrimination was computed at each location.

The outcome of this mapping revealed extensive sensitivity to orthographic stimuli bilaterally in the ventral cortex (Fig. 4A). A large ventral swath including parts of the inferior occipital and fusiform gyri (as well as the ventral ROIs uncovered by our univariate tests) was found in both hemispheres.



**Figure 4.** Category-level information-based maps (group results). Extensive and largely overlapping areas of the bilateral ventral cortex support category-level discrimination of (A) words (relative to false-font strings) and (B) faces (relative to houses). Crosshairs mark the discriminability peak of each map (words:  $-39$ ,  $-46$ ,  $-7$ ; faces:  $-24$ ,  $-61$ ,  $-11$ ).

Interestingly, sensitivity peaked in the ventral cortex in the left fusiform gyrus ( $x, y, z = -39, -46, -7$ ) within the conventional boundaries of the vWFA (Cohen et al. 2002).

To assess the spatial specificity of this mapping, we computed a different type of discrimination using faces and houses. The resulting map revealed a swath of ventral cortex comparable in size and location with that found above (Fig. 4B). However, face information was somewhat more medial than orthographic stimulus information as evidenced by the overall placement of the maps and the position of the peaks. Also, houses, while serving as a conventional control for faces (Kriegeskorte et al. 2006, 2007), are considerably dissimilar from them at the perceptual level, unlike words relative to false-font strings. Thus, the perceptual difference between these 2 types of discrimination makes our ability to localize orthographic information all the more notable.

In summary, unlike univariate mapping, its multivariate counterpart was able to uncover extensive category-level information regarding orthographic stimuli in the ventral cortex. Specifically, this type of mapping revealed that a significant expanse of the ventral cortex, including areas traditionally associated with visual word processing, encodes information able to support the discrimination of orthographic stimuli.

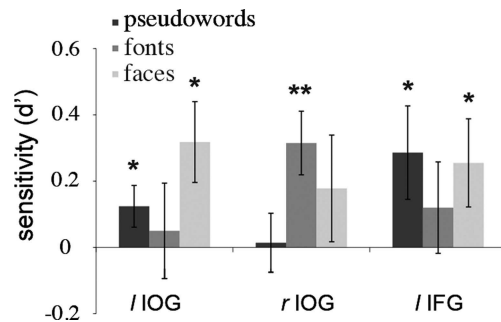
#### Identity-level Discrimination of Pseudowords

A finer-grained, though more challenging, type of discrimination was carried out by comparing orthographic stimuli with each other (rather than with other stimulus categories). Furthermore, to eliminate a potential semantic or lexical basis for discrimination, we used pseudowords (Fig. 1B) instead of actual words. Specifically, we computed identity-level pseudoword discrimination across changes in font.

Multivariate “searchlight” mapping was unable to locate identity-level discriminability even at a liberal threshold ( $q < 0.10$ ). This result is not surprising in that within-category discrimination attempts to capture subtle fine-grained differences (especially when compared with between-category discrimination) and such differences may escape whole-volume mapping following standard corrections for multiple comparisons. However, it is possible that smaller and/or more local effects can be captured by ROI analyses as confirmed by our results below.

Each of 3 individually defined ROIs identified by means of univariate analysis (Table 1) was tested with respect to discrimination performance (see Fig. 5). Specifically, the level of performance of each ROI was compared against chance ( $d' = 0$ ) using 1-group  $t$ -tests. These analyses revealed that both of the left-hemisphere ROIs produced significant levels of discrimination (left IOG:  $t_7 = 1.98, P < 0.05$ ; left IFG:  $t_7 = 2.03, P < 0.05$ ). Unlike its left homologue, the right IOG’s performance did not differ from chance ( $P > 0.10$ ); however, repeated-measures analyses found no significant differences among the 3 ROIs.

Additional analyses examined whether low-level image similarity among stimuli can account for the results above (see Supplementary Material). The absence of significant correlations between estimates of image similarity and discrimination performance for any ROI disconfirmed this hypothesis. The chance-level discrimination found for a control EVC region also suggests that low-level visual representations are



**Figure 5.** Individual-level discrimination for different pseudowords (across font), for different fonts (across pseudoword), and for different facial identities (across expression). Different ROIs show sensitivity to font and pseudoword discrimination but the same left ROIs exhibit sensitivity for both pseudowords and faces (\* $P < 0.05$ , \*\* $P < 0.01$ ).

not sufficient to explain the performance of our left-hemisphere ROIs.

In summary, the results above provide direct evidence for the presence of identity-level sublexical representations within the left ventral cortex at the location of the vWFA.

#### Specificity of Orthographic Processing

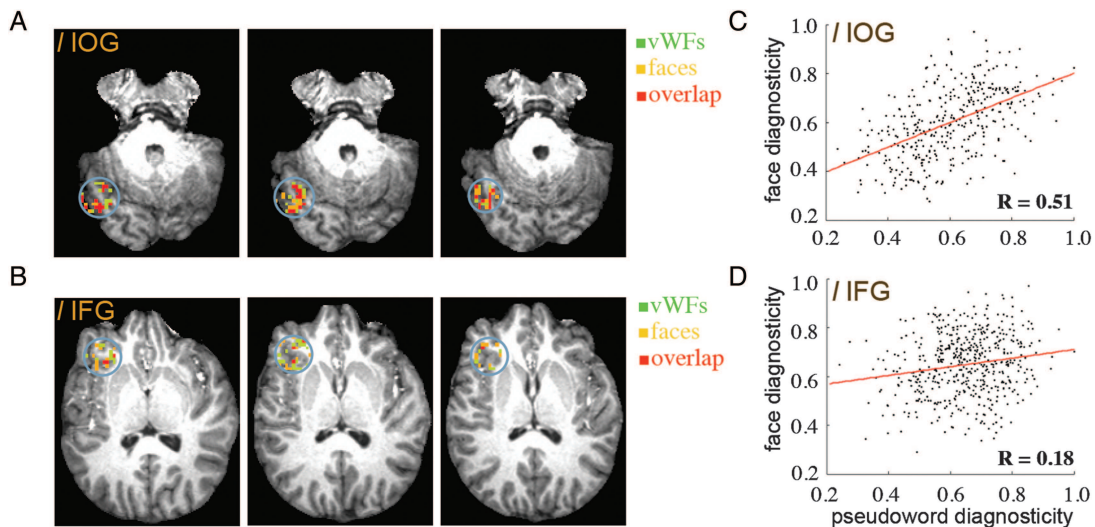
To assess the functional and domain specificity of orthographic processing, we conducted a number of multivariate analyses across the same ROIs.

First, we discriminated different types of fonts across different pseudowords, a type of discrimination orthogonal to that investigated above. This analysis revealed above-chance performance in the right IOG ( $t_7 = 3.19, P < 0.01$ ) but not in the remaining ROIs (Fig. 5).

Secondly, to examine domain specificity, we computed identity-level face discrimination across variation in expression. The results of this analysis (Fig. 5) revealed significant levels of discrimination in left-hemisphere ROIs (left IOG:  $t_7 = 2.60, P < 0.05$ ; left IFG:  $t_7 = 1.92, P < 0.05$ ) but not the right IOG ( $P > 0.10$ ).

Thirdly, the 3 sets of discrimination results (i.e. with respect to pseudowords, fonts and facial identities) were examined by means of a 2-way repeated-measures ANOVA (ROI  $\times$  discrimination type). No significant main effects or interaction were found by this analysis. Thus, although comparisons with chance suggest that pseudoword and font discrimination exploit oppositely lateralized resources, this conclusion would need to be tempered by the absence of an interaction between ROI and type of discrimination (see Discussion).

Fourthly, in order to clarify the relationship between orthographic stimulus and face processing, we correlated voxel-specific diagnosticity scores for the 2 types of discrimination (i.e. between pseudowords and between facial identities). Voxel-specific scores derived by RFE analysis (De Martino et al. 2008) were separately computed for each of the left-hemisphere ROIs shown to support both types of identity-level discrimination. We expect the value of these correlations to be positive if the 2 types of discrimination rely on similar groups of voxels, negative if they rely on distinct sets of voxels and equivalent to chance if they rely on partly overlapping groups of voxels.



**Figure 6.** (A, B) Distribution of individual-level orthographic/face information in a representative subject. The maps show voxels diagnostic for pseudoword discrimination, face discrimination and their overlap (in native space). For either discrimination the top 100 voxels (ranked by recursive feature elimination) are selected and color-labeled separately in 2 ROIs: the left IOG and the left IFG. The maps show extensive overlap in the IOG but not in the IFG. (C, D) Correlation of voxel diagnosticity (across all ROI voxels) for orthographic and face discrimination in the same subject.

Correlation coefficients for the left IOG scored positive values for all subjects (mean  $r = 0.34 \pm 0.13$  SD). A map of diagnostic voxels in a representative participant are shown in Figure 6A and a scatter plot of the 2 types of voxel diagnosticity is shown in Figure 6C. Following conversion to z-scores, correlations were tested against chance ( $r = 0$ ) at the group level and were found to be significantly higher than chance ( $t_7 = 6.97$ ,  $P < 0.001$ ). In contrast, left IFG values (mean  $r = 0.19 \pm 0.30$  SD) were not different from chance ( $P > 0.10$ ; see Figures 6B and D). However, the 2 regions did not perform significantly differently from each other (paired  $t$ -test across participants:  $t_7 = 0.64$ ,  $P > 0.10$ ).

While the results above suggest a close relationship between orthographic and face processing in the vWFA, it is possible that the source of the correlations is more general in nature. In particular, the signal-to-noise ratio (SNR) or the overall responsiveness of the voxels to visual stimuli may factor into the ranking of the voxels independent of domain. For instance, it is possible that more responsive voxels are ranked higher by RFE. If so, the current results would be less informative of a specific relationship between orthographic and face processing.

To test this hypothesis, we estimated both temporal SNR (Murphy et al. 2007) and voxel responsiveness with the goal of factoring out their contribution from RFE rank correlations. Voxelwise temporal SNR was computed as the ratio between signal mean and standard deviation—off-stimulus signal was estimated by regressing out both variables of interest (corresponding to different stimulus types) and nuisance variables (motion and linear trend) (Murphy et al. 2007). Next, voxel responsiveness was computed as the average activation of each voxel in response to 5 different visual categories (presented in our block-design scans) relative to fixation. The correlation of voxel-based orthographic and face diagnosticity was then recomputed while regressing out these 2 estimates. The outcome of this analysis replicated our initial findings: only voxels in the left IOG ( $t_7 = 6.15$ ,  $P < 0.001$ ) were related to each other in terms of diagnosticity for the 2 domains but

correlation values across the 2 ROIs were not significantly different from each other ( $P > 0.10$ ). Thus, the relationship found between orthographic and face processing is unlikely to be accounted for by a simple general signal/response property.

Taken together, these results show that orthographic processing in the ventral cortex may not extend equivalently and uniformly to all aspects of visual orthographic processing. Concretely, (pseudo)word and font discrimination may exploit differentially neural resources in the occipitotemporal cortex of the 2 hemispheres. More importantly, they suggest that visual processing at the level of the vWFA is not specific to orthographic stimuli, but that faces and orthographic stimuli may share a common processing, and representational basis.

## Discussion

The present work investigated the nature and specificity of visual word form processing in the human ventral cortex. This investigation is timely, as an adequate characterization of orthographic processing continues to generate debate. Some have argued that a left ventral cortical region plays a dedicated role in the representation of visual word forms (although the exact form of these representations is also controversial). Others have suggested that this region is responsive to a broader array of inputs and is, therefore, not domain-specific at all.

To elucidate the functional role of the left ventral cortex, we employed a variety of multivariate techniques applied to fMRI data. Critically, we addressed this issue with regard to category and identity word form processing, allowing us to map and characterize the distribution of relevant information at 2 different levels. In addition, we conducted a similar investigation in a parallel domain, namely face recognition, motivated by its similarity in terms of visual expertise (James et al. 2005) and by its potential relationship with word form processing (Dehaene et al. 2010; Plaut and Behrmann 2011).

### **Category-level Mapping of Orthographic Processing**

Standard univariate analyses did not yield evidence of an advantage for words relative to false-font strings in the left fusiform cortex. The absence of this advantage is significant in that it replicates a general result (Price and Devlin 2003; Wright et al. 2008; Vogel et al. in press) and may account for some of the controversy lying at the core of the debate mentioned above. Importantly here, this serves as a first step in our inquiry: can multivariate analysis map visual word form information where univariate analysis fails to find an advantage? And if so, why?

In response to the first question, our results show that multivariate mapping can indeed locate category-level discriminability in the ventral cortex when comparing words with false-font controls. Moreover, the sensitivity peak of this mapping was located at the traditional location of the vWFA (Cohen et al. 2002). Overall, these findings are consistent with a critical role of this area in visual word form processing.

Beyond the presence of category-level discriminability, the current results are notable in 2 other respects. First, unlike typical visual categories compared with each other (e.g. faces and houses) (Haxby et al. 2001; Kriegeskorte et al. 2006), words and false-font controls were designed to be as similar as possible in terms of featural makeup and visual complexity. Thus, large differences in visual appearance could not serve as a basis for discrimination. Secondly, the above-chance discriminability was far more extensively distributed than anticipated: it encompassed a large swath of the bilateral ventral cortex. While semantic, lexical, and/or phonological processing may help our ability to discriminate the 2 categories in the IFG (Fiez and Petersen 1998; Jobard et al. 2003) and the anterior FG (Glezer et al. 2009; Kronbichler et al. 2004), such sources are considerably less plausible to account for IOG or posterior FG activation. Therefore, we conclude that the basis of the discrimination in these posterior areas is likely higher level visual information of a sublexical nature.

As far as the discrepancy between our univariate and multivariate results is concerned, 2 different hypotheses need to be considered. First, as noted above, the absence of an activation advantage for words relative to unfamiliar or false-fonts is not uncommon (Xue and Poldrack 2007; Wang et al. 2011; but see Ben-Shachar et al. 2007; Szwed et al. 2011). Thus, one possibility is that the vWFA, while critical for orthographic stimulus processing, is equally involved in the processing of other visual categories. If so, univariate category effects, when observed, may not reflect representational differences but, instead, may stem from other sources such as the specifics of the experimental design (Starrfelt and Gerlach 2007). At the same time, pattern analyses may accurately capture the involvement of the vWFA in visual word form processing, and may exhibit more robustness across different experimental manipulations. For instance, orthographic stimuli have been shown to elicit robust response patterns at the category level across different experimental tasks (Ma et al. 2011).

Another possibility is that a genuine activation advantage of orthographic stimuli relative to other categories does exist but the standard resolution of fMRI may be too coarse to capture it (Baker et al. 2007). If so, this problem is likely to be aggravated by the spatial smoothing preceding standard univariate analysis (Cohen and Dehaene 2004). In line with this explanation, multivariate analyses that are able to exploit

higher resolution subvoxel information (Swisher et al. 2010) were shown here to uncover extensive sensitivity to orthographic stimuli.

One obstacle in arbitrating between such competing explanations is the scale of these investigations. First, category-level comparisons, while informative, may be too coarse to clarify the precise involvement of ventral areas in visual word form processing. For instance, if the primary function of the vWFA is distinguishing and selecting among different visual word forms, then identity-level comparisons are critical in establishing and characterizing this function. Secondly, whole-ROI analyses suffer from a lack of spatial specificity in that different subregions or, more generally, different populations of neurons may serve different functions. This issue is particularly relevant in the case of visual word form processing given the evidence for an entire spectrum of sensitivity to orthographic stimuli across the left ventral cortex (Cohen and Dehaene 2004; Vinckier et al. 2007; Levy et al. 2008). Our work deals with these issues as discussed below.

### **Identity-level Discrimination of Pseudowords**

Three different constraints guided our investigation. First, discrimination was performed among pseudowords rather than actual words in the attempt to eliminate a lexical or semantic basis for discrimination. Secondly, the analysis was carried out across variations in font, considering that genuine orthographic encoding should exhibit font invariance (McCandliss et al. 2003; Hillis et al. 2005). Thirdly, to align this analysis with traditional fMRI findings, we performed the discrimination within our conventionally defined ROIs.

Our results revealed identity-level discrimination within the left ventral cortex as well as the left IFG. While the basis of the latter is likely phonological processing (Fiez and Petersen 1998; Jobard et al. 2003), the former serves as evidence for visual word form encoding. Importantly, we rule out a low-level image-based account of these findings given the variability of our stimuli (across font) and the absence of an image similarity effect on neural discriminability scores. Thus, the basis of these results appears to be higher level sublexical visual information.

One recent set of findings corroborates our results and their interpretation. Braet et al. (2012) report discriminable patterns of activation in the vWFA for different orthographic stimuli across variation in position and visual appearance. Importantly, in this study, orthographic similarity, but not semantic relationship or lexicality, modulated the size of the correlation between patterns of activation associated with different stimuli. Specifically, the authors found that words and pseudowords with similar orthographic structure (i.e. shared letters in most positions) led to similar patterns of activation while semantically related words (i.e. synonyms) did not elicit similar patterns. Furthermore, these results were specific to the vWFA and were not replicated in other retinotopic or object-selective areas. Of note, this study was considerably different from ours with respect to both experimental manipulations (e.g. design and task) and type of multivariate analyses (e.g. classifier and cross-validation procedure). Thus, identity-level discrimination between different orthographic stimuli appears to capture a genuine difference in neural representation robust across the specifics of experimental manipulation and analysis.



In summary, the findings above confirm an important implication of current theories of visual word form processing (Cohen and Dehaene 2004; Dehaene and Cohen 2011). Specifically, they provide a direct demonstration that the left ventral cortex responds with different neural patterns to different orthographic stimuli and, thus, support the critical role of the left ventral cortex in sublexical word form processing.

### **Functional Specificity of Orthographic Processing**

While the main goal of orthographic stimulus processing is presumably visual word form encoding and discrimination, this does not imply lack of sensitivity to other aspects of orthographic stimuli. In particular, it is of interest to consider whether this sensitivity extends to properties orthogonal to visual word form identity, such as font. An analysis of our 3 functionally localized ROIs found above-chance discrimination in the right IOG but not in our left ROIs for font discrimination (across pseudowords). While the absence of an interaction between ROIs and type of discrimination precludes a strong claim regarding hemispheric asymmetry, the results above are suggestive of different functional priorities for the vWFA and its right homologue.

Interestingly, neuropsychological evidence (Barton et al. 2010) suggests opposite lateralization for the 2 types of processing considered here: visual word form identity in the left vWFA and style (i.e. type of font/handwriting) in its right homologue. Recent neuroimaging results (Qiao et al. 2010) also show that reading handwriting versus printed font boosts activation in the ventral cortex, particularly in the right hemisphere. To account for these results, the authors of the study suggest that the right-hemisphere homologue of the vWFA is involved in writer identification rather than letter or word identification. Overall, our results are consistent with these suggestions: the processing of orthographic stimuli in the 2 hemispheres appears to prioritize different functions.

What is the source of this functional difference? A general explanation could be couched in terms of different modes of processing characterizing the 2 hemispheres. For instance, holistic versus featural processing are advocated as characteristic of the right- versus the left-hemisphere, particularly in the context of face processing (Rossion et al. 2000). Relevantly here, holistic visual processing may rely on global properties such as overall curvature and spacing characteristic of script “style” or font while featural processing would serve a critical role in discriminating fine visual differences essential for visual letter and word identification. To be clear, featural processing here does not imply that words are processed sequentially letter-by-letter, which is certainly not true of normal reading—the fact that words are processed as whole units is evidenced by a number of classical results such as the word superiority effect (Wheeler 1970; McClelland and Johnston 1977). Instead, it suggests that, unlike its right homologue, the vWFA prioritizes the identification of diagnostic letter features to discriminate efficiently letters and words from each other. Future research will be needed to test the validity of this hypothesis as well as the full extent of the asymmetry noted above.

### **Domain Specificity of Visual Word Form Processing**

To address the issue of domain specificity in the ventral cortex, we related visual word form processing to face processing both at the category and identity levels.

As expected, category-level face information was mapped across an extensive portion of the ventral cortex (Haxby et al. 2001; Kriegeskorte et al. 2006) similar to that hosting visual word form information. Overall, this attests to the sensitivity of our multivariate mapping. However, as previously noted, the spatial extent of this sensitivity and the contribution of additional factors such as semantics make its results less instrumental in constraining an interpretation.

Identity-level analyses addressed this concern by comparing pseudoword and face processing relative to finer within-category discriminations both at the ROI and at the voxel level. First, at the ROI level, 2 of the regions examined, in the left IOG and IFG, supported pseudoword as well as face identification. Secondly, at the voxel level, diagnosticity for face and pseudoword discrimination correlated positively within the left IOG but not within the left IFG.

We note that the ability of the left IFG to support face discrimination may appear surprising at first sight given the involvement of this region in phonological and lexical processing (Fiez and Petersen 1998; Jobard et al. 2003; Hauk et al. 2008). However, this result is likely to reflect the silent naming of different facial identities—while the stimuli in the study were not explicitly associated with any names or verbal identifiers, the subjects confirmed using specific labels in order to keep track of the stimuli (e.g. “face one” and “face two”). The absence of a correlation between face and pseudoword diagnosticity can be explained on this basis as a lexicality effect since faces but not orthographic stimuli were labeled with actual words.

More interestingly here, we found that pseudoword and face discrimination were related to each other both at the ROI and at the voxel level in the left IOG. A number of recent empirical and computational results also document this relationship. For instance, Dehaene et al. (2010) found that acquisition of literacy in normal adult populations leads to decreased face responses at the location of the vWFA and increased responses in the right fusiform gyrus. Further examination of these results revealed that orthographic representations compete with face representations by limiting their expansion in the left ventral cortex. Importantly, the competition between orthographic representations and faces was more pronounced than that between orthographic representations and other categories such as houses and tools (see also Cantlon et al. 2011). Another neuroimaging study (Mei et al. 2010) found that a higher vWFA activation was associated with better recognition memory for words and faces alike. Finally, a recent model of orthographic and face processing (Plaut and Behrmann 2011) illustrated the dynamics of this relationship in terms of resource competition and sharing within the left ventral cortex.

Overall, our results are consistent with these previous findings, and suggest that common shape processing underlies both visual word form and face identification within the vWFA. However, what is less clear is the extent of this relationship across cortical neural networks including but not limited to the left ventral cortex. In a previous study, we have found that the right fusiform face area (FFA) supported the 2

types of discrimination (Nestor et al. 2011). In contrast though, other regions critical for face identification (e.g. in the anterior fusiform gyrus) did not support pseudoword identification. Neuropsychological data (Barton et al. 2010) also show that damage to the right ventral cortex leading to face recognition deficits (i.e. prosopagnosia) may be accompanied by deficits of orthographic processing, but not necessarily. Thus, it appears that the networks for orthographic and face processing have a significant overlap but are clearly not identical. Furthermore, regions of overlap are likely to play differential roles within the 2 networks. Specifically, while the vWFA is critical for orthographic processing, it presumably serves only a supporting role in face recognition and the converse holds for the FFA. Future connectivity analyses will be needed to tease apart the relative contribution of these regions concurrently within the networks subserving orthographic and face perception.

Finally, a critical issue concerns the nature of the relationship between orthographic and face perception. This relationship is particularly intriguing, given that words and faces are vastly different in terms of visual appearance and that face processing is associated with heavier reliance upon low-frequency information than orthographic stimuli (Woodhead et al. 2011). Arguably, this relationship goes beyond an expertise-based account involving primarily the right fusiform gyrus (Gauthier et al. 2000). Perhaps, the most promising account in this sense is suggested by the common reliance upon high-acuity central vision of both face and orthographic processing (Levy et al. 2001; Hasson et al. 2002). Specifically, faces along with words and letter strings, unlike other categories such as building and tools, recruit areas of the high-level visual cortex associated predominantly with central rather than peripheral vision. To be clear, this is not at odds with the importance of low-frequency information in face processing (Woodhead et al. 2011) since face recognition can involve a broad range of low and high frequencies (Halit et al. 2006). This common reliance upon central vision along with pressure for left hemispheric language lateralization (Cai et al. 2008, 2010) can, in principle, account for the concurrent involvement of the vWFA and FFA in face and orthographic processing as well as for their differential role within the corresponding networks (Plaut and Behrmann 2011). Quantitative comparisons between model predictions and neural data could clarify in the future the validity and the scope of this explanation.

To conclude, the absence of an activation as well as discrimination advantage for orthographic stimuli relative to other visual categories argues against domain specificity in the left ventral cortex. At the same time, the similar multivariate profile of face and orthographic discrimination within the left ventral cortex provides novel evidence concerning the relationship between these 2 domains.

## Summary

Our investigation demonstrates that orthographic information can be discriminated in the ventral cortex at both the category and identity levels. These findings are instrumental in advancing the debate regarding the nature and specificity of neural orthographic processing. Specifically, our work supports the existence of sublexical orthographic representations within the left ventral cortex. At the same time, it argues against

claims of dedicated circuitry by showing that orthographic processing and face processing rely on common neural resources. More generally, it suggests that multivariate analyses serve as a critical research component in elucidating the neural basis of visual word form encoding.

## Supplementary Material

Supplementary material can be found at: <http://www.cercor.oxfordjournals.org/>.

## Funding

This work was supported by the National Science Foundation (grant number SBE-0542013 to the Temporal Dynamics of Learning Center and grant number BCS0923763 to M.B. and D.P.).

## References

- Baker CI, Liu J, Wald LL, Kwong KK, Benner T, Kanwisher N. 2007. Visual word processing and experiential origins of functional selectivity in human extrastriate cortex. *Proc Natl Acad Sci USA*. 104:9087–9092.
- Barton JJ, Sekunova A, Sheldon C, Johnston S, Iaria G, Scheel M. 2010. Reading words, seeing style: the neuropsychology of word, font and handwriting perception. *Neuropsychologia*. 48:3868–3877.
- Ben-Shachar M, Dougherty RF, Deutsch GK, Wandell BA. 2007. Differential sensitivity to words and shapes in ventral occipito-temporal cortex. *Cereb Cortex*. 17:1604–1611.
- Binder JR, Medler DA, Westbury CF, Liebenthal E, Buchanan L. 2006. Tuning of the human left fusiform gyrus to sublexical orthographic structure. *Neuroimage*. 33:739–748.
- Bolger DJ, Perfetti CA, Schneider W. 2005. Cross-cultural effect on the brain revisited: universal structures plus writing system variation. *Hum Brain Mapp*. 25:92–104.
- Braet W, Wagemans J, Op de Beeck HP. 2012. The visual word form area is organized according to orthography. *Neuroimage*. 59:2751–2759.
- Brainard DH. 1997. The psychophysics toolbox. *Spat Vis*. 10:433–436.
- Brem S, Bach S, Kucian K, Guttorm TK, Martin E, Lyytinen H, Brandeis D, Richardson U. 2010. Brain sensitivity to print emerges when children learn letter-speech sound correspondences. *Proc Natl Acad Sci USA*. 107:7939–7944.
- Cai Q, Lavidor M, Brysbaert M, Paulignan Y, Nazir TA. 2008. Cerebral lateralization of frontal lobe language processes and lateralization of the posterior visual word processing system. *J Cogn Neurosci*. 20:672–681.
- Cai Q, Paulignan Y, Brysbaert M, Ibarrola D, Nazir TA. 2010. The left ventral occipito-temporal response to words depends on language lateralization but not on visual familiarity. *Cereb Cortex*. 20:1153–1163.
- Cantlon JF, Pinel P, Dehaene S, Pelphrey KA. 2011. Cortical representations of symbols, objects, and faces are pruned back during early childhood. *Cereb Cortex*. 21:191–199.
- Chang CC, Lin CJ. 2011. LIBSVM: a library for support vector machines. *ACM Trans Intell Syst Technol*. 2:27.
- Cohen L, Dehaene S. 2004. Specialization within the ventral stream: the case for the visual word form area. *Neuroimage*. 22:466–476.
- Cohen L, Dehaene S, Naccache L, Lehéric S, Dehaene-Lambertz G, Hénaff MA, Michel F. 2000. The visual word form area: spatial and temporal characterization of an initial stage of reading in normal subjects and posterior split-brain patients. *Brain*. 123(Pt 2): 291–307.
- Cohen L, Lehéric S, Chochon F, Lemer C, Rivaud S, Dehaene S. 2002. Language-specific tuning of visual cortex? Functional properties of the visual Word Form Area. *Brain*. 125:1054–1069.

- Coltheart M. 1981. The MRC psycholinguistic database. *Q J Exp Psychol A*. 33:497–505.
- Cox RW. 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput Biomed Res*. 29:162–173.
- Dehaene S, Cohen L. 2011. The unique role of the visual word form area in reading. *Trends Cogn Sci*. 15:254–262.
- Dehaene S, Cohen L, Sigman M, Vinckier F. 2005. The neural code for written words: a proposal. *Trends Cogn Sci*. 9:335–341.
- Dehaene S, Le Clec'h G, Poline JB, Le Bihan D, Cohen L. 2002. The visual word form area: a prelexical representation of visual words in the fusiform gyrus. *Neuroreport*. 13:321–325.
- Dehaene S, Naccache L, Cohen L, Bihan DL, Mangin JF, Poline JB, Rivière D. 2001. Cerebral mechanisms of word masking and unconscious repetition priming. *Nat Neurosci*. 4:752–758.
- Dehaene S, Pegado F, Braga LW, Ventura P, Nunes Filho G, Jobert A, Dehaene-Lambertz G, Kolinsky R, Morais J, Cohen L. 2010. How learning to read changes the cortical networks for vision and language. *Science*. 330:1359–1364.
- De Martino F, Valente G, Staeren N, Ashburner J, Goebel R, Formisano E. 2008. Combining multivariate voxel selection and support vector machines for mapping and classification of fMRI spatial patterns. *Neuroimage*. 43:44–58.
- Fiez JA, Petersen SE. 1998. Neuroimaging studies of word reading. *Proc Natl Acad Sci USA*. 95:914.
- Friston KJ, Jezzard P, Turner R. 1994. Analysis of functional MRI time-series. *Hum Brain Mapp*. 1:153–171.
- Gauthier I, Skudlarski P, Gore JC, Anderson AW. 2000. Expertise for cars and birds recruits brain areas involved in face recognition. *Nat Neurosci*. 3:191–197.
- Glazer LS, Jiang X, Riesenhuber M. 2009. Evidence for highly selective neuronal tuning to whole words in the “visual word form area”. *Neuron*. 62:199–204.
- Green DM, Swets JA. 1966. *Signal detection theory and psychophysics*. New York: John Wiley and Sons.
- Guyon I, Weston J, Barnhill S, Vapnik V. 2002. Gene selection for cancer classification using support vector machines. *Mach Learn*. 46:389–422.
- Halit H, de Haan M, Schyns PG, Johnson MH. 2006. Is high-spatial frequency information used in the early stages of face detection? *Brain Res*. 1117:154–161.
- Hanson SJ, Halchenko YO. 2008. Brain reading using full brain support vector machines for object recognition: there is no “face” identification area. *Neural Comput*. 20:486–503.
- Hasson U, Levy I, Behrmann M, Hendler T, Malach R. 2002. Eccentricity bias as an organizing principle for human high-order object areas. *Neuron*. 34:479–490.
- Hauk O, Davis MH, Pulvermüller F. 2008. Modulation of brain activity by multiple lexical and word form variables in visual word recognition: A parametric fMRI study. *Neuroimage*. 42:1185–1195.
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*. 293:2425–2430.
- Hillis AE, Newhart M, Heidler J, Barker P, Herskovits E, Degaonkar M. 2005. The roles of the “visual word form area” in reading. *Neuroimage*. 24:548–559.
- James KH, James TW, Jobard G, Wong AC, Gauthier I. 2005. Letter processing in the visual system: different activation patterns for single letters and strings. *Cogn Affect Behav Neurosci*. 5:452–466.
- Jobard G, Crivello F, Tzourio-Mazoyer N. 2003. Evaluation of the dual route theory of reading: a meta-analysis of 35 neuroimaging studies. *Neuroimage*. 20:693–712.
- Kanwisher N. 2010. Functional specificity in the human brain: a window into the functional architecture of the mind. *Proc Natl Acad Sci USA*. 107:11163–11170.
- Kriegeskorte N, Formisano E, Sorger B, Goebel R. 2007. Individual faces elicit distinct response patterns in human anterior temporal cortex. *Proc Natl Acad Sci USA*. 104:20600–20605.
- Kriegeskorte N, Goebel R, Bandettini P. 2006. Information-based functional brain mapping. *Proc Natl Acad Sci USA*. 103:3863–3868.
- Kronbichler M, Hutzler F, Wimmer H, Mair A, Staffen W, Ladurner G. 2004. The visual word form area and the frequency with which words are encountered: evidence from a parametric fMRI study. *Neuroimage*. 21:946–953.
- Levy I, Hasson U, Avidan G, Hendler T, Malach R. 2001. Center-periphery organization of human object areas. *Nat Neurosci*. 4:533–539.
- Levy J, Pernet C, Treserras S, Boulanouar K, Berry I, Aubry F, Demonet JF, Celsis P. 2008. Piecemeal recruitment of left-lateralized brain areas during reading: a spatio-functional account. *Neuroimage*. 43:581–591.
- Ma L, Jiang Y, Bai J, Gong Q, Liu H, Chen HC, He S, Weng X. 2011. Robust and task-independent spatial profile of the visual word form activation in fusiform cortex. *PLoS ONE*. 6:e26310.
- McCandliss BD, Cohen L, Dehaene S. 2003. The visual word form area: expertise for reading in the fusiform gyrus. *Trends Cogn Sci*. 7:293–299.
- McClelland JL, Johnston JC. 1977. The role of familiar units in perception of words and nonwords. *Percept Psychophys*. 22:249–261.
- Mechelli A, Gorno-Tempini ML, Price CJ. 2003. Neuroimaging studies of word and pseudoword reading: consistencies, inconsistencies, and limitations. *J Cogn Neurosci*. 15:260–271.
- Mei L, Xue G, Chen C, Xue F, Zhang M, Dong Q. 2010. The “visual word form area” is involved in successful memory encoding of both words and faces. *Neuroimage*. 52:371–378.
- Mourão-Miranda J, Friston KJ, Brammer M. 2007. Dynamic discrimination analysis: a spatial-temporal SVM. *Neuroimage*. 36:88–99.
- Murphy K, Bodurka J, Bandettini PA. 2007. How long to scan? The relationship between fMRI temporal signal to noise ratio and necessary scan duration. *Neuroimage*. 34:565–574.
- Nestor A, Plaut DC, Behrmann M. 2011. Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. *Proc Natl Acad Sci USA*. 108:9998–10003.
- Pelli DG. 1997. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis*. 10:437–442.
- Pereira F, Botvinick M. 2011. Information mapping with pattern classifiers: a comparative study. *Neuroimage*. 56:476–496.
- Plaut DC, Behrmann M. 2011. Complementary neural representations for faces and words: a computational exploration. *Cogn Neuropsychol*. 28:251–275.
- Polk TA, Farah MJ. 2002. Functional MRI evidence for an abstract, not perceptual, word-form area. *J Exp Psychol General*. 131:65–72.
- Price CJ, Devlin JT. 2011. The interactive account of ventral occipito-temporal contributions to reading. *Trends Cogn Sci*. 15:246–253.
- Price CJ, Devlin JT. 2003. The myth of the visual word form area. *Neuroimage*. 19:473–481.
- Qiao E, Vinckier F, Szwed M, Naccache L, Valabrègue R, Dehaene S, Cohen L. 2010. Unconsciously deciphering handwriting: subliminal invariance for handwritten words in the visual word form area. *Neuroimage*. 49:1786–1799.
- Rauschecker AM, Bowen RF, Perry LM, Kevan AM, Dougherty RF, Wandell BA. 2011. Visual feature-tolerance in the reading network. *Neuron*. 71:941–953.
- Reinke K, Fernandes M, Schwindt G, O’Craven K, Grady CL. 2008. Functional specificity of the visual word form area: general activation for words and symbols but specific network activation for words. *Brain Lang*. 104:180–189.
- Rossion B, Dricot L, Devolder A, Bodart JM, Crommelinck M, De Gelder B, Zootjes R. 2000. Hemispheric asymmetries for whole-based and part-based face processing in the human fusiform gyrus. *J Cogn Neurosci*. 12:793–802.
- Schurz M, Sturm D, Richlan F, Kronbichler M, Ladurner G, Wimmer H. 2010. A dual-route perspective on brain activation in response to visual words: evidence for a length by lexicality interaction in the visual word form area (VWFA). *Neuroimage*. 49:2649–2661.
- Spiridon M, Kanwisher N. 2002. How distributed is visual category information in human occipito-temporal cortex?. An fMRI study. *Neuron*. 35:1157–1165.
- Starrfelt R, Gerlach C. 2007. The visual what for area: words and pictures in the left fusiform gyrus. *Neuroimage*. 35:334–342.

- Swisher JD, Gatenby JC, Gore JC, Wolfe BA, Moon CH, Kim SG, Tong F. 2010. Multiscale pattern analysis of orientation-selective activity in the primary visual cortex. *J Neurosci*. 30:325–330.
- Szwed M, Dehaene S, Kleinschmidt A, Eger E, Valabrègue R, Amadon A, Cohen L. 2011. Specialization for written words over objects in the visual cortex. *Neuroimage*. 56:330–344.
- Turkeltaub PE, Eden GF, Jones KM, Zeffiro TA. 2002. Meta-analysis of the functional neuroanatomy of single-word reading: method and validation. *Neuroimage*. 16:765–780.
- Van Doren L, Dupont P, De Grauwe S, Peeters R, Vandenberghe R. 2010. The amodal system for conscious word and picture identification in the absence of a semantic task. *Neuroimage*. 49:3295–3307.
- Vigneau M, Jobard G, Mazoyer B, Tzourio-Mazoyer N. 2005. Word and non-word reading: what role for the Visual Word Form Area?. *Neuroimage*. 27:694–705.
- Vinckier F, Dehaene S, Jobert A, Dubus JP, Sigman M, Cohen L. 2007. Hierarchical coding of letter strings in the ventral stream: dissecting the inner organization of the visual word-form system. *Neuron*. 55:143–156.
- Vogel AC, Petersen SE, Schlaggar BL. Forthcoming. The left occipito-temporal cortex does not show preferential activity for words. *Cereb Cortex*. in press.
- Wager TD, Nichols TE. 2003. Optimization of experimental design in fMRI: a general framework using a genetic algorithm. *Neuroimage*. 18:293–309.
- Wang X, Yang J, Shu H, Zevin JD. 2011. Left fusiform BOLD responses are inversely related to word-likeness in a one-back task. *Neuroimage*. 55:1346–1356.
- Wheeler DD. 1970. Processes in word recognition. *Cognit Psychol*. 1:59–85.
- Woodhead ZV, Wise RJ, Sereno M, Leech R. 2011. Dissociation of sensitivity to spatial frequency in word and face preferential areas of the fusiform gyrus. *Cereb Cortex*. 21:2307–2312.
- Wright ND, Mechelli A, Noppeney U, Veltman DJ, Rombouts SA, Glensman J, Haynes JD, Price CJ. 2008. Selective activation around the left occipito-temporal sulcus for words relative to pictures: individual variability or false positives? *Hum Brain Mapp*. 29:986–1000.
- Xue G, Poldrack RA. 2007. The neural substrates of visual perceptual learning of words: implications for the visual word form area hypothesis. *J Cogn Neurosci*. 19:1643–1655.

## Supplementary Methods

### Behavioral procedures and results

Stimulus familiarization and task practice were carried out prior to the experiment on a non-scanning day (along with subject screening) and on scanning days at the beginning of each experimental session. Each time subjects were presented with 32 pseudoword and face stimuli and were instructed how to map the stimuli to 8 different keys as a function of stimulus identity (e.g. facial identity). Task practice followed largely the same procedure described for the main experiment except that trials ended as soon as the subject made a response. Also, each stimulus was presented 6 times during a practice block (instead of a single time during scans). Practice blocks were repeated until identification accuracy surpassed 95% (usually one or two blocks sufficed).

Behavioral markers of performance (accuracy and RT) were also collected during scanning – these data were not recorded for one subject due to technical difficulties.

Overall identification accuracy across participants reached ceiling: 99% ( $\pm 1\%$  SD). With respect to pseudoword stimuli, no difference was found for different identities either in accuracy ( $F_{4, 18} = 0.99$ ,  $p > 0.10$ ) or in RT ( $F_{4, 18} = 0.27$ ,  $p > 0.10$ ). Similar null results were found for face stimuli. These results are instrumental in confirming that pattern discrimination of stimulus identity is not based on differential levels of difficulty associated with the recognition of different stimuli.

## **Design**

Two types of design were deployed during each experimental session. A block design was used mainly for the purpose of category-level functional localization while a widely spaced event-related design was used in order to test individual-level processing of orthographic stimuli and faces. One of the benefits of a widely spaced design is that it provides us with temporal information regarding the shape and the characteristics of the hemodynamic response. Such information would be largely missed by a block design (which mainly aims to capture the response amplitude associated with an entire group of stimuli) as well as by a fast event-related design (e.g. due to deconvolution with a suboptimal estimate of the hemodynamic response function). Thus, by providing access to the raw (i.e. non-deconvolved) time course of each stimulus, a widely spaced design appears better suited for the use of spatiotemporal pattern classification (Nestor et al., 2011). This method involves concatenation of patterns at different time points (instead of their averaging) with the purpose of allowing classification to exploit both spatial and temporal information (Mourão-Miranda et al., 2007). In sum, the widely spaced design aims to boost the sensitivity of our analyses in the attempt to capture subtle effects associated with individual orthographic and face representations.

The considerations above also account for our use of conventional (i.e. average-based) pattern construction for the purpose of multivariate analyses with our block-design data. The application of spatiotemporal analyses to such data would provide little benefit (Mourão-Miranda et al., 2007) and would also be impractical due to the number of time

points to be considered (i.e. 8 or more). Specifically, the use of concatenation (instead of averaging) would lead to significantly higher pattern dimensionality and, thus, to longer training times and more overfitting. Thus, the two designs have largely different goals and our analyses are adjusted to reflect their different characteristics.

### **ROI masks**

Spherical masks with a radius of 5 voxels were used both for the purpose of ‘searchlight’ whole-volume analysis (Kriegeskorte et al., 2006) and for that of ROI analyses. The overall size of a mask was 515 voxels (for a total volume of 8, 047 mm<sup>3</sup>). The size of the mask was selected so as to balance a number of different and opposing goals: the size trades off amount of information versus spatial specificity and classification reliability. Specifically, the larger the size of the mask, the more (potentially relevant) information is included in the analysis. At the same time, the source of relevant information is less precisely spatially localized for larger masks and classification is more prone to overfitting due to the increase in dimensionality (i.e. due to the inclusion of more voxels) (Pereira et al., 2009). In order to estimate the optimal size of the mask, our analyses (conducted across 3 experimental session with a single pilot subject) varied the radius of the mask from 2 to 6 voxels (in steps of one voxel) and applied multivariate searchlight analysis to category-level discrimination (e.g. faces vs houses). Overall, we found that the levels of discrimination asymptote for masks with a 5-voxel radius. Future investigations will be instrumental in confirming these findings and in determining if / how the optimal size of the mask varies as a function of the effect

investigated and the type of classification procedure – for instance, the optimal size of the mask is likely to depend on how well the classifier scales with pattern dimensionality. For the purpose of the analyses reported in the present study, we used our initial explorations to select the size of the masks in a non-arbitrary fashion.

### **Multi-class pattern classification**

BOLD activation patterns were classified using a ‘one-against-one’ procedure. The term ‘one-against-one’ refers to a common extension of binary pattern classification (e.g. as accomplished by traditional SVM) to multi-class classification (Knerr et al., 1990). This procedure involves constructing a separate classifier for each pair of classes rather than a single classifier per class – the latter strategy involves discriminating a given class from all other classes at the same time (‘one-against-all’). ‘One-against-one’ and ‘one-against-all’ produce in most cases comparable results but the former is associated with shorter training times (Hsu and Lin, 2002). In our case, ‘one-against-one’ involves discriminating each of four different stimulus identities (e.g. every pseudoword) separately from every other one.

### **Low-level image similarity and its role in discrimination**

Distances were computed across pairs of pseudoword images using an L2 (Euclidean) metric, a common measure of low-level image similarity (Moon and Phillips, 2001). More precisely, each image of a pseudoword (i.e. displayed with different types of font) was compared with every image of another pseudoword and an average estimate



was computed across all pairings. Thus, a single estimate was produced for each pair of different pseudowords.

If low-level similarity is the primary source of information exploited by classification, we expect higher discrimination estimates for pairs of pseudowords more dissimilar from each other. In order to test this possibility, the estimates above were correlated with the corresponding discrimination levels for each pair of pseudowords separately for each ROI and each participant. The correlations scores obtained were converted to z-scores using Fisher's z transform and were compared against chance across participants.

## **Supplementary Results**

### **Univariate mapping of object and face-selective areas**

To assess the overall effectiveness of our block-design data for mapping purposes, we examined the results of additional pairwise contrasts. First, a contrast of objects versus faces showed an activation advantage for objects in the ventral cortex in the lateral occipital cortex (Grill-Spector et al., 2001). And second, a contrast between faces and houses revealed the standard face-selective regions of the ventral cortex (e.g. the bilateral fusiform face area) – univariate and multivariate analyses of these areas have been detailed elsewhere (Nestor et al., 2011).

### **Univariate ROI analysis**

A two-way repeated-measures ANOVA (ROI x category) was conducted across individually-defined ROIs for faces, houses and objects (see Fig. 3 for average response amplitudes). These categories were analyzed independently of words and false-font strings due to their vastly different visual structure and due to the absence of low-level property matching – the analysis of the latter two categories is detailed in the main text.

The analysis found significant main effects for both factors (ROI:  $F_{2,14}=25.23$ ,  $p<0.001$ ; category:  $F_{2,14}=4.30$ ,  $p<0.05$ ) and a significant interaction ( $F_{4,28}=6.34$ ,  $p<0.01$ ). Pairwise contrasts (Bonferroni-corrected) showed that the left IFG responded less robustly than IOG regions (left IOG:  $t_7=4.77$ ,  $p<0.01$ ; right IOG:  $t_7=5.59$ ,  $p<0.01$ ) but there was no difference between the latter two ( $p>0.10$ ). Also, the comparison among different categories revealed that houses elicited higher activation than faces in the right IOG ( $t_7=4.42$ ,  $p<0.05$ ). No other contrasts were significant among the three categories tested ( $p>0.10$ ).

### **Contribution of low-level image similarity to identity-level discrimination**

Our identity-level analyses were aimed at discriminating among invariant representations of visual word forms. This constraint was enforced by using one common and considerable source of image variability: font type. However, it is possible that discrimination was still able to exploit certain low-level image properties diagnostic of stimulus identity.

To examine this possibility, we computed both pseudoword and font discrimination in a control region of the early visual cortex (EVC). This region should be

able to take advantage of diagnostic low-level properties for the purpose of linear discrimination; critically though, it has no direct access to invariant visual properties derived through nonlinear processing and characteristic of higher-level cortical regions (Kamitani and Tong, 2005). For either type of discrimination (i.e. pseudoword and font) our analysis found that performance was at chance level ( $p > 0.10$ ).

Another test of the contribution of low-level similarity explored whether discrimination in our target ROIs exploited gross image similarity. To examine this, we computed overall image distances for each pair of pseudowords and we correlated them with their corresponding discriminability values separately for each ROI. If low-level image differences drive our classification results, we expect image distances and performance estimates to be negatively correlated with each other. However, none of these correlations produced significant results for any of the ROIs ( $p > 0.10$ ). A similar result was obtained for font discrimination.

In conclusion, the considerations above suggest that the basis of our discrimination results relies on higher-level representations that are (at least partly) invariant to low-level image properties.

## References

Grill-Spector K, Kourtzi Z, Kanwisher N. 2001. The lateral occipital complex and its role in object recognition. *Vision Res.* 41:1409-1422.

Hsu CW, Lin CJ. 2002. A comparison of methods for multi-class support vector machines . *IEEE Transactions on Neural Networks.*:415-425.

Kamitani Y, Tong F. 2005. Decoding the visual and subjective contents of the human brain. *Nat Neurosci.* 8:679-685.

Knerr S, Personnaz L, Dreyfus G (1990) Single-layer learning revisited: a stepwise procedure for building and training a neural network. In: *Neurocomputing: Algorithms, Architectures and Applications* (Fogelman-Soulie F, Herault J, eds), pp41-50. Springer-Verlag.

Kriegeskorte N, Goebel R, Bandettini P. 2006. Information-based functional brain mapping. *Proc Natl Acad Sci U S A.* 103:3863-3868.

Moon H, Phillips PJ. 2001. Computational and performance aspects of PCA-based face-recognition algorithms. *Perception.* 30:303-321.

Mourão-Miranda J, Friston KJ, Brammer M. 2007. Dynamic discrimination analysis: a spatial-temporal SVM. *Neuroimage.* 36:88-99.

Nestor A, Plaut DC, Behrmann M. 2011. Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. *Proc Natl Acad Sci U S A.* 108:9998-10003.

Pereira F, Mitchell T, Botvinick M. 2009. Machine learning classifiers and fMRI: a tutorial overview. *Neuroimage.* 45:S199-S209.

## Supplementary figure caption

Supplementary Figure 1. Experimental face stimuli (four individuals X four expressions). Original face images courtesy of the Face-Place Face Database Project (<http://www.face-place.org/>) copyright 2008, Michael J. Tarr (funding provided by NSF Award 0339122).

