Journal Articles

Donald and Barbara Zucker School of Medicine
Academic Works

2014

Exemplar Selectivity Reflects Perceptual Similarities in the Human Fusiform Cortex

I. Davidesco

E. Zion-Golumbic

S. Bickel
Northwell Health

M. Harel

D. M. Groppe
Northwell Health

See next page for additional authors

Recommended Citation

Authors

I. Davidesco, E. Zion-Golumbic, S. Bickel, M. Harel, D. M. Groppe, C. J. Keller, C. A. Schevon, G. M. McKhann, R. R. Goodman, G. Goelman, C. E. Schroeder, A. D. Mehta, and R. Malach

Exemplar Selectivity Reflects Perceptual Similarities in the Human Fusiform Cortex

Ido Davidesco[1,2], Elana Zion-Golumbic[3,5], Stephan Bickel[6], Michal Harel[7], David M. Groppe[6], Corey J. Keller[6,8], Catherine A. Schevon[4], Guy M. McKhann[4], Robert R. Goodman[4], Gadi Goelman[9], Charles E. Schroeder[3,5], Ashesh D. Mehta[6] and Rafael Malach[7]

[1]Interdisciplinary Center for Neural Computation, [2]Edmond and Lily Safra Center for Brain Sciences, Hebrew University of Jerusalem, Jerusalem, Israel, [3]Department of Psychiatry, [4]Department of Neurology, Columbia University College of Physicians and Surgeons, New York, NY, USA, [5]Cognitive Neuroscience and Schizophrenia Program, Nathan Kline Institute, Orangeburg, NY, USA, [6]Department of Neurosurgery, Hofstra North Shore LIJ School of Medicine and Feinstein Institute for Medical Research, Manhasset, NY, USA, [7]Department of Neurobiology, Weizmann Institute of Science, Rehovot, Israel, [8]Department of Neuroscience, Albert Einstein College of Medicine, Bronx, NY, USA, and [9]Medical Biophysics, Hadassah Hebrew University Medical Center, Jerusalem, Israel

Address correspondence to Prof. Rafael Malach, Department of Neurobiology, Weizmann Institute of Science, Rehovot 76100, Israel.
Email: rafi.malach@gmail.com

While brain imaging studies emphasized the category selectivity of face-related areas, the underlying mechanisms of our remarkable ability to discriminate between different faces are less understood. Here, we recorded intracranial local field potentials from face-related areas in patients presented with images of faces and objects. A highly significant exemplar tuning within the category of faces was observed in high-Gamma (80–150 Hz) responses. The robustness of this effect was supported by single-trial decoding of face exemplars using a minimal ($n = 5$) training set. Importantly, exemplar tuning reflected the psychophysical distance between faces but not their low-level features. Our results reveal a neuronal substrate for the establishment of perceptual distance among faces in the human brain. They further imply that face neurons are anatomically grouped according to well-defined functional principles, such as perceptual similarity.

Keywords: ECoG, face perception, high-gamma, perceptual similarity

Introduction

A major advance in our understanding of human visual perception has been the discovery of category-selective regions, and in particular face-selective ones, in high-order visual cortex of both human and nonhuman primates (Kanwisher et al. 1997; McCarthy et al. 1997; Ishai et al. 1999; Hasson et al. 2003; Tsao et al. 2006). In nonhuman primates, it has recently been shown that these "face patches" are capable of discriminating between various face exemplars (Tsao et al. 2006; Freiwald et al. 2009). Paradoxically, although human face recognition is superb, the evidence for such exemplar selectivity in the human brain has been minimal and based mainly on fMR-adaptation (Rotshtein et al. 2004; Gilaie-Dotan and Malach 2007; Dricot et al. 2008; Gilaie-Dotan et al. 2010; Malach 2012), which is an indirect measure of selectivity (Sawamura et al. 2006; Mur et al. 2010). Other studies using more direct univariate and multivariate analyses on face-exemplar differentiation in the fusiform face area (FFA) yielded contradicting results (Kriegeskorte et al. 2007; Op de Beeck et al. 2010; Nestor et al. 2011; Mur et al. 2012).

Moreover, the relationship between such neuronal "face tuning" and perception is unclear. Thus, it is presently unknown whether perceptual metrics—i.e., the likeness and dissimilarity among different face exemplars, are coded in the neural responses. A hint that neuronal face tuning may indeed be related to perceptual metrics has been provided for

the category of abstract shapes in a few imaging and single unit studies. These studies reported a link between neuronal tuning on the one hand and perceptual and physical shape similarity on the other hand (Op de Beeck et al. 2001, 2008; Haushofer et al. 2008; Drucker and Aguirre 2009). However, objects have a wide range of different physical shapes compared with the highly uniform category of faces. Thus, it is not clear to what extent principles relevant for diverse sets of abstract shapes can be extrapolated to the more physically homogeneous face domain.

In the present study, we addressed these questions using electrocorticographic (ECoG) recordings in epilepsy patients monitored for presurgical evaluation. We have previously found that such recordings show superior selectivity (and hence spatial resolution) compared with fMRI (Privman et al. 2007), and therefore, ECoG recordings may be more suitable for studying within-category tuning. Previous ECoG studies indeed found clear category-selective responses in high-order visual cortex but have not shown discrimination between different exemplars within a certain category (Allison et al. 1994, 1999; Liu et al. 2009). Most of these studies have focused on "evoked" activity, time-locked to the stimulus. Here we extend this approach by examining "induced" activity, which jitters in latency between trials. Induced ECoG power at various frequencies—and particular high-frequency (Broadband Gamma) bands—provides an informative index to the underlying population firing rates (Kreiman et al. 2006; Nir et al. 2007; Manning et al. 2009; Burns et al. 2010) and is tightly linked to the perceptual states of the observers (Mukamel et al. 2005; Fisch et al. 2009).

Our results reveal robust and consistent exemplar selectivity in face-selective electrodes placed over ventral temporal cortex (fusiform gyrus). Importantly, the results show that the "neural" discriminability of individual faces reflected their "perceptual" distinctness but not their low-level image features. Thus, our results uncover an organizing principle in human face areas that could explain why certain faces appear similar while others are perceived as more distinct.

Materials and Methods

Subjects and Recordings

Fourteen patients (5 males, 13 right handed, age 29.1 ± 2.7 years (mean $\pm$ SEM), see Table 1) with pharmacologically intractable epilepsy, monitored for presurgical evaluation, participated in the

present study. The recordings were conducted at the Long Island Jewish Medical Center and at Columbia University College of Physicians and Surgeons, NY, USA. All patients provided fully informed consent according to the US National Institute of Health guidelines, as monitored by the local institutional review boards.

Each patient was implanted with 84–187 intracranial electrodes for 5–10 days. The electrodes were arranged in subdural grids, strips, and/or depth arrays (Integra Lifesciences Corp.). In the subdural grids and strips, each recording site was 2 mm in diameter with 1 cm separation, whereas in the depth electrodes each recording site was 1 mm in diameter with 0.5 cm separation. The location and number of electrodes were based solely on clinical criteria. The signals were filtered electronically between 0.1 Hz and 1 kHz and sampled at a rate of 2 kHz (XLTEK EMU 128 LTM System). A strip electrode screwed into the frontal bone near the bregma was used as common mode ground and reference. Stimulus-triggered electrical pulses were recorded along with the ECoG data for precise synchronization of the stimuli with the neural responses.

The recordings were conducted at the patients' quiet bedside. Stimuli were presented via a standard LCD screen and keyboard responses were recorded for measurement of behavioral performance.

Stimulus Presentation
The patients viewed grayscale digital photographs of faces, man-made tools, buildings, and geometric patterns (of ~15° × 15° visual angle), which were superimposed with a small white fixation dot. The images were presented for 250 ms in pseudorandom order at a rate of 1 Hz, while the patients performed a 1-back memory task (i.e., pressing a mouse button each time a specific image repeated twice in a row). Stimulus repetitions were infrequent (~10% of the trials) and were mainly used to keep the patient alert. The image set contained 14 different faces, 10 houses, 10 tools, and 5 patterns. Each exemplar was presented 6 times throughout the experiment.

Quantitative Definitions of Electrode Responses
(A) Visually responsive electrodes: Following Fisch et al. (2009), these electrodes were defined as those with short- to mid-latency responses (up to 250-ms poststimulus onset). Latency was defined as the time point in which high-Gamma (80–150 Hz) power first became significantly greater than its prestimulus baseline value if the response remained significant for at least 15 successive time points (Fisch et al. 2009).
(B) Category-selective electrodes: These electrodes were defined as follows: 4 exemplars were selected from each category. For each electrode, the mean area under the curve (AUC) of the high-Gamma response was computed in a time window of 50–350 ms poststimulus. Next, the response for each category was examined and electrodes for which at least one of the exemplars evoked a

significantly higher response compared with all the exemplars from the other categories were defined as category-selective. Note that the 4 face exemplars used for this analysis were excluded from the classification procedure, thus equalizing the number of exemplars in the face category to the tool and house categories.
(C) Early visual electrodes: An electrode was defined as an early visual one when meeting ALL the following 3 criteria: 1) Anatomical location in retinotopic regions, as defined by an fMRI retinotopy experiment (in 7 of the 14 patients); 2) Significantly stronger response to geometric patterns compared with each one of the other categories (Levy et al. 2001); 3) Response latency of ≤100 ms (Fisch et al. 2009).

Data Analysis

Preprocessing and High-Gamma Responses Calculation
In the preprocessing stage, the signals were downsampled to 500 Hz and potential 60-Hz electrical interference was removed from the raw signals using a linear-phase notch finite impulse response filter of order 4. In addition, each electrode was dereferenced by subtraction of the averaged signal of all the electrodes, thus discarding non-neuronal contributions (Privman et al. 2007). The induced high-Gamma responses were computed in the following manner: Time–frequency spectrogram decomposition was based on Fourier transform amplitude spectrum with Hanning window tapering, calculated per trial in a 64-ms sliding window with a step size of 5.96 ms and averaged across trials (Lachaux et al. 2005). The time–frequency decomposition of each trial was divided by a prestimulus baseline of 250 ms, averaged across all the trials of a given category. In order to avoid dependencies between trials (Pereira et al. 2009), in the classification analysis the normalization was based on the trial-specific baseline, averaged across all the electrodes. The induced high-Gamma responses were calculated as the average of the spectrogram's frequency rows in the range of 80–150 Hz (Canolty et al. 2006; Ray et al. 2008; Vidal et al. 2010). In contrast, evoked responses were computed by averaging the raw ECoG signal according to the onset of the visual stimuli.

Grand-Averaged Exemplar Tuning
The calculation of the mean exemplar tuning curve (Fig. 3, left panels) was done as follows (see Supplementary Fig. 3A): for each electrode, the exemplars were first rank-ordered according to the high-Gamma AUC in a time window of 50–350 ms poststimulus, from the exemplar that elicited the strongest response to the exemplar that elicited the weakest response. This short time window was selected so as to be largely below the average time of saccade initiation and in accordance with previous ECoG studies (Liu et al. 2009). The ranking was done based on the responses in odd trials ($n = 3$) only. Then, according to this preset order, the high-Gamma AUC for each exemplar was computed based on the responses in even trials ($n = 3$). The procedure was then reversed, this time using the even trials to order the exemplars and the odd trials to compute the responses. Given the high resemblance of the 2 independent tuning curves (Supplementary Fig. 3B), they were averaged together. The resulting tuning curve was normalized according to the top-ranked exemplar response and then averaged across electrodes, yielding the grand-averaged tuning curve presented in Figure 3.

We also defined an exemplar selectivity index for each electrode as follows:

$$\text{Exemplar selectivity idx} = \frac{\text{response to 2 top exemplars} - \text{response to 2 bottom exemplars}}{\text{response to 2 top exemplars} + \text{response to 2 bottom exemplars}}$$

As a control, we applied this analysis also to the evoked responses (Supplementary Fig. 4B). Following Liu et al. (2009), the evoked response was defined as the range of the raw intracranial field potential signal (MAX–MIN) in the same time window used in the main analysis (50–350 ms poststimulus).

Table 1
Patient characteristics

Patient ID	Gender	Age (years)	Handedness	Implanted hemisphere	No. of implanted electrodes	Seizure localization
S1	F	36	Right	Bilateral	128	Left F-T
S2	M	22	Right	R	128	Right F-T
S3	M	21	Right	R	124	Right T
S4	F	23	Right	L	128	Left O
S5	F	23	Right	R	128	Right T-P
S6	M	31	Right	R	128	Right T
S7	F	55	Right	Bilateral	106	Left T
S8	M	22	Right	R	127	Right T
S9	M	25	Left	L	127	Left T
S10	F	21	Right	R	104	Right O
S11	F	45	Right	L	94	Left T
S12	F	32	Right	L	84	N/A
S13	F	25	Right	L	127	Left mT
S14	F	26	Right	Bilateral	187	Left aT
					N = 1720	

F, female; F-T, fronto-temporal; M, male; O, occipital; T, temporal; T-P, temporal-parietal; P, parietal; m, mesial; a, anterior.

As a second control, we repeated the exemplar tuning analysis without first ranking the exemplars for each individual electrode. Instead, we computed the response to each exemplar, averaged across the entire electrode ensemble, based on even trials ($n = 3$). The exemplars were then rank-ordered according to this "global order", and the tuning curves were computed based on the responses in odd trials ($n = 3$). The responses were normalized according to the top-ranked exemplar response. The procedure was then reversed, this time using the odd trials to generate the global order and the even trials to plot the responses. Next, the 2 independent tuning curves were averaged, and the resulting tuning curves were averaged across electrodes, yielding the grand-averaged "global tuning curve" presented in Figure 3 (right panels).

Data processing was carried out using MATLAB (The Mathworks, Inc.) and EEGLAB (Delorme and Makeig 2004).

Classification of Exemplars in Single Trials

A linear discriminant analysis (Duda et al. 2001) was used to compute the hyperplane, which optimally separated the responses to the various exemplars. The classifier was cross-validated using a leave-one-out scheme. In each fold of the cross-validation procedure, the classifier was trained based on 5 of the 6 repetitions of each exemplar and the classification performance was assessed using the remaining repetition.

The classifier was trained to distinguish between the 35 exemplars of the image set (10 exemplars from each category, except of the patterns category, which contained only 5 exemplars). Note that 4 face exemplars were excluded from this analysis because they were used to define the face selectivity of the electrodes. In order to assess whether the classification performance was dependent on the specific face exemplars that were excluded, 4 face exemplars were randomly selected and excluded from the exemplar classification, repeating this process 100 times. The mean decoding rate of faces in ventral face electrodes, obtained in this process, was very similar that that reported in the main text (29.1% and 33.3%, respectively).

Following Liu et al. (2009), this classification approach was implemented both at the level of single electrodes (Supplementary Fig. 6) and at the level of an ensemble of electrodes from different patients by concatenating the responses of individual electrodes (Fig. 5A) (Liu et al. 2009).

Classifier Input

The classifier input was either the AUC of the high-Gamma response at a time window of 50- to 350-ms poststimulus onset or the time course of the high-Gamma response within this time range. To avoid overfitting the high-Gamma responses were downsampled by a factor of 4, yielding 13 time points between 50 and 350-ms poststimulus onset (Pereira et al. 2009).

The classification performance of high-Gamma responses was compared with that obtained with low-Gamma (25–40 Hz) responses and with evoked responses. For the classification of evoked responses, the AUC measure was replaced by the range (MAX–MIN) of the raw signal within the same time range used for the Gamma responses (Liu et al. 2009). The duration of the time window was varied systematically between 300 and 500 ms (see Supplementary Fig. 5).

In addition, a novel data-driven approach was employed in which the classifier was trained based on the time–frequency decomposition of each trial at a time window of 50–350 ms poststimulus. Given the high-dimensionality of the time–frequency decompositions, within each training fold the following dimensionality reduction procedure was used: the classifier was first trained to distinguish between the 4 categories (faces, tools, buildings, and patterns) based on each time–frequency bin. To that aim, 4 exemplars were randomly selected from each category and the data was divided into nonoverlapping training set (60% of the trials for each category) and testing set (the remaining 40% of the trials). The P-value of the decoding performance of each time–frequency bin was estimated using the binomial distribution (Pereira et al. 2009). Note that this feature selection procedure was univariate: the category classifier was trained and tested on each time–frequency bin individually. Time–frequency bins that showed

significant ($P < 0.01$; uncorrected) category classification were then used as the input for the exemplar classifier. The distribution of these significant time–frequency bins is presented in Supplementary Figure 7. Note that this distribution followed the response profile of each electrode (compare Fig 1B and Supplementary Fig 7).

Note that the dimensionality reduction approach described above did not bias the results of the exemplar classification. First, the time–frequency bins were selected within each training set, independently from the test set (Pereira et al. 2009). Second, the 4 face exemplars that were used in the category classification were excluded from the exemplar classification. Owing to the lower number of exemplars in nonface categories, this exclusion was performed in the face category only. Thus, this feature selection procedure was applied for face electrodes only.

Reproducibility Across Repetitions

Following Liu et al. (2009), the degree of reproducibility of evoked responses and high-Gamma responses across repetitions of the same image was assessed. For each electrode and each exemplar, the raw signal or high-Gamma response between 0 and 800 ms poststimulus was extracted. Then, the correlation coefficient between every pair of repetitions for each exemplar was computed. Statistical significance was assessed by shuffling the exemplar labels 500 times. For each electrode, the z-score of the average correlation value relative to this null distribution was computed.

Neuronal Distance Analyses

In the neuronal distance analyses (Figs 5 and 6), a population vector was constructed for each face exemplar by concatenating the mean high-Gamma responses at a time window of 50- to 350-ms poststimulus onset (see Fig. 5A). The neural distance between 2 face exemplars was defined as the Euclidean distance between the corresponding population vectors. The neuronal distances between each pair of face exemplars were then correlated with several measures: 1) perceptual distance between the faces; 2) square root distance of low-level image features; and 3) square root distance of facial features (see below). The correlation between the neuronal distances in the "ventral face ensemble" responses and in the "early visual ensemble" was examined as well.

Statistical Tests

In all cases, unless mentioned otherwise, statistical significance was determined by means of a nonparametric Wilcoxon rank-sum test. In order to assess the significance of visual responses (see above) the comparisons were made per individual time points. In other cases (testing for category selectivity and for exemplar selectivity), they were carried out on the high-Gamma AUC at a window of 50–350 ms poststimulus. All these tests were 1 tailed.

The significance level of the grand-averaged exemplar tuning was assessed through a permutation test as follows (see Supplementary Fig. 4A): exemplar labels were randomly shuffled 10 000 times. For each iteration, the mean exemplar selectivity index was computed through the same procedure that was used for the real data, generating the distribution of this index under the null hypothesis.

The significance of exemplar classification was assessed through a permutation procedure, in which the exemplar labels were randomly shuffled 10 000 times.

The significance levels of correlation coefficients were assessed with a student's t-test. Bonferroni procedure was applied to correct for multiple comparisons.

The significance level of the difference in overall decoding rate between high-, low-Gamma, and evoked responses was first assessed with a 1-way repeated-measures ANOVA followed by post hoc t-tests (see Supplementary Fig. 5).

Perceptual Similarity Measurement

Perceptual similarity was measured in 2 different psychophysical experiments performed on groups of healthy observers. In the first experiment 20 healthy observers (9 males, 17 right handed, mean age 32.5 ± 2.1 years) were asked to subjectively cluster the face images
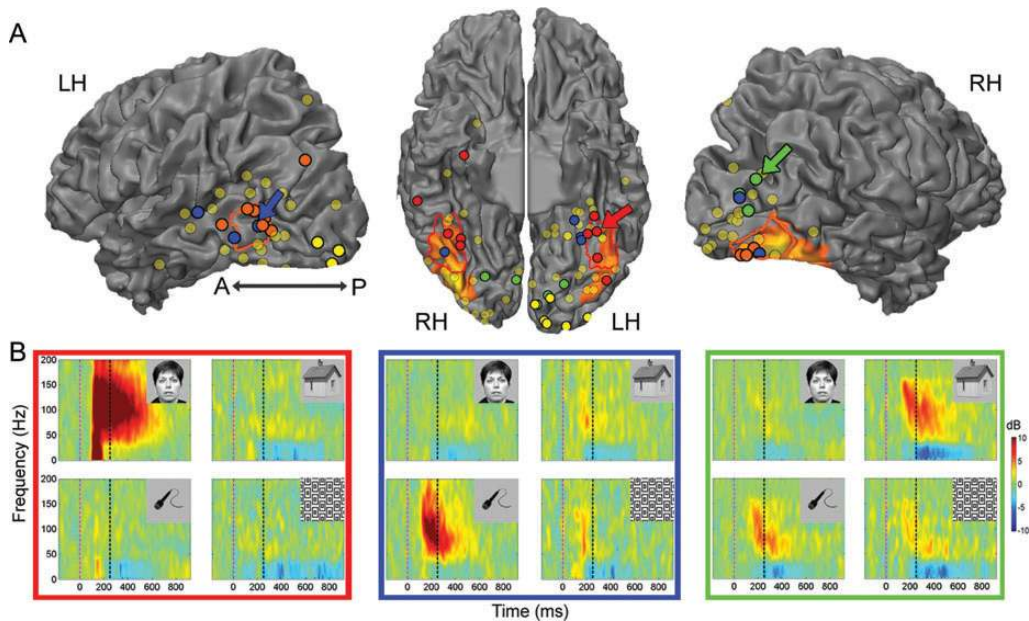
Figure 1. Category selectivity in high-order visual cortex. (*A*) Location of all visually responsive ECoG recording sites (131 electrodes in 14 patients) superimposed on a cortical reconstruction of a healthy subject from a previous fMRI study. Electrode positions are shown on a ventral view of the cortical hemispheres (middle) and on a lateral view of the cortical hemispheres (sides). Visual but noncategory-selective electrodes are marked in shaded yellow and category-selective electrodes are marked in red (ventral face-selective electrodes), orange (lateral face-selective electrodes), green (house-selective electrodes), and blue (tool-selective electrodes). Early visual electrodes are marked in bright yellow. The colored arrows point to the electrodes shown in (*B*). The electrodes are superimposed on the peak face-related activation measured in an fMRI mapping experiment. The borders of the fusiform face area (FFA, red) and of the occipital face area (OFA, orange) based on a previous fMRI mapping experiment in control subjects are included for reference as well. (*B*) Time–frequency spectrograms showing induced spectral responses to the 4 image categories of a ventral face-selective electrode (red frame), a tool-selective electrode (blue frame), and a house-selective electrode (green frame). Spectrogram's colors indicate power increase (red) or decrease (blue) relative to prestimulus baseline. Dashed vertical lines denote stimulus onset and offset.

used in the main experiment according to their similarity. The face images were printed on 14×14 cm cardboards. The observers arranged these cardboards on a table—such that similar looking faces were placed close to each other and distinct looking faces were placed further apart from each other. Pairwise distances between the centers of each face were measured and normalized according to the observer's mean distance between faces. Next, for each pair of faces, we computed the mean perceptual distance rating as well as a coefficient of variation across individuals (mean divided by standard deviation). This coefficient reflects the consistency of the ratings across observers. Consistent face pairs were defined as the 50% face pairs with the lowest coefficients of intersubject variance, that is, highest consistency across observers. Note that these cross-subject consistent perceptual ratings should in principle be more reliably indicative of the patients' own rating.

The experiment described above provides a more "ecological" (i.e., close to natural vision) measure of perceptual similarity. However, the viewing conditions were completely different from those used in the ECoG experiment. For example, in the ECoG experiment, the images were presented on a computer screen for 250 ms, whereas in the psychophysical experiment, the images were presented on cardboards and subjects were allowed to inspect them for unlimited duration. To better simulate the viewing conditions during the ECoG experiment, we performed a second psychophysical experiment in which 16 healthy observers (8 males, all right handed, mean age 28.1 ± 0.6 years) performed a same/different task on pairs of faces from the main experiment. In each trial, 2 face images were presented sequentially on a computer screen, each one for 250 ms with an interstimulus interval (ISI) of 500 ms. Each pair of different faces (91 pairs in total) was repeated 12 times throughout the experiment in a random order. There were 1092 "different" trials in total as well as 364 "same" trials (that were excluded from the analysis), in which the same face was presented twice. For each different face pair and for each participant, the inverse of the median of correct reaction times was computed (Kahn et al. 2010). Next, this measure was averaged across

participants. Similar to the first psychophysical experiment—we also computed for each face pair a coefficient of variation across observers.

Face-Related Features

Inspired by previous studies in the macaque (Freiwald et al. 2009) for each face image, the following parameters were measured: 1) face eccentricity (aspect ratio) defined the eccentricity of a solid ellipse constituting the face outline; 2) total face area— the area of an ellipse constituting the face outline; 3) hair intensity— mean intensity of hair pixels; 4) hair width—width of the hair at the level of the eyes, normalized by total face width; 5) hair height—height of the hair above the forehead, normalized by total face height; 6) hair pixels—total number of hair pixels; 7) eye–eyebrow distance—vertical distance between iris center and eyebrow, normalized by face height; 8) eyebrow thickness at the center of the eyebrow, normalized by face height; 9) eyebrow intensity—mean intensity of eyebrow pixels; 10) intereye distance—distance between iris centers, normalized by face width; 11) Iris diameter—diameter of a solid circle around the iris, normalized by face width; 12) eye width—major radius of a solid ellipse around the eye, normalized by face width; 13) eye height— minor radius of a solid ellipse around the eye, normalized by face height; 14) nose base width, normalized by face width; 15) nose altitude, normalized by face height; 16) nose–mouth distance—distance of the mouth below the nose, normalized by face height; 17) lower lip thickness, normalized by face height; 18) upper lip thickness, normalized by face height; 19) mouth width—horizontal distance between the 2 corners of the mouth, normalized by face width; 20) mouth–chin distance—distance of the chin below the mouth, normalized by face height. We limited the analysis to these 20 face-related features both due to statistical power considerations and due to the fact that these measures were performed manually for each face image.
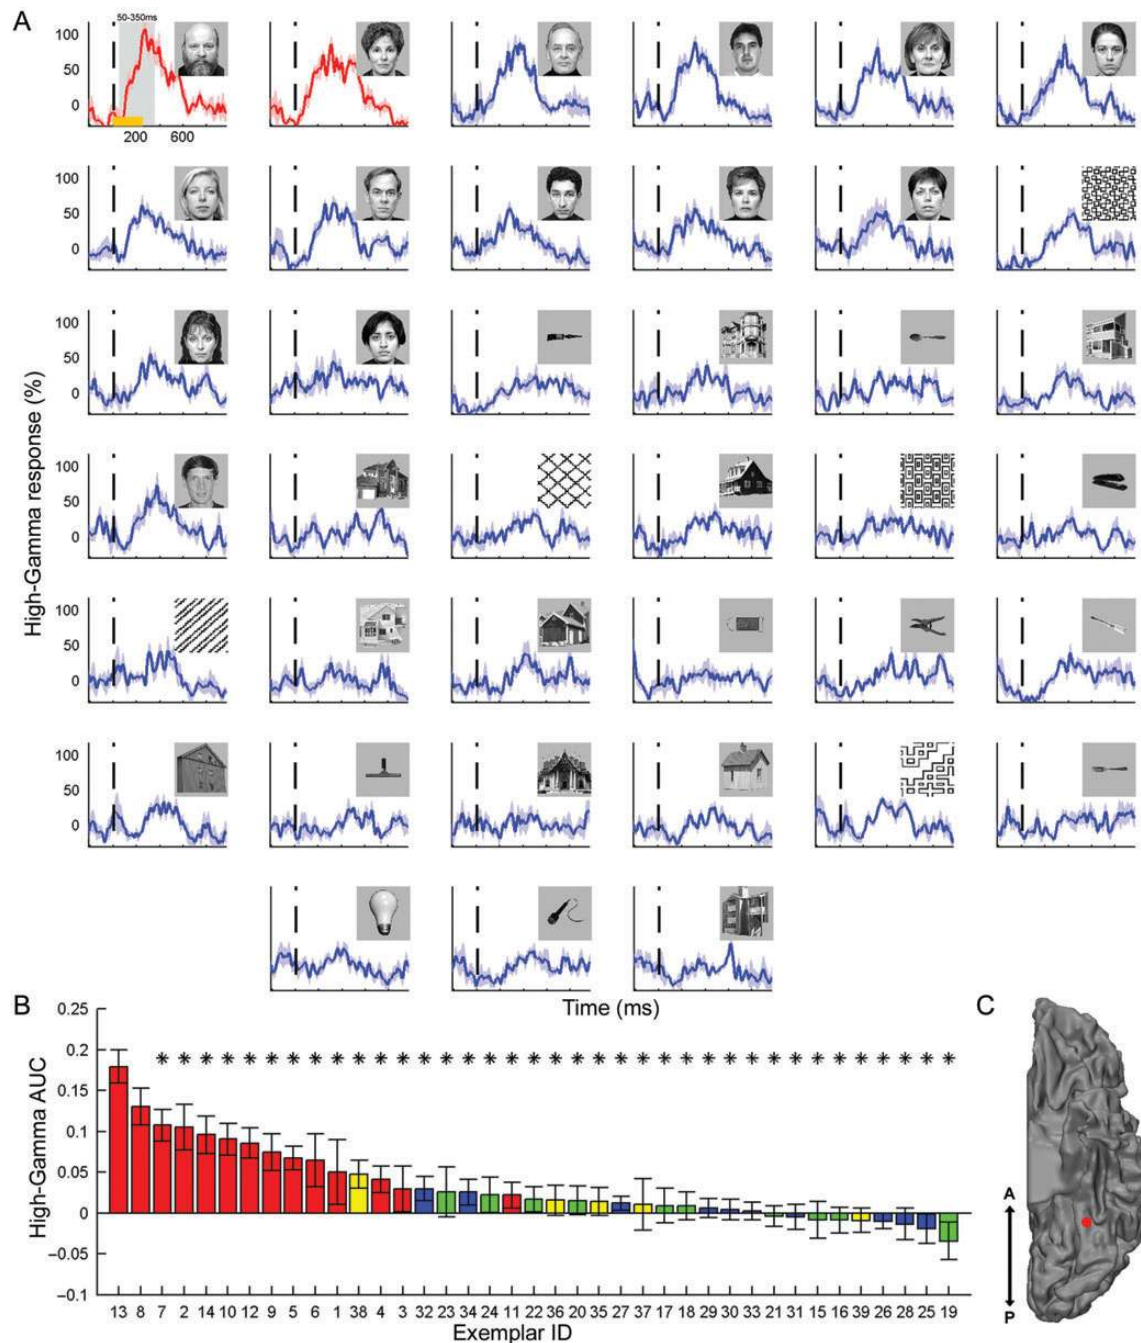
Figure 2. Exemplar selectivity in a face-selective electrode. (A) Induced high-Gamma (80–150 Hz) responses to each of the exemplars in the image set. The responses are arranged according to the AUC during a time window of 50–350 ms poststimulus from the strongest response (top left) to the weakest response (bottom right). The vertical dashed line marks stimulus onset and the shaded areas denote SEM across the 6 repetitions of each exemplar. Responses which are significantly ($P < 0.05$; uncorrected) lower than the top response are shown in blue, whereas nonsignificantly different responses are plotted in red. Orange horizontal bar (top left response) indicates image presentation duration and the shaded gray area denotes the time window in which the AUC was computed. (B) Exemplar tuning curve: AUC of each exemplar, sorted by response strength. The color of the bars denotes the category: faces (red), tools (blue), houses (green), and patterns (yellow). Asterisks mark significantly ($P < 0.05$; uncorrected) lower responses than the top exemplar. Error bars denote SEM. (C) Location of the electrode on a ventral view of the left hemisphere.

Low-Level Features

To quantify the extent to which the responses in face-selective electrodes could be explained by low-level characteristics of the images, a list of 12 basic image properties was considered (Liu et al. 2009): 1) mean pixel grayscale value; 2) standard deviation of the pixel grayscale values; 3) minimum pixel grayscale value; 4) number of pixels different from the background gray; 5) number of pixels below the background gray; 6) number of pixels above the background gray; 7) number of very dark pixels (grayscale value <64); 8) number of very bright pixels (grayscale >192); 9) number of boxes above background gray (box size = 20 pixels); 10) number of boxes below background gray; 11) number of very dark boxes (mean intensity within the 20 × 20 pixel box <64); and 12) number of very bright boxes (mean intensity within the 20 × 20 pixel box >192).

Electrode Localization

Computed tomography (CT) scans following electrode implantation were co-registered to the postoperative MRI. Pre- and postoperative MRIs were both skull-stripped using the BET algorithm from the
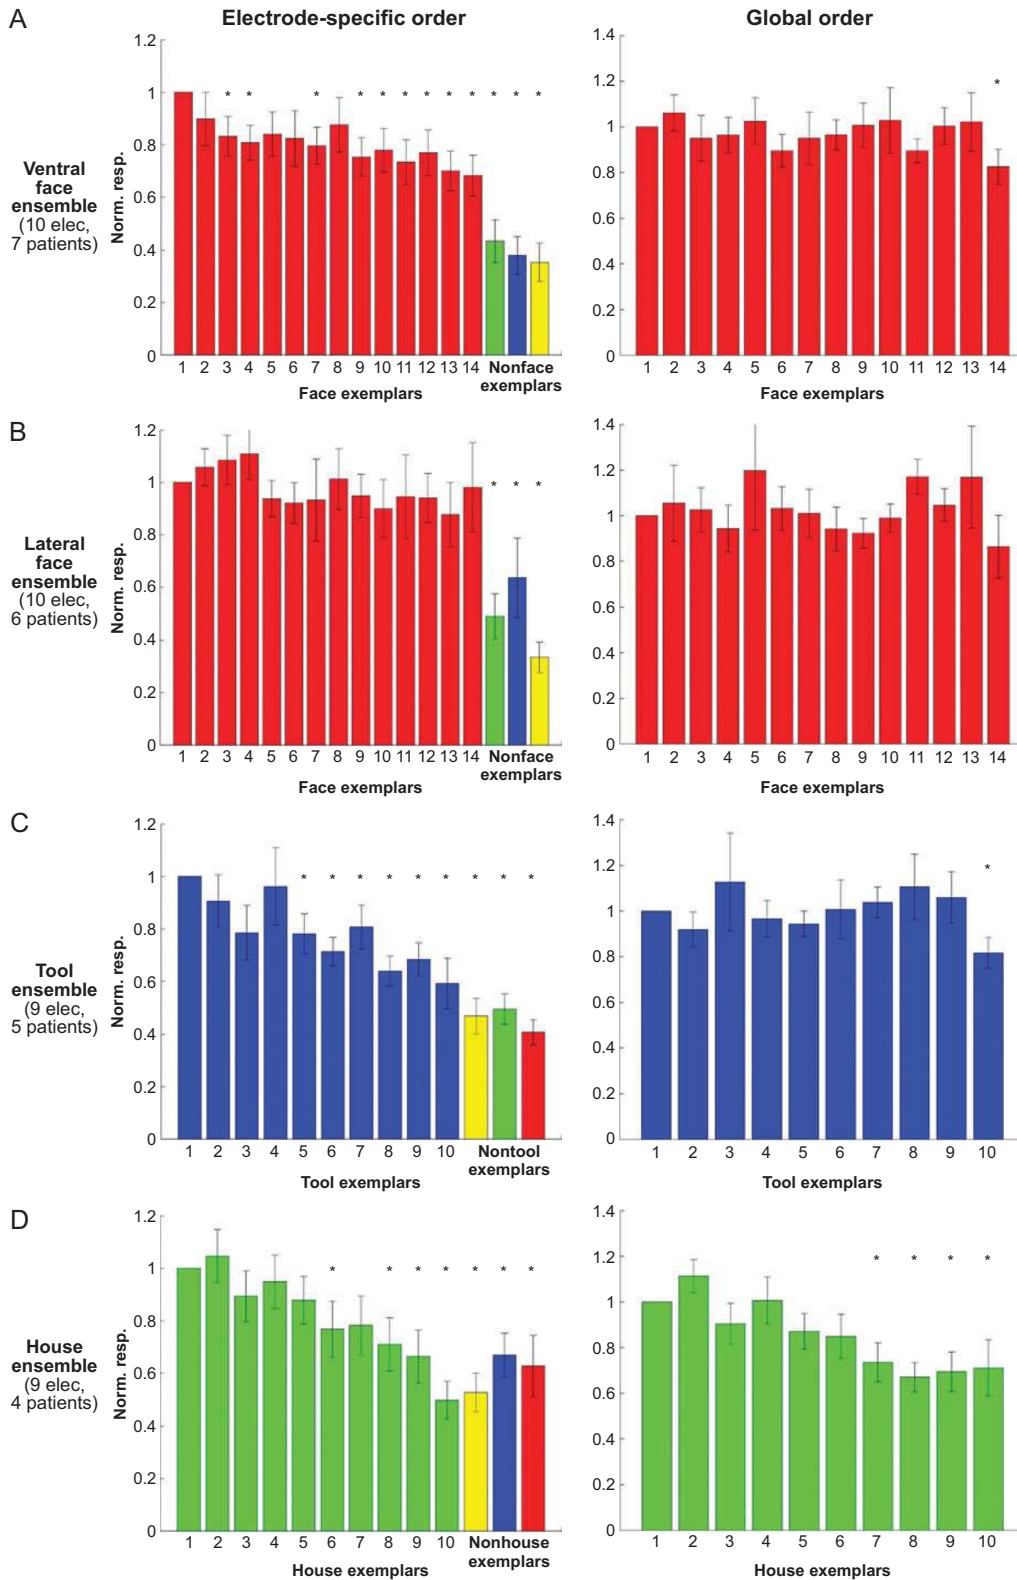
Figure 3. Grand-averaged exemplar tuning in 10 ventral face-selective electrodes (*A*), 10 lateral face-selective electrodes (*B*), 9 tool-selective electrodes (*C*), and 9 house-selective electrodes (*D*). Left panel: for each electrode, the exemplars were rank-ordered according to high-Gamma response AUC in even trials ($n = 3$), and the tuning curve was computed based on the responses in odd trials ($n = 3$). The procedure was then reversed, this time using the odd trials to order the exemplars and the even trials to plot the responses, and the 2 independent tuning curves were averaged (see also Supplementary Fig. 3). One exemplar from each nonpreferred category that yielded the strongest response is also shown for comparison. The responses were normalized according to the response to the "optimal" exemplar. Same asterisks and error bar notations as in Figure 2*B*. Right panel: "global tuning": the exemplars were rank-ordered according to their mean response across the entire electrode population in half of the trials ($n = 3$), and the tuning curve was computed based on the responses in the other half of trials ($n = 3$). Note the lack of global tuning, indicating that unlike the electrode-specific order—the global order was not consistent across the 2 datasets.
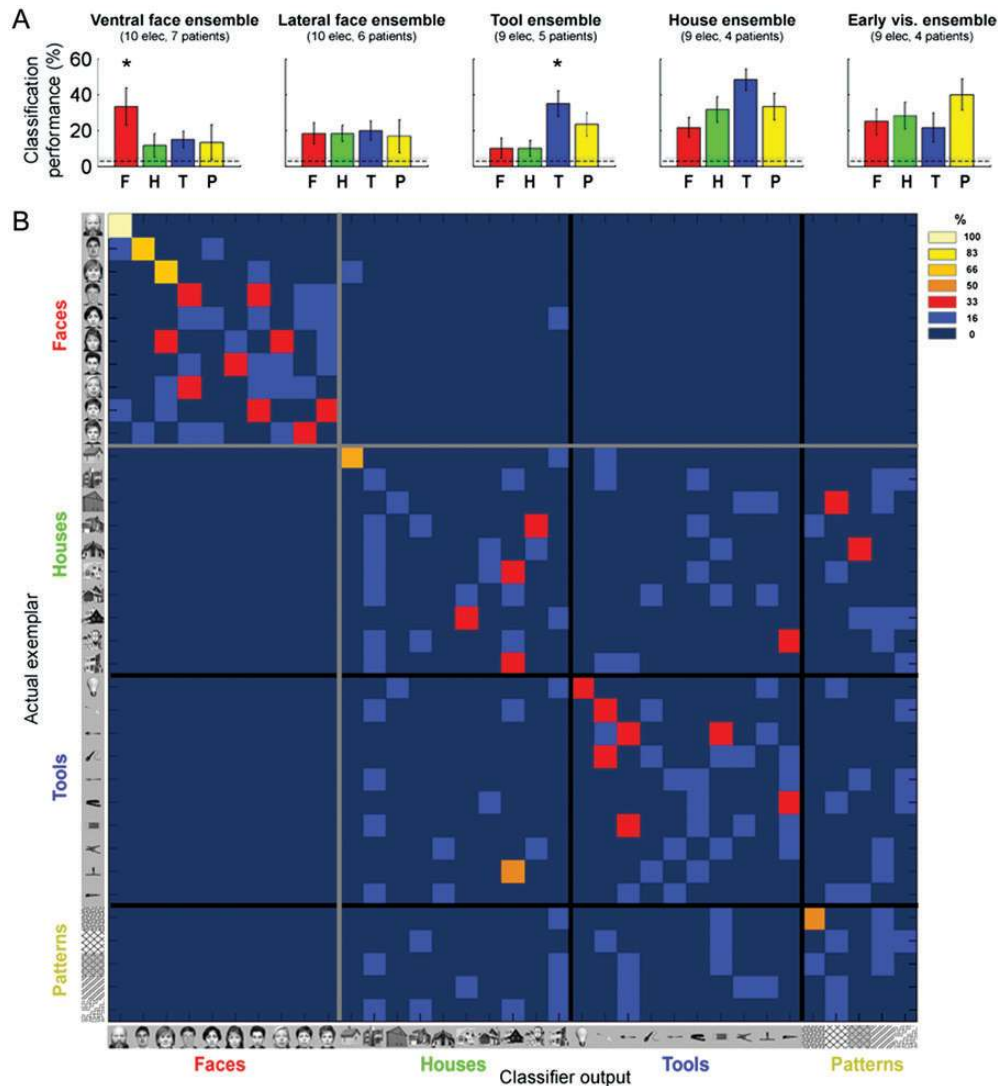
Figure 4. Exemplar decoding in category-selective electrodes. (*A*) From left to right: exemplar classification performance of the ventral face-, lateral face-, tool-, house-, and early visual ensembles. Chance level is shown by the dashed line (1 of 35 exemplars = 2.8%) and significance level is marked by the dotted line ($P < 0.05$). Error bars denote SEM across exemplars in each category. A significant difference between the decoding rates of exemplars from the preferred and nonpreferred categories is marked by an asterisk. (*B*) Misclassification pattern (confusion matrix) of the "ventral face ensemble", constructed by concatenating the high-Gamma responses of the 10 ventral face-selective electrodes. Each row denotes the actual exemplar that was presented, and each column denotes the classifier prediction (color coded from 0% in dark blue to 100% in bright yellow). The black and gray lines mark the borders between the categories.

Oxford Centre for Functional MRI of the Brain (FMRIB) software library (FSL; www.fmrib.ox.ac.uk/fsl/) followed by co-registration to account for possible brain shift caused by electrode implantation and surgery. Electrodes were identified in the CT using BioImagesuite (www.bioimagesuite.org). The coordinates of the electrodes were then normalized to Talairach coordinates (Talairach and Tournoux 1988) and rendered in BrainVoyager software in 2 dimensions as a surface mesh, enabling precise localization of the electrodes both with relation to the patient's anatomical MRI scan and in standard coordinate space. For joint presentation of all patients' electrodes, electrode locations were projected onto a cortical reconstruction of a healthy subject from a previous fMRI study of our group (Hasson et al. 2003).

fMRI Mappings in Patients

Prior to electrode implantation, 5 of the 14 patients underwent a standard fMRI category-localizer experiment used to delineate the borders of category-selective regions within high-order visual cortex. In this experiment, blocks of faces, buildings, man-made objects, and geometric patterns were presented in epochs of 10 s, followed by a 6-s blank screen. Each image was presented for 250 ms, followed by a 750-ms blank screen, while the patients performed a 1-back memory task (Levy et al. 2004). Seven patients performed a second fMRI localizer experiment, which mapped the borders of retinotopic areas according to the representations of the vertical and horizontal visual field meridians (Engel et al. 1994; Sereno et al. 1995; DeYoe et al. 1996). Patients were scanned on a General Electric Signa HDx 3-T scanner. Functional imaging using Blood-oxygenation-level–dependent contrast were obtained with gradient-echo echo-planar imaging sequence [field of view (FOV) = 220 mm, 3.5 × 3.5 × 4 mm voxel size, 64 × 64 matrix, flip angle = 70, repetition time (TR) = 2000 ms, echo time (TE) = 30 ms, axial acquisition plane, 210 contiguous volumes in the category-localizer experiment and 256 in the retinotopy experiment]. An anatomical T_1-weighted image was acquired using a spoiled gradient recalled sequence [FOV = 256 mm, 1 × 1 × 1 mm voxel size, 256 × 256 matrix, flip angle = 8, TR = 2500 ms, TE = 30 ms, inversion time = 650 ms, axial acquisition plane, 180 slices]. Face-related activation was identified in each patient separately using the contrast of faces versus baseline.
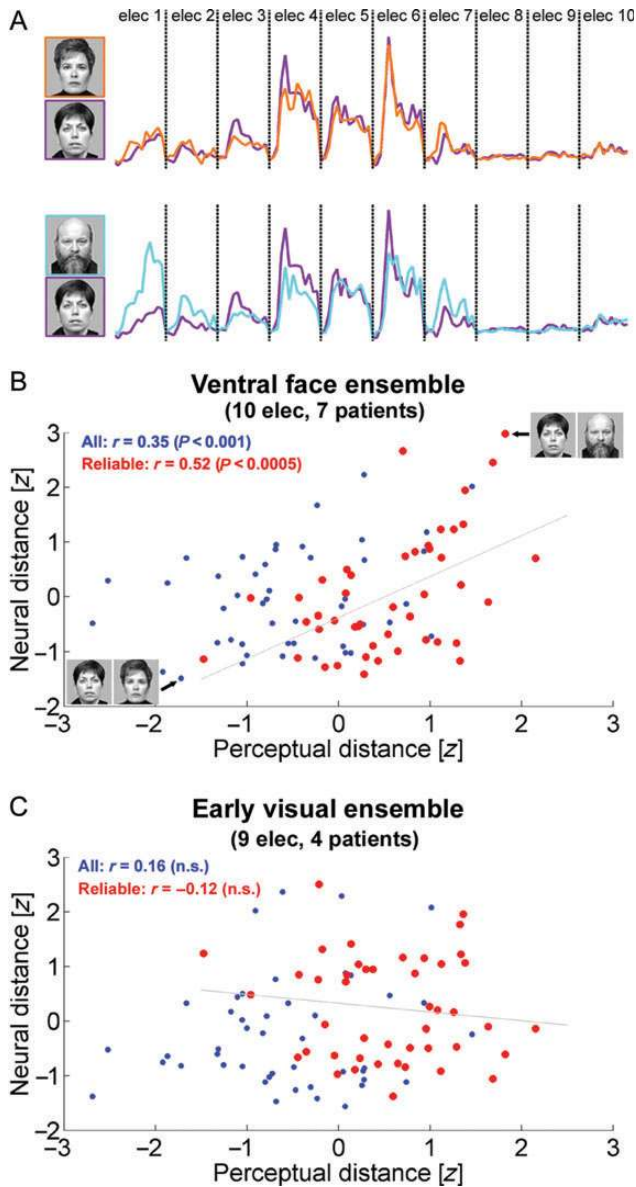
Figure 5. Exemplar selectivity in ventral face-selective electrodes reflects perceptual similarity between face exemplars. (*A*) Calculation of neural distance between face pairs. For each face exemplar, a population vector was constructed by concatenating the high-Gamma responses of ventral face-selective electrodes in a time window of 50- to 350-ms poststimulus onset. The neural distance between a pair of face exemplars was defined as the Euclidean distance between the corresponding population vectors. This panel depicts the face pair with the minimal neural distance (top) and the face pair with the maximal neural distance (bottom). These 2 face pairs are also marked in (*B*). (*B* and *C*) For each pair of faces, 2 measures were correlated: neuronal distance (*y*-axis) of the "ventral face ensemble" (*A*) or the "early visual ensemble" (*B*) and psychophysical distance (*x*-axis), based on ratings of 20 healthy observers. Each dot denotes a pair of faces. The 50% face pairs that received the most reliable ratings across observers (i.e., lowest coefficient of variation) are depicted in red.

fMRI Mappings in Control Subjects

The borders of the FFA and of the occipital face area (OFA) (see Fig. 1) were depicted based on a previous fMRI study conducted in our group (Hasson et al. 2003). Briefly, 12 healthy subjects participated in a blocked-design experiment that included line drawings of faces, buildings, man-made objects, and geometric patterns. Each image was presented for 800 ms followed by a 200-ms blank screen containing a fixation point. Each block lasted 9 s and was followed by a 6 s blank with fixation point screen. The borders of

face-related areas were set based on the contrast of faces > buildings (random-effects analysis, FDR corrected) (Kanwisher et al. 1997).

fMRI Data Analysis

fMRI data were analyzed with the BrainVoyager software package (R. Goebel, Brain Innovation, Maastricht, The Netherlands). The first 3 images of each functional scan were discarded. The functional images were superimposed on 2D anatomical images and incorporated into the 3D datasets through trilinear interpolation. The complete dataset was transformed into Talairach space (Talairach and Tournoux 1988). Preprocessing of functional scans included 3D motion correction, slice scan time correction, linear trend removal, filtering out of low frequencies up to 4 cycles per experiment, and 8-mm spatial smoothing.

Results

Exemplar Selectivity in Face-Related Areas

Fourteen patients (5 males, 13 right handed, mean age 29.1 ± 2.7 years, see Table 1), each implanted with 84–187 intracranial electrodes for presurgical evaluation, participated in the present study. The patients viewed images of faces, man-made tools, buildings, and patterns, presented for 250 ms in pseudorandom order at a rate of 1 Hz, while performing a 1-back memory task. Each category of images consisted of 5–14 different exemplars, which were repeated 6 times throughout the experiment (see Materials and Methods section for further details).

Figure 1*A* depicts the distribution of all visually responsive ECoG recording sites (131 electrodes in 14 patients) projected on a cortical reconstruction of a healthy subject from a previous fMRI study (Hasson et al. 2003). Peak face-related fMRI activation, measured in 5 patients prior to electrodes implantation (see Materials and Methods section), is also presented as well as the estimated borders of the FFA (red) and the OFA (orange) based on a previous fMRI mapping experiment in healthy participants (see Materials and Methods section). Our analysis focused on induced high-Gamma (80–150 Hz) changes in the spectral power of the ECoG signals (Canolty et al. 2006; Ray et al. 2008). First, we examined whether each electrode showed a significant visual response. This analysis revealed that 131 of the 1720 electrodes (7.6%) were visually responsive. Next, we examined to what extent each electrode showed a preferential response to exemplars from one category over exemplars from the other categories (see Materials and Methods section). Twenty electrodes in 9 of the patients showed a significant selectivity for faces. These electrodes could be divided anatomically into 2 clusters: a ventral cluster along the fusiform gyrus (10 electrodes in 7 patients; depicted in red) and a lateral occipital cluster (10 electrodes in 6 patients; orange) (Allison et al. 1999). Nine electrodes in 5 patients were tool selective (blue) and 9 electrodes in 4 patients were house selective (green). As a control, we also identified 9 early visual electrodes in 4 of the patients based on fMRI retinotopy and functional criteria (depicted in yellow with thick black contours; see Materials and Methods section). The distribution of electrodes in individual patients is presented in Supplementary Figure 1 and in Supplementary Table 1.

Examples of time–frequency decompositions, which depict the trial-averaged spectral power changes induced by the various object categories, are shown in Figure 1*B* for a ventral face-selective electrode (red frame), for a tool-selective
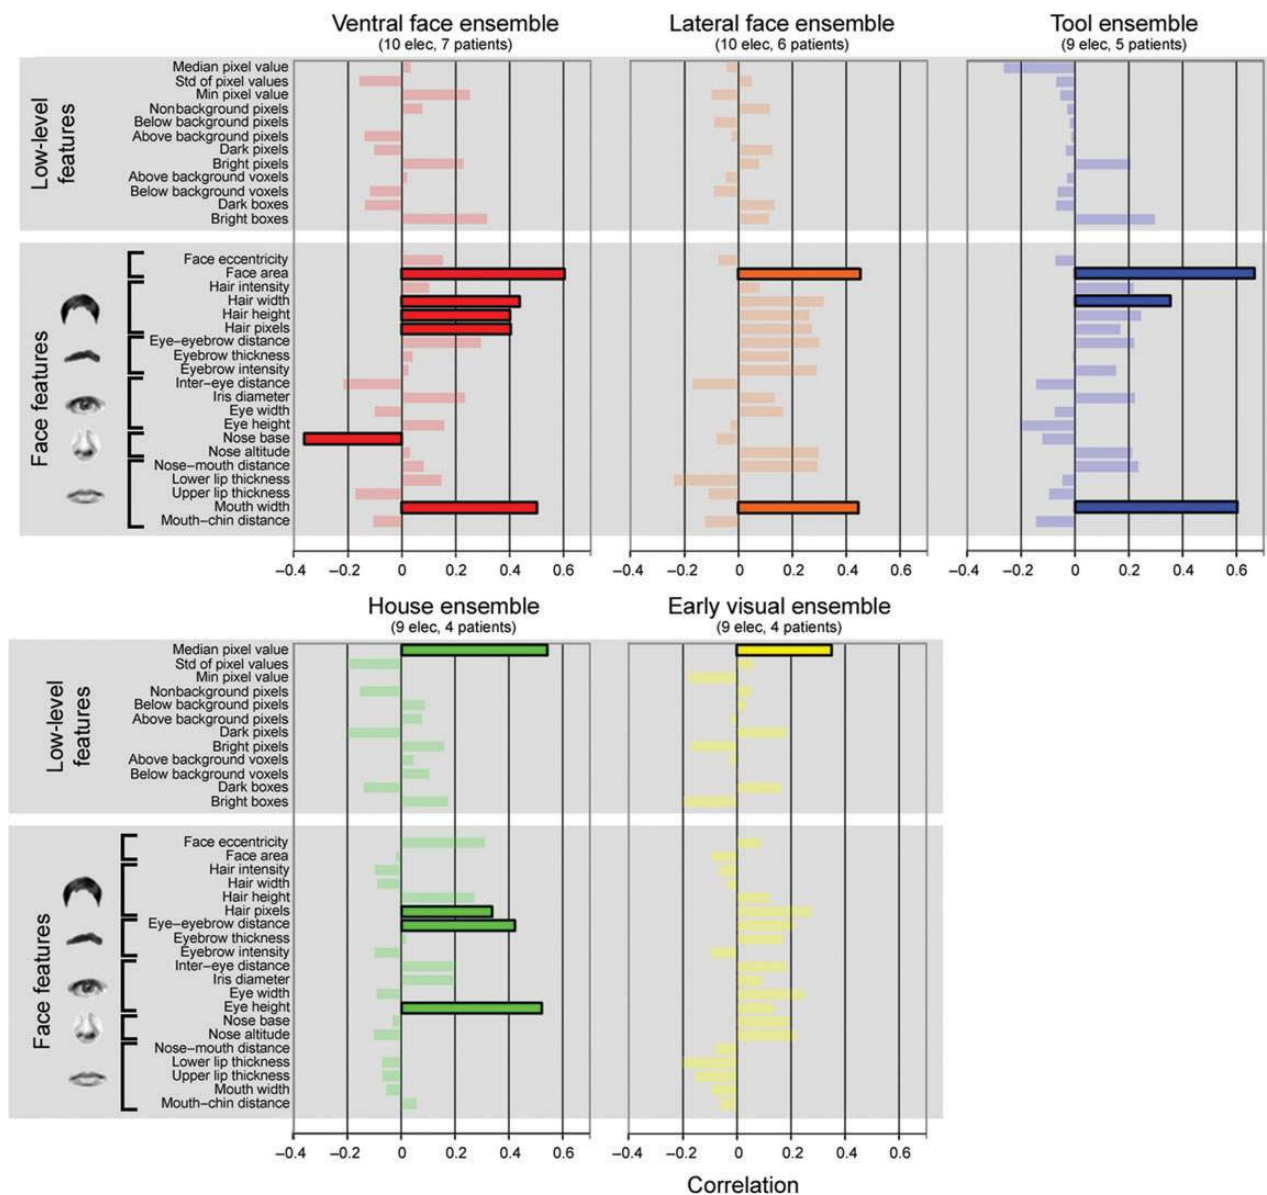
Figure 6. Organizing principles underlying exemplar selectivity. Correlations between the neural distance of each electrode ensemble and a set of low-level image features and face-related features. Significant correlations ($P < 0.05$, Bonferroni corrected) are highlighted.

electrode (blue), and for a house-selective electrode (green). The spectral decompositions show the typical broadband and fast rising response followed by a sustained elevation of Gamma-band power and a corresponding reduction in low frequencies (5–25 Hz) power (also termed event-related de-synchronization (Crone et al. 2006; Privman et al. 2011).

Next, we examined to what extent the category-selective face electrodes also showed exemplar selectivity— i.e., showed a consistent difference in the response to various face exemplars. Figure 2 depicts the induced high-Gamma responses of a face-selective electrode to each of the exemplars in the image set. The high-Gamma responses were arranged according to the AUC during a time window of 50–350 ms poststimulus. Note that although exemplar selectivity extended well beyond the 350 ms boundary, this short time window was selected in order to remain within the typical

time of saccade initiation and in accordance with previous ECoG studies (Liu et al. 2009). As expected in a category-selective electrode, all top responses were driven by images of faces. Critically, a robust and significant exemplar tuning was found when comparing the responses across different face images. Thus, almost all face exemplars (12 of 13) produced a significantly ($P < 0.05$; uncorrected) lower response compared with the strongest face response (see Supplementary Fig. 2 for exemplar selectivity in a tool-selective electrode).

To examine if exemplar selectivity was a general phenomenon in face-related areas, we calculated the grand-averaged exemplar tuning of the 10 ventral and 10 lateral face-selective electrodes (left panels of Fig. 3A,B, respectively). The exemplar tuning was determined based on independent trial sets. For each electrode, we ranked-order the exemplars based on

3 trials and computed the exemplar tuning curve based on the remaining 3 trials (see Materials and Methods section and Supplementary Fig. 3). As can be seen, significant exemplar selectivity could be discerned also in the group responses. This effect was particularly robust in ventral face electrodes, in which the response to the last ranked face exemplar was reduced to a level of $68.2 \pm 7.7\%$ of the response to the "optimal" face exemplar. In order to test the significance of this effect we computed an exemplar selectivity index for each electrode and averaged this index separately for ventral and lateral face-selective electrodes. Then, we recomputed this index while randomly shuffling the exemplar labels (see Supplementary Fig. 4A and Materials and Methods section for further details). This permutation procedure showed that exemplar selectivity was significant both in ventral face electrodes ($P < 0.01$) and in lateral face electrodes ($P < 0.05$). Exemplar tuning was evident also in tool-selective electrodes ($P < 0.05$; Fig. 3C) and in house-selective electrodes ($P < 10^{-4}$; Fig. 3D). Interestingly, no exemplar tuning was evident in the evoked responses (Supplementary Fig. 4B).

An important question concerns the source of the exemplar selectivity. Such selectivity may be global, that is, a subset of the exemplars produced consistently stronger responses across patients and electrodes. Alternatively, the exemplar selectivity profiles may have been unique to each electrode, that is, for each recording site a different set of face images were the optimal stimuli. One way to distinguish between these alternatives can be obtained by recomputing the grand-averaged exemplar tuning curves without first ranking the exemplars for each electrode separately. Instead, we ranked-order the exemplars according to the mean response to each exemplar in 3 trials, averaged across the entire electrode population, and computed the exemplar tuning curve based on the remaining 3 trials. If all face-selective electrodes exhibit similar tuning profiles (i.e., respond the most to a particular subset of face exemplars), the population exemplar tuning curve should remain intact. However, as can be seen in the right panels of Figure 3A,B, ventral and lateral face electrodes failed to reveal a significant "global" tuning. In other words, there was no global order of face exemplars, which was consistent across the 2 independent trial sets. By contrast, in house-selective electrodes there were 4 specific house exemplars that consistently evoked weaker responses compared with the "optimal" house exemplar ($P < 0.05$; Fig. 3D, right panel).

Decoding Face Exemplars in Single Trials

Up to this point, we analyzed the mean exemplar tuning across several repetitions. In ventral face-selective electrodes, both the evoked responses and the high-Gamma responses were highly reproducible across repetitions of the same image (evoked responses: z-score = 5.21 ± 1.37; high-Gamma responses: z-score = 4.29 ± 0.95; see Materials and Methods section for details). Therefore, we set to examine the tuning at a single-trial level using a multivariate pattern classifier (linear discriminant analysis) (Duda et al. 2001). The classifier was trained to differentiate between the 35 exemplars from all categories (4 face exemplars were used to define the face-selective electrodes and thus were excluded from the current analysis), and it was cross-validated using a leave-one-out scheme (see Materials and Methods section). The classifier

was trained based on either the high-Gamma AUC (used in previous analyses) or the high-Gamma time course in a time window of 50- to 350-ms poststimulus onset. As higher decoding rates were obtained using the high-Gamma time course (see Supplementary Fig. 5), we chose that signal as the input of the classifier in the remaining analyses. Statistical significance of classification performance was assessed through a permutation test (see Materials and Methods section).

Despite the minimal training set ($n = 5$), 9 of the 10 ventral face-selective electrodes showed a significant classification performance (Bootstrap $P < 0.05$, uncorrected; Supplementary Fig. 6). The overall decoding rate was $6.2 \pm 0.7\%$ and the decoding rate for faces was $10.3 \pm 1.5\%$ (mean ± SEM across electrodes, chance level = 1 of 35 exemplars = 2.8%). Note that in most electrodes the highest accuracy was obtained in the discrimination between face exemplars, although several electrodes showed above chance classification for exemplars from other categories as well.

Next, following Liu et al. (2009), we constructed an ensemble vector by concatenating the high-Gamma time courses of all the ventral face electrodes and applied the same classification approach (Liu et al. 2009). Using this "ventral face ensemble" (10 electrodes in 7 patients), the overall accuracy rate was $19.1 \pm 4.2\%$ (mean ± SEM across exemplars, chance level = 2.8%). The accuracy in discriminating between different face exemplars was $33.3 \pm 10.5\%$ (chance = 2.8%, Bootstrap $P < 10^{-4}$; Fig. 4A). Importantly, the classification performance clearly exceeded chance level even when considering a more conservative threshold, taking into account the a priori category selectivity of the electrodes (10%, i.e., 1 of 10 face exemplars). Classification performance was significantly higher for face exemplars compared with nonface exemplars ($P < 0.05$). By contrast, the ensemble of lateral face electrodes (10 electrodes in 6 patients) did not exhibit higher classification rate for faces compared with nonfaces ($P = 0.47$).

Inspection of the misclassification or confusion patterns of the "ventral face ensemble" (Fig. 4B) revealed 2 aspects. First, faces were almost always (in 97% of trials) decoded as faces. By contrast, there were many cross-category misclassifications for houses, tools, and patterns. Note that these electrodes were defined as face-selective based on 4 face exemplars which were excluded from the current analysis, and therefore, these estimates were not biased. Second, the decoding accuracy rates for different face exemplars were not constant—some face exemplars were highly differentiated and some were not.

In addition to high-Gamma responses, we applied the classifier for low-Gamma (25–40 Hz) and for evoked responses (Supplementary Fig. 5). The overall decoding rate was significant in all signal types ($19.1 \pm 4.2\%$, $9.1 \pm 2.3\%$, and $11.4 \pm 2.2\%$ for high-, low-Gamma, and evoked responses, respectively, Bootstrap, $P < 0.05$). One-way repeated-measures ANOVA revealed a significant effect of signal type ($F_{2,68} = 3.78$, $P < 0.05$). The overall decoding rate for high-Gamma responses was significantly higher than these obtained by low-Gamma and evoked responses (post hoc t-test, $P < 0.05$). More specifically, the decoding rate of faces was much higher for high-Gamma responses compared with evoked responses ($33.3 \pm 10.5\%$ vs. $13.3 \pm 3.3\%$ respectively, $P < 0.05$).

We also applied a novel data-driven approach in which the classifier was trained based on features extracted from the

time–frequency decomposition (Supplementary Fig. 7). Instead of selecting a priori a specific frequency band, in each fold of the cross-validation procedure we selected time–frequency bins that were category-selective and then used them in order to decode the specific exemplar that was presented (see Materials and Methods section) (Pereira et al. 2009). Applying this method to ventral face-selective electrodes yielded an overall decoding rate of $19.5 \pm 3.5\%$ and a face decoding rate of $28.3 \pm 7.1\%$, further validating the previous results.

Organizing Principles Underlying Exemplar Selectivity

What could be the principles that underlie exemplar tuning in face-related areas? More specifically why did certain face exemplars produce similar neural responses while others produced highly different patterns? A particularly important parameter of interest, suggested in the object perception literature (Olson 2001; Op de Beeck et al. 2001, 2008; Haushofer et al. 2008; Drucker and Aguirre 2009) is perceptual similarity. Ideally, one would like to directly compare the neural responses and the perceptual similarity measures obtained from the patients themselves. However, due to constraints imposed by the clinical setting, collection of behavioral data from the patients was not feasible.

To examine whether the neuronal responses in the patients may have represented more "canonical" perceptual similarities, we performed a separate psychophysical experiment in which 20 healthy observers (9 males, 17 right handed, mean age 32.5 ± 2.1 years) were asked to subjectively cluster the face exemplars according to their similarity (see Materials and Methods section for details). For each pair of face exemplars, we compared the neuronal distance and the perceptual distance evaluated by the observers. The neuronal distance was computed in the following manner: we first constructed a population vector for each face exemplar by concatenating the single electrodes mean high-Gamma responses in a time window of 50–350 ms poststimulus. The neuronal distance between 2 face exemplars was defined as the Euclidean distance between the corresponding population vectors (Fig. 5*A*). In the "ventral face ensemble", we obtained a highly significant correlation between neuronal distance and perceptual similarity ($r = 0.35$; $P < 0.001$; Fig. 5*B* and Supplementary Movie 1 demonstrating misclassification of perceptually similar faces). Restricting this analysis to the more "canonical" pairs, that is, the 50% face pairs which received the most consistent ratings across observers, yielded a stronger correlation coefficient ($r = 0.52$; $P < 0.0005$; depicted in red in Fig. 5*B*; see Materials and Methods section). Note that these cross-subject consistent perceptual ratings should in principle be more indicative of the patients' rating. The correlation between neuronal distance and perceptual distance was significant, albeit weaker, in the "lateral face ensemble" ($r = 0.29$, $P < 0.005$; reliable face pairs only: $r = 0.35$; $P < 0.05$). In contrast, no significant correlation between neural and perceptual distances was observed in an ensemble of early visual electrodes (9 electrodes in 4 patients: $r = 0.16$, $P = 0.13$; reliable face pairs only: $r = -0.13$, $P = 0.40$; Fig. 5*C*), nor in an ensemble of tool-selective electrodes (9 electrodes in 5 patients: $r = 0.11$, $P = 0.49$) and in an ensemble of house-selective electrodes (9 electrodes in 4 patients: $r = 0.05$, $P = 0.72$).

It could be argued that the method that was used to evaluate perceptual similarity between faces, despite being "ecological" (i.e., close to natural vision), differed substantially from stimulus presentation conditions during the main ECoG experiment. For example, in the main ECoG experiment images were presented for 250 ms, whereas in the behavioral experiment subjects were allowed to inspect the images for unlimited duration. Therefore, we performed a control experiment on an independent sample of 16 observers (8 males, all right handed, mean age 28.1 ± 0.6 years), in which face pairs were presented sequentially while the participants performed a same/different task. Each image was presented for 250 ms with an ISI of 500 ms. For each pair of faces, the inverse of the median reaction time was computed as a measure of perceptual similarity (see Materials and Methods section) (Kahn et al. 2010). The results using this more controlled measure replicated the significant correlation found between neural distance and perceptual distance in the "ventral face ensemble" ($r = 0.23$, $P < 0.05$; reliable face pairs only: $r = 0.32$, $P < 0.05$). Importantly, the "early visual ensemble" did not exhibit any significant correlation between neural and perceptual distances ($r = -0.07$, $P = 0.52$; reliable face pairs only: $r = -0.04$, $P = 0.82$).

In order to further examine the neural representation of face exemplars, we measured a set of image properties for each face exemplar, which included 12 low-level features (Liu et al. 2009) (such as contrast, number of bright and dark pixels) and 20 face features, inspired by previous studies in the macaque (Freiwald et al. 2009) (e.g., total face area and intereye distance, see Materials and Methods section). Next, we computed the correlation coefficient between the neuronal distance of each pair of face images and the difference between these 2 images in each of the above parameters. Bonferroni procedure was applied in order to correct for multiple comparisons. The results of this analysis are depicted in Figure 6. As can be seen, the response of the "ventral face ensemble" (red) was significantly correlated only to face-related features (primarily to the total face area and hair length), and not to any of the low-level features. By contrast, the response of the "early visual ensemble" (yellow) was significantly correlated only to low-level features (mean intensity of the image) but not to any of the face-related features. The "tool ensemble" (9 electrodes in 5 patients) and "house ensemble" (9 electrodes in 4 patients) exhibited an intermediate pattern that was partially related to low-level features (house ensemble) and partially to face features.

Role of Low-Level Features

An important question concerns the possible role of low-level features in the observed exemplar tuning. We examined this issue using 2 approaches. First, as described above, there was a clear dissociation between the "ventral face ensemble" and the "early visual ensemble" in the parameters that predicted the neuronal distance between face exemplars. Second, examining the correlation coefficient between the responses of the "ventral face ensemble" and the "early visual ensemble" revealed that it was close to zero ($r = -0.07$, $P = 0.48$, see Materials and Methods section). Thus, a low-level explanation for the exemplar tuning in face-related areas seems unlikely.

Discussion

Exemplar Selectivity in Face-Related Cortex

Our results show that human face-related regions not only show a clear category preference, as has been shown in fMRI (Kanwisher et al. 1997; McCarthy et al. 1997; Ishai et al. 1999; Hasson et al. 2003), scalp-EEG (Bentin et al. 1996) and ECoG (Allison et al. 1994, 1999; Liu et al. 2009), but also exhibit a significant exemplar selectivity within the face category. Remarkably, despite the exceedingly small training set ($n = 5$), this neuronal selectivity was sufficiently robust and reliable to allow a significant, single-trial classification for face exemplars, even at the level of individual electrodes (Supplementary Fig. 6). The classification performance and the fact that the tuning for different face exemplars was consistent across independent trial sets (Fig. 3) argue against the possibility that the tuning merely reflected random fluctuations in the neuronal responses.

It is important to emphasize that 75% of the face-selective electrodes showing exemplar selectivity were localized within the borders of the well-known fMRI-defined FFA and the more posterior OFA (Kanwisher et al. 1997; Ishai et al. 1999; Hasson et al. 2003), as assessed by fMRI mappings in some of the patients and in control subjects (Fig. 1A). Thus, our results demonstrate exemplar selectivity in regions that were previously considered category-selective.

These results are in line with previous single unit recordings in nonhuman primates (Baylis et al. 1985; Hasselmo et al. 1989; Rolls et al. 1992; Tsao et al. 2006). Of particular importance is the study of Tsao et al. (2006), showing that face exemplars can be decoded at very high rates from single unit activity in the middle face patch (ML), which is assumed to be the homolog of the human FFA (Tsao et al. 2008). It is important to emphasize that the "exemplar" discrimination that was shown in the present study does not necessarily imply "identity" discrimination. Assessing identity selectivity necessitates comparing responses across transformations that clearly change the visual appearance of a face (such as viewpoint) while maintaining its identity, which were beyond the scope of this study. Such view-invariant representation was found in more anterior temporal areas than those recorded in the present study (Freiwald and Tsao 2010).

Our finding of exemplar selectivity in face-related cortex is also compatible with fMR-adaptation studies in humans (Grill-Spector et al. 1999; Rotshtein et al. 2004; Gilaie-Dotan and Malach 2007; Gilaie-Dotan et al. 2010). These studies have demonstrated that the FFA shows release from adaptation when facial identity is changed. However, fMR-adaptation is an indirect and coarse measure of selectivity (Sawamura et al. 2006; Mur et al. 2010). fMRI multivariate pattern analysis studies that addressed this issue more directly have yielded mixed results. Recently, Nestor et al. (2011) have shown that facial identity can be significantly decoded based on FFA multivoxel activity, whereas 2 previous studies have reached the opposite conclusion (Kriegeskorte et al. 2007; Op de Beeck et al. 2010). Of particular interest is a recent fMRI study demonstrating graded responses within the preferred category of the FFA and parahippocampal place area (Mur et al. 2012). As discussed below, our results extend this study by showing that within-category response differences in face-selective areas are related to perceptual similarity

between faces and to face-related features rather than to low-level features (Figs 5 and 6). Our results also demonstrate that there was no global tuning shared by all face-selective electrodes (Fig. 3). An intriguing possibility is that a common tuning preference may still exist within individual patients rather than across the patient population. However, in our dataset the number of face electrodes per patient was far too small (1.43 ± 0.43 electrodes) to allow a proper analysis of this question and thus it should be addressed in future studies. Altogether these findings suggest that it is unlikely that exemplar tuning in face-related areas can be attributed to global (e.g., attentional) or low-level effects.

While we focused in our analysis mainly on face exemplar selectivity, it is important to note that this phenomenon could also be observed in tool-selective electrodes and house-selective electrodes (Fig. 3 and Supplementary Fig. 2). This is consistent with a previous report of a reliable decoding of object identity in the macaque inferior temporal cortex (Hung et al. 2005) and with fMRI multivariate pattern analysis in the human lateral occipital complex (Eger et al. 2008). Thus, our results support the notion that exemplar selectivity is a general property in high-order visual cortex.

Low-Level Contribution?

An important question concerns the role that low-level visual aspects such as contrast, size, or spatial frequency, may have played in the exemplar tuning we observed. This concern is particularly significant for the present study because we used a diverse set of face exemplars, which allowed a wide range of face-related features and perceptual distances. Indeed the finding that early visual cortex electrodes as well as tool- and house-selective electrodes significantly discriminated between different exemplars likely indicates a substantial difference in the low-level aspects of these images.

Could it be then that the exemplar selectivity in high-order face-selective electrodes simply reflected differences in the low-level image statistics? The fact that we found a clear dissociation in the parameters that explain the neuronal selectivity in low- versus high-order electrodes argues against this interpretation (Fig. 6). Thus, the responses of early visual electrodes were correlated with low-level image features, such as image contrast, but not with face-related features, whereas the responses of face-selective electrodes were correlated with facial features but not with elementary image parameters. House- and tool-selective electrodes showed an intermediate pattern that was related both to low-level and to face features. However, only ventral face-selective electrodes showed a significant correlation to the perceptual similarity between face images (Fig. 5). Finally, the lack of correlation between early visual and face electrodes' responses further renders a low-level explanation unlikely. Thus, we can conclude that the exemplar selectivity we find in high-order visual areas is unlikely to be merely the consequence of low-level physical differences between face images.

Principles of Neuronal Anatomical Organization Revealed by Exemplar Selectivity

The finding of exemplar selectivity in the ECoG responses offers an important new insight into the anatomical principles by which neurons in human high-order visual cortex are grouped. The logic behind this notion stems from

accumulating evidence that the ECoG signal, and particularly the broadband Gamma response, likely reflects the approximate aggregate of firing activity generated by a large population of neurons located in the vicinity of the recording electrode (Mukamel et al. 2005; Kreiman et al. 2006; Nir et al. 2007; Manning et al. 2009). The fact that despite this massive signal averaging, the ECoG signal still manifests exemplar selectivity bears important implications for the functional organization of neurons in the recording sites. Note that the observed exemplar selectivity in the ECoG response can be generated only if 2 separate constraints are satisfied. First, each neuron must show exemplar selectivity on its own, and second, neurons sharing similar exemplar tuning properties are likely to be, at least partially, grouped together. By identifying the stimuli that produced similar activation profiles in the ECoG responses, we could gain insight into the anatomical grouping principles of the underlying neurons.

Our results revealed a highly significant correlation between the neural response of face-selective electrodes and the perceptual similarity between faces. Simply put, faces that appeared more similar also elicited similar neural responses (Fig. 5B). Perceptual similarity was measured in the current study in 2 ways: an "ecological" manner, in which the participants clustered the face images according to their perceptual similarity, and a more controlled manner, in which the face images were briefly presented sequentially while the participants performed a same-different task. The correlation with the neural responses of face electrodes was significant for both measures but higher for the "ecological" measure. Note that in the sequential presentation subjects were required to bridge the percept across an interruption, a situation that occasionally may have given rise to a "change blindness" type of disruption (Beck et al. 2001). This potential difference should be addressed in future studies.

The correlation between neuronal activity and perceptual similarity reported here must be an underestimate of the true correlation. This is because our perceptual measures of face similarity were derived from a different subject population (Israeli students) compared with the patient population (American adults). However, it is important to emphasize that despite major potential differences between these 2 populations the correlation between perceptual and neuronal similarities was highly significant. This result is compatible with our previous research showing a similar across-group correlation in sensory-driven responses (Mukamel et al. 2005). It is also in line with previous nonhuman primates electrophysiology and human fMRI studies, showing that object shape representations reflect physical and perceptual similarities (Olson 2001; Op de Beeck et al. 2001, 2008; Haushofer et al. 2008; Drucker and Aguirre 2009). We now extend these findings to the category of faces, which are much more perceptually homogeneous and hence pose a more challenging visual problem. This result raises the intriguing possibility that the perceptual sense of similarity may be linked to the anatomical proximity (i.e., anatomical grouping) of high-order cortical neurons or to the tuning profile of individual neurons (Edelman 1998; Wang et al. 1998).

The responses of the ensemble of face-selective electrodes were also correlated with several face-related features, including total face area and hair length (Fig. 6). This is compatible with findings in the macaque emphasizing face

aspect ratio and iris size as important parameters (Freiwald et al. 2009). While external facial features (such as hair) are commonly excluded in face recognition studies, it has been shown that they exert an important role in face perception (Axelrod and Yovel 2010). Our study points to the important role that such external features play in face exemplar tuning.

Induced versus Evoked Activity

Previous ECoG studies have emphasized the category selectivity of high-order visual cortex, but have not shown within-category exemplar differentiation. It has been reported that face-specific evoked potentials in occipitotemporal cortex were not affected by face familiarity, by identification of the face or by image type (colored, grayscale, or line drawing of a face) (McCarthy et al. 1999). Of particular interest is a recent ECoG study by Liu et al. (2009) which reported that category but not exemplar information can be significantly decoded from high-order visual cortex. There are a few possible explanations for the inconsistency between Liu et al. and the present study. Most importantly, Liu et al. focused on evoked potentials, time-locked to the stimulus, and not on induced changes, in which the modulations in spectral power within a certain time window are measured. Our results show that exemplar selectivity was mainly manifested in the induced high-Gamma responses (compare Fig. 3 and Supplementary Fig. 4B) although a more sensitive classification analysis showed that evoked responses did discriminate to some extent between individual exemplars (Supplementary Fig. 5). The inconsistency between the 2 studies can also be accounted by different electrode selection criteria (category selectivity in our study vs. variance between categories divided by variance within each category in Liu et al.) and the number of exemplars per category (5–14 in our study vs. 5 in Liu et al.).

The higher exemplar sensitivity of induced high-Gamma responses relative to evoked responses is compatible with reports that the N170 potential, recorded in scalp-EEG, is insensitive to individual faces, leading to the notion that face individuation only occurs at later stages (Bentin and Deouell 2000; Amihai et al. 2011). Our data thus argue against this notion and emphasize the importance of induced activity in assessing perceptually related neural responses (Engell and McCarthy 2010, 2011; Vidal et al. 2010; Privman et al. 2011).

Ventral versus Lateral Face-Related Regions

The face selective electrodes identified in our study were divided anatomically into 2 clusters: a ventral cluster along the fusiform gyrus and a lateral occipital cluster (Allison et al. 1999). There were several important distinctions between these 2 clusters. First, the exemplar tuning was much sharper in ventral face electrodes compared with lateral face electrodes (Fig. 3). Second, the exemplar classification rate for faces was significantly higher than nonface exemplars only in ventral face electrodes. Finally, the correlation between neuronal responses and perceptual similarity was stronger for ventral face electrodes (Fig. 6). These findings are compatible with previous fMRI findings, which reported that the OFA is more sensitive to low-level features compared with the FFA (Grill-Spector et al. 1999). Other studies have found that the FFA and OFA differ in their category selectivity (Levy et al. 2001; Andrews and Ewbank 2004) and in their release from adaptation following gradual morphing between faces (Gilaie-

Dotan and Malach 2007). This ventral versus lateral distinction has also been demonstrated in the object domain (Haushofer et al. 2008).

In conclusion, our results reveal that high-order visual areas exhibit selectivity not only at the category level but also at the exemplar level. Exemplar selectivity follows general principles, such as perceptual distance, and offers important insights into the anatomical clustering of neurons in high-order visual areas.

Supplementary Material

Supplementary material can be found at: http://www.cercor.oxford journals.org/.

References

Allison T, McCarthy G, Nobre A, Puce A, Belger A. 1994. Human extrastriate visual cortex and the perception of faces, words, numbers, and colors. Cereb Cortex. 4:544–554.

Allison T, Puce A, Spencer DD, McCarthy G. 1999. Electrophysiological studies of human face perception. I: Potentials generated in occipito-temporal cortex by face and non-face stimuli. Cereb Cortex. 9:415–430.

Amihai I, Deouell LY, Bentin S. 2011. Neural adaptation is related to face repetition irrespective of identity: a reappraisal of the N170 effect. Exp Brain Res. 1–12.

Andrews TJ, Ewbank MP. 2004. Distinct representations for facial identity and changeable aspects of faces in the human temporal lobe. Neuroimage. 23:905–913.

Axelrod V, Yovel G. 2010. External facial features modify the representation of internal facial features in the fusiform face area. Neuroimage. 52:720–725.

Baylis G, Rolls ET, Leonard C. 1985. Selectivity between faces in the responses of a population of neurons in the cortex in the superior temporal sulcus of the monkey. Brain Res. 342:91–102.

Beck DM, Rees G, Frith CD, Lavie N. 2001. Neural correlates of change detection and change blindness. Nat Neurosci. 4:645–650.

Bentin S, Allison T, Puce A, Perez E, McCarthy G. 1996. Electrophysiological studies of face perception in humans. J Cogn Neurosci. 8:551–565.

Bentin S, Deouell LY. 2000. Structural encoding and identification in face processing: ERP evidence for separate mechanisms. Cogn Neuropsychol. 2000:35–54.

Burns SP, Xing D, Shapley RM. 2010. Comparisons of the dynamics of local field potential and multiunit activity signals in macaque visual cortex. J Neurosci. 30:13739–13749.

Canolty RT, Edwards E, Dalal SS, Soltani M, Nagarajan SS, Kirsch HE, Berger MS, Barbaro NM, Knight RT. 2006. High gamma power is phase-locked to theta oscillations in human neocortex. Science. 313:1626–1628.

Crone NE, Sinai A, Korzeniewska A. 2006. High-frequency gamma oscillations and human brain mapping with electrocorticography. Prog Brain Res. 159:275–295.

Delorme A, Makeig S. 2004. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J Neurosci Methods. 134:9–21.

DeYoe EA, Carman GJ, Bandettini P, Glickman S, Wieser J, Cox R, Miller D, Neitz J. 1996. Mapping striate and extrastriate visual areas in human cerebral cortex. Proc Natl Acad Sci. 93:2382–2386.

Dricot L, Sorger B, Schiltz C, Goebel R, Rossion B. 2008. The roles of "face" and "non-face" areas during individual face perception: evidence by fMRI adaptation in a brain-damaged prosopagnosic patient. Neuroimage. 40:318–332.

Drucker DM, Aguirre GK. 2009. Different spatial scales of shape similarity representation in lateral and ventral LOC. Cereb Cortex. 19:2269–2280.

Duda RO, Hart PE, Stork DG. 2001. Pattern Classification. A Wiley-Interscience Publication. New York: John Wiley & Sons, Inc.

Edelman S. 1998. Representation is representation of similarities. Behav Brain Sci. 21:449–467.

Eger E, Ashburner J, Haynes JD, Dolan RJ, Rees G. 2008. fMRI activity patterns in human LOC carry information about object exemplars within category. J Cogn Neurosci. 20:356–370.

Engel SA, Rumelhart DE, Wandell BA, Lee AT. 1994. fMRI of human visual cortex. Nature. 369:525.

Engell AD, McCarthy G. 2011. The relationship of gamma oscillations and face-specific ERPs recorded subdurally from occipitotemporal cortex. Cereb Cortex. 21:1213–1221.

Engell AD, McCarthy G. 2010. Selective attention modulates face-specific induced gamma oscillations recorded from ventral occipitotemporal cortex. J Neurosci. 30:8780–8786.

Fisch L, Privman E, Ramot M, Harel M, Nir Y, Kipervasser S, Andelman F, Neufeld MY, Kramer U, Fried I. 2009. Neural "ignition": enhanced activation linked to perceptual awareness in human ventral stream visual cortex. Neuron. 64:562–574.

Freiwald WA, Tsao DY. 2010. Functional compartmentalization and viewpoint generalization within the macaque face-processing system. Science. 330:845–851.

Freiwald WA, Tsao DY, Livingstone MS. 2009. A face feature space in the macaque temporal lobe. Nat Neurosci. 12:1187–1196.

Gilaie-Dotan S, Gelbard-Sagiv H, Malach R. 2010. Perceptual shape sensitivity to upright and inverted faces is reflected in neuronal adaptation. Neuroimage. 50:383–395.

Gilaie-Dotan S, Malach R. 2007. Sub-exemplar shape tuning in human face-related areas. Cereb Cortex. 17:325–338.

Grill-Spector K, Kushnir T, Edelman S, Avidan G, Itzchak Y, Malach R. 1999. Differential processing of objects under various viewing conditions in the human lateral occipital complex. Neuron. 24:187–203.

Hasselmo ME, Rolls ET, Baylis GC. 1989. The role of expression and identity in the face-selective responses of neurons in the temporal visual cortex of the monkey. Behav Brain Res. 32:203–218.

Hasson U, Harel M, Levy I, Malach R. 2003. Large-scale mirror-symmetry organization of human occipito-temporal object areas. Neuron. 37:1027–1041.

Haushofer J, Livingstone MS, Kanwisher N. 2008. Multivariate patterns in object-selective cortex dissociate perceptual and physical shape similarity. PLoS Biol. 6:e187.

Hung CP, Kreiman G, Poggio T, DiCarlo JJ. 2005. Fast readout of object identity from macaque inferior temporal cortex. Science. 310:863–866.

Ishai A, Ungerleider LG, Martin A, Schouten JL, Haxby JV. 1999. Distributed representation of objects in the human ventral visual pathway. Proc Natl Acad Sci. 96:9379–9384.

Kahn DA, Harris AM, Wolk DA, Aguirre GK. 2010. Temporally distinct neural coding of perceptual similarity and prototype bias. J Vision. 10:1–12.

Kanwisher N, McDermott J, Chun MM. 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. J Neurosci. 17:4302–4311.

Kreiman G, Hung CP, Kraskov A, Quiroga RQ, Poggio T, DiCarlo JJ. 2006. Object selectivity of local field potentials and spikes in the macaque inferior temporal cortex. Neuron. 49:433–445.

Kriegeskorte N, Formisano E, Sorger B, Goebel R. 2007. Individual faces elicit distinct response patterns in human anterior temporal cortex. Proc Natl Acad Sci. 104:20600–20605.

Lachaux JP, George N, Tallon-Baudry C, Martinerie J, Hugueville L, Minotti L, Kahane P, Renault B. 2005. The many faces of the gamma band response to complex visual stimuli. Neuroimage. 25:491–501.

Levy I, Hasson U, Avidan G, Hendler T, Malach R. 2001. Center-periphery organization of human object areas. Nat Neurosci. 4:533–539.

Levy I, Hasson U, Harel M, Malach R. 2004. Functional analysis of the periphery effect in human building related areas. Hum Brain Mapp. 22:15–26.

Liu H, Agam Y, Madsen JR, Kreiman G. 2009. Timing, timing, timing: fast decoding of object information from intracranial field potentials in human visual cortex. Neuron. 62:281–290.

Malach R. 2012. Targeting the functional properties of cortical neurons using fMR-adaptation. Neuroimage. 62:1163–1169.

Manning JR, Jacobs J, Fried I, Kahana MJ. 2009. Broadband shifts in local field potential power spectra are correlated with single-neuron spiking in humans. J Neurosci. 29:13613–13620.

McCarthy G, Puce A, Belger A, Allison T. 1999. Electrophysiological studies of human face perception. II: Response properties of face-specific potentials generated in occipitotemporal cortex. Cereb Cortex. 9:431–444.

McCarthy G, Puce A, Gore JC, Allison T. 1997. Face-specific processing in the human fusiform gyrus. J Cogn Neurosci. 9:605–610.

Mukamel R, Gelbard H, Arieli A, Hasson U, Fried I, Malach R. 2005. Coupling between neuronal firing, field potentials, and FMRI in human auditory cortex. Science. 309:951–954.

Mur M, Ruff DA, Bodurka J, Bandettini PA, Kriegeskorte N. 2010. Face-identity change activation outside the face system: "release from adaptation" may not always indicate neuronal selectivity. Cereb Cortex. 20:2027–2042.

Mur M, Ruff DA, Bodurka J, De Weerd P, Bandettini PA, Kriegeskorte N. 2012. Categorical, yet graded–single-image activation profiles of human category-selective cortical regions. J Neurosci. 32:8649–8662.

Nestor A, Plaut D, Behrmann M. 2011. Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. Proc Natl Acad Sci USA. 108:9998–10003.

Nir Y, Fisch L, Mukamel R, Gelbard-Sagiv H, Arieli A, Fried I, Malach R. 2007. Coupling between neuronal firing rate, gamma LFP, and BOLD fMRI is related to interneuronal correlations. Curr Biol. 17:1275–1285.

Olson CR. 2001. Object-based vision and attention in primates. Curr Opin Neurobiol. 11:171–179.

Op de Beeck H, Wagemans J, Vogels R. 2001. Inferotemporal neurons represent low-dimensional configurations of parameterized shapes. Nat Neurosci. 4:1244–1252.

Op de Beeck HP, Brants M, Baeck A, Wagemans J. 2010. Distributed subordinate specificity for bodies, faces, and buildings in human ventral visual cortex. Neuroimage. 49:3414–3425.

Op de Beeck HP, Torfs K, Wagemans J. 2008. Perceived shape similarity among unfamiliar objects and the organization of the human object vision pathway. J Neurosci. 28:10111–10123.

Pereira F, Mitchell T, Botvinick M. 2009. Machine learning classifiers and fMRI: a tutorial overview. Neuroimage. 45:199–209.

Privman E, Fisch L, Neufeld MY, Kramer U, Kipervasser S, Andelman F, Yeshurun Y, Fried I, Malach R. 2011. Antagonistic relationship between gamma power and visual evoked potentials revealed in human visual cortex. Cereb Cortex. 21:616–624.

Privman E, Nir Y, Kramer U, Kipervasser S, Andelman F, Neufeld MY, Mukamel R, Yeshurun Y, Fried I, Malach R. 2007. Enhanced category tuning revealed by intracranial electroencephalograms in high-order human visual areas. J Neurosci. 27:6234–6242.

Ray S, Niebur E, Hsiao SS, Sinai A, Crone NE. 2008. High-frequency gamma activity (80–150Hz) is increased in human cortex during selective attention. Clin Neurophysiol. 119:116–133.

Rolls ET, Cowey A, Bruce V. 1992. Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas [and discussion]. Philos Trans R Soc Lond B Biol Sci. 335:11–21.

Rotshtein P, Henson RNA, Treves A, Driver J, Dolan RJ. 2004. Morphing Marilyn into Maggie dissociates physical and identity face representations in the brain. Nat Neurosci. 8:107–113.

Sawamura H, Orban GA, Vogels R. 2006. Selectivity of neuronal adaptation does not match response selectivity: a single-cell study of the FMRI adaptation paradigm. Neuron. 49:307–318.

Sereno M, Dale A, Reppas J, Kwong K, Belliveau J, Brady T, Rosen B, Tootell R. 1995. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. Science. 268:889–893.

Talairach J, Tournoux P. 1988. Co-planar stereotaxic atlas of the human brain: 3-dimensional proportional system: an approach to cerebral imaging. New York: Thieme.

Tsao DY, Freiwald WA, Tootell RB, Livingstone MS. 2006. A cortical region consisting entirely of face-selective cells. Science. 311:670–674.

Tsao DY, Moeller S, Freiwald WA. 2008. Comparing face patch systems in macaques and humans. Proc Natl Acad Sci. 105:19514–19519.

Vidal JR, Ossandón T, Jerbi K, Dalal SS, Minotti L, Ryvlin P, Kahane P, Lachaux JP. 2010. Category-specific visual responses: an intracranial study comparing gamma, beta, alpha, and ERP response selectivity. Front Hum Neurosci. 4:1–23.

Wang G, Tanifuji M, Tanaka K. 1998. Functional architecture in monkey inferotemporal cortex revealed by in vivo optical imaging. Neurosci Res. 32:33–46.