



A face is more than just the eyes, nose, and mouth: fMRI evidence that face-selective cortex represents external features



Frederik S. Kamps, Ethan J. Morris, Daniel D. Dilks*

Department of Psychology, Emory University, Atlanta, GA 30322, USA

ARTICLE INFO

Keywords:

Face perception
fMRI
Fusiform face area (FFA)
Occipital face area (OFA)
Superior temporal sulcus (pSTS)

ABSTRACT

What is a face? Intuition, along with abundant behavioral and neural evidence, indicates that internal features (e.g., eyes, nose, mouth) are critical for face recognition, yet some behavioral and neural findings suggest that external features (e.g., hair, head outline, neck and shoulders) may likewise be processed as a face. Here we directly test this hypothesis by investigating how external (and internal) features are represented in the brain. Using fMRI, we found highly selective responses to external features (relative to objects and scenes) within the face processing system in particular, rivaling that observed for internal features. We then further asked *how* external and internal features are represented in regions of the cortical face processing system, and found a similar division of labor for both kinds of features, with the occipital face area and posterior superior temporal sulcus representing the parts of both internal and external features, and the fusiform face area representing the coherent arrangement of both internal and external features. Taken together, these results provide strong neural evidence that a “face” is composed of both internal and external features.

1. Introduction

Faces are the gateway to our social world. A face alone is enough to reveal a person's identity, gender, emotional state, and more. But what is a “face”, precisely? Common sense suggests that internal features like the eyes, nose, and mouth are particularly important, and dictionaries typically define a face based on these features. Moreover, behavioral experiments have widely demonstrated our remarkable sensitivity to internal features (e.g., Thompson, 1980; Tanaka and Sengco, 1997; Farah et al., 1998), computer scientists have designed vision systems that primarily process internal features (e.g., Brunelli and Poggio, 1993), and the vast majority of fMRI studies have tested representation of internal features only (Tong et al., 2000; Yovel and Kanwisher, 2004; Schiltz and Rossion, 2006; Maurer et al., 2007; Schiltz et al., 2010; Arcurio et al., 2012; Zhang et al., 2012; James et al., 2013; Lai et al., 2014; Zhao et al., 2014; de Haas et al., 2016; Nestor et al., 2016). Intriguingly, however, some behavioral evidence suggests that a face is more than the internal features alone (Young et al., 1987; Rice et al., 2013; Abudarham and Yovel, 2016; Hu et al., 2017). For example, in the classic “presidential illusion” (Sinha and Poggio, 1996), the same internal features are placed within the heads and bodies of Bill Clinton and Al Gore, yet viewers readily recognize “Bill Clinton” and “Al Gore” using the external features

only. Further work suggests that external and internal features are processed in a similar manner; for example, external features, like internal features, are particularly difficult to recognize when inverted (Moscovitch and Moscovitch, 2000; Brandman and Yovel, 2012). Taken together, these findings suggest that a face is composed of external features, along with internal features.

However, despite such behavioral work suggesting that both internal and external features are part of face representation, the possibility remains that external features are not represented in the same neural system as internal features (i.e., within the cortical face processing system), but rather are represented in a different neural system (e.g., for object or body processing), and consequently, that only internal features, not external features, are represented as part of a face. Accordingly, a promising approach toward unraveling which features make up a face would be to test directly whether and how both internal and external features are represented in the brain. Indeed, a handful of studies have taken this approach, and claimed to have found external feature representation in face-selective cortex (Liu et al., 2009; Andrews et al., 2010; Axelrod and Yovel, 2010; Betts and Wilson, 2010). Critically, however, none of these studies has established whether external features, like internal features, are represented *selectively* within the cortical face processing system, leaving open the question of whether external features,

* Corresponding author. Department of Psychology Emory University 36 Eagle Row Atlanta, GA 30322, USA.

E-mail address: dilks@emory.edu (D.D. Dilks).

like internal features, are part of face representation. More specifically, the majority of these studies did not compare responses to internal and external features with those to a non-face control condition (e.g., objects). Given that these studies generally find weaker responses to external than internal features, it is unclear then whether the weaker response to external features is nevertheless a selective response (i.e., with face-selective cortex responding significantly more to external features than objects). In fact, it could be the case that the response to external features is similar to the response to objects, and consequently that only internal features are selectively represented in face-selective cortex. Beginning to address this question, one of the above studies (Liu et al., 2009) found a greater response to external features than scenes. However, given that face-selective cortex is known to respond more to objects than scenes (of course, with a greater response to faces than to either of these categories), these findings still do not answer the question of whether this pattern reflects face selectivity per se, or a more general preference for any object over scenes. Closer still, one EEG study (Eimer, 2000) found larger N170 amplitudes for isolated internal and external features than for houses, and crucially, hands; however, this study investigated only the face-selective temporal electrodes (T5 and T6), and no control electrodes; thus, it is unclear whether this finding reflects actual face selectivity or general attention. Taken together then, the selectivity of the cortical face processing system for external features is not yet established, and thus the question remains whether external features, like internal features, compose a face.

Here we present the strongest test of the hypothesis that external features, not just internal features, are part of face representation, by comparing fMRI activation in the cortical face processing system to isolated internal and external features with that to objects and scenes. We predicted that if a “face” includes both internal and external features, then face-selective regions, including the occipital face area (OFA), fusiform face area (FFA), and posterior superior temporal sulcus (pSTS), should respond strongly and selectively to both isolated internal and isolated external features, compared to objects and scenes.

Finally, in order to ultimately understand face processing, we need to understand not only which features (e.g., internal and/or external) make up a face, but also the more precise nature of the representations extracted from those features. To our knowledge, no previous study has explored *how* external features are represented in the cortical face processing system. By contrast, studies of internal feature representation have found a division of labor across the three face-selective regions, with OFA and pSTS representing the parts of faces, and FFA representing the canonical, “T-shape” configuration of face parts (Pitcher et al., 2007; Harris and Aguirre, 2008, 2010; Liu et al., 2009). While this division of labor has been shown for internal features, it has never been tested for external features, allowing us to explore for the first time whether face-selective cortex exhibits a similar division of labor for external features as internal features, further supporting the hypothesis that internal features, like external features, are part of face representation.

2. Methods

2.1. Participants

Twenty participants (Age: 21–38; mean age: 27.6; 8 male, 11 female, 1 other) were recruited for this experiment. All participants gave informed consent and had normal or corrected-to-normal vision. Procedures for the study were approved by the Emory Institutional Review Board.

2.2. Design

We used a region of interest (ROI) approach in which we used one set of runs to localize category-selective regions (Localizer runs), and a second set of runs to investigate the responses of these same voxels to the experimental conditions (Experimental runs). For both Localizer and

Experimental runs, participants performed a one-back task, responding every time the same image was presented twice in a row. In addition to the standard ROI analysis, we conducted a novel “volume-selectivity function” (VSF) analysis, which is described in the Data Analysis Section.

For the Localizer runs, a blocked design was used in which participants viewed images of faces (including internal and external features), bodies, objects, scenes, and scrambled objects. Each participant completed 3 localizer runs. Each run was 400s long and consisted of 4 blocks per stimulus category. The order of blocks in each run was palindromic (e.g., faces, bodies, objects, scenes, scrambled objects, scrambled objects, scenes, objects, bodies, faces, etc.) and the order of blocks in the first half of the palindromic sequence was pseudorandomized across runs. Each block contained 20 images from the same category for a total of 16 s blocks. Each image was presented for 300 ms, followed by a 500 ms interstimulus interval, and subtended $8 \times 8^\circ$ of visual angle. We also included five 16 s fixation blocks: one at the beginning, three in the middle interleaved between each palindrome, and one at the end of each run.

For the Experimental runs, participants viewed runs during which 16 s blocks (20 stimuli per block) of 8 categories of images were presented (six conditions were used for the current experiment, while two additional “scene” categories tested unrelated hypotheses about scene processing) (Fig. 1). Each image was presented for 300 ms, followed by a 500 ms interstimulus interval, and subtended $8 \times 8^\circ$ of visual angle. Participants viewed 8 runs, and each run contained 21 blocks (2 blocks of each condition, plus 5 blocks of fixation), totaling 336s. The order of blocks in each run was palindromic, and the order of blocks in the first half of the palindromic sequence was pseudorandomized across runs. As depicted in Fig. 1, the six categories of interest were: (1) arranged internal features with no external features (i.e. eyes, nose, and mouth only, arranged into their canonical “T” configuration); (2) rearranged internal features with no external features (i.e. the same eyes, nose, and mouth, but rearranged such that they no longer form a coherent T-shape); (3) arranged external features with no internal features (i.e. hair, head outline, and neck/shoulders only, arranged in a coherent configuration); (4) rearranged external features with no internal features (i.e. the same hair, head outline, and neck/shoulders, but rearranged such that they no longer form a coherent configuration); (5) objects (multiple objects, matching the multiple face parts shown in the internal and external feature conditions); and (6) scenes (empty apartment rooms). Images used to create stimuli for the four face conditions were drawn from the Radboud Faces Database (Langner et al., 2010), while the object and scene stimuli were the same as those used two previous studies (Epstein and Kanwisher, 1998; Kamps et al., 2016). Internal and external features were parcellated based on linguistic conventions and natural physical boundaries, which we established with pilot behavioral data. For example, there are clear words and natural physical boundaries between the hair, head, and neck/shoulders; indeed, pilot participants who were instructed to simply “label particular features on the image” (when viewing versions of our “face” stimuli that now included both internal and external features) spontaneously labeled the “neck,” “shoulders,” “chin,” and “hair.” Next, given that there is no clear physical boundary between the “neck” and “shoulders,” these features were grouped as a single unit. Likewise, given that there is no clear physical boundary between the chin and the rest of the head (sans internal features), this entire extent was treated as a single unit. Finally, the hair could clearly be separated from the rest of the head, leaving this as a third distinct unit.

2.3. fMRI scanning

All scanning was performed on a 3T Siemens Trio scanner in the Facility for Education and Research in Neuroscience at Emory. Functional images were acquired using a 32-channel head matrix coil and a gradient-echo single-shot echoplanar imaging sequence (35 slices, TR = 2 s, TE = 30 ms, voxel size = $3 \times 3 \times 3$ mm, and a 0.3 mm interslice gap). For all scans, slices were oriented approximately between

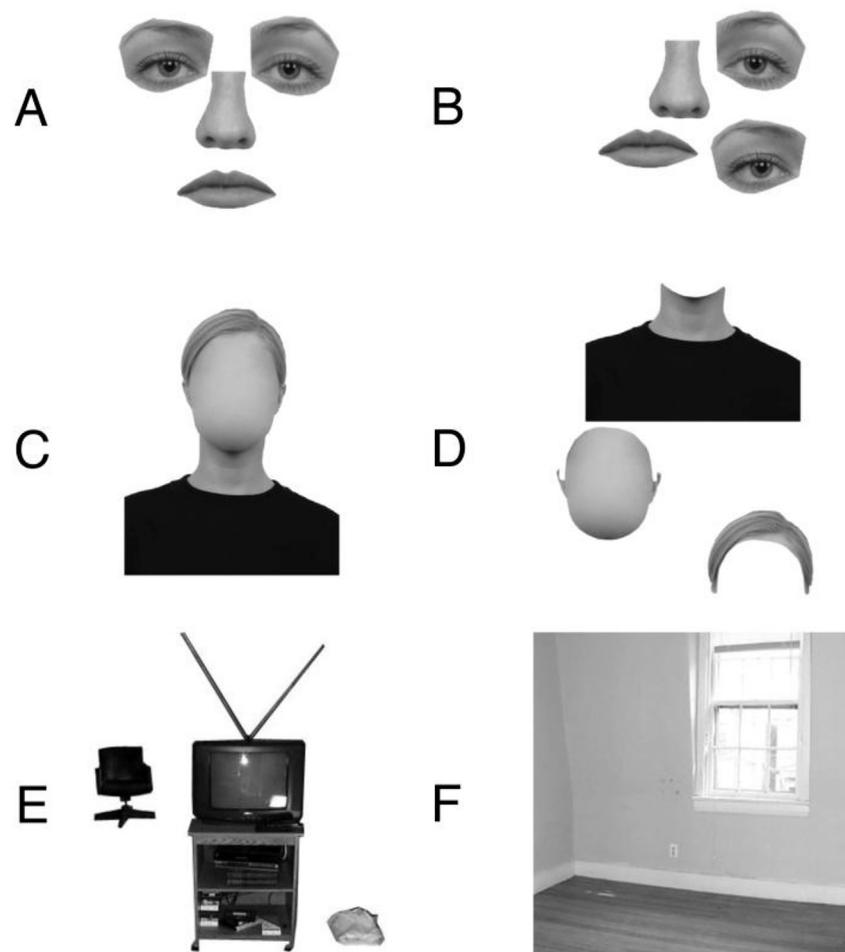


Fig. 1. Example stimuli used in the Experimental runs. A) Isolated internal features (arranged). B) Isolated internal features (rearranged). C) Isolated external features (arranged). D) Isolated external features (rearranged). E) Objects. F) Scenes.

perpendicular and parallel to the calcarine sulcus, covering all of the occipital and temporal lobes, as well as the majority of the parietal and frontal lobes. Whole-brain, high-resolution anatomical images were also acquired for each participant for purposes of registration and anatomical localization.

2.4. Data analysis

MRI data analysis was conducted using a combination of tools from the FSL software (FMRIB's Software Library: www.fmrib.ox.ac.uk/fsl) (Smith et al., 2004), the FreeSurfer Functional Analysis Stream (FS-FAST; <http://surfer.nmr.mgh.harvard.edu/>), and custom written MATLAB code. Before statistical analysis, images were motion corrected (Cox and Jesmanowicz, 1999), detrended, intensity normalized, and fit using a double gamma function. Localizer data, but not Experimental data, were spatially smoothed with a 5 mm kernel (except where noted otherwise). Experimental data were not smoothed to prevent inclusion of information from adjacent, non-selective voxels, an approach taken in many previous studies (e.g., Yovel and Kanwisher, 2004; Pitcher et al., 2011; Persichetti and Dilks, 2016), although the results did not qualitatively change when analyses were performed on smoothed data. After pre-processing, face-selective regions OFA, FFA, and pSTS were defined in each participant (using data from the independent Localizer scans) as those regions that responded more strongly to faces than objects ($p < 10^{-4}$, uncorrected), as described previously (Kanwisher et al., 1997). OFA and pSTS were identified in at least one hemisphere in all participants, while FFA was identified in at least one hemisphere in

19/20 participants. Importantly, since FFA is adjacent to and sometimes overlapping with the body-selective fusiform body area (FBA; Downing et al., 2001), we additionally defined the FBA as the region responding significantly more to bodies than objects ($p < 10^{-4}$, uncorrected), and then subtracted any overlapping FBA voxels from FFA, ensuring that the definition of FFA included face-selective voxels only. Finally, we functionally defined three control regions: the lateral occipital complex (LOC), defined as the region responding significantly more to objects than scrambled objects ($p < 10^{-4}$, uncorrected) (Grill-Spector et al., 1998); the parahippocampal place area (PPA), defined as the region responding significantly more to scenes than objects ($p < 10^{-4}$, uncorrected) (Epstein and Kanwisher, 1998); and foveal cortex (FC), defined as the region along the calcarine sulcus responding significantly more to all conditions than fixation ($p < 10^{-4}$, uncorrected). Localization of each category-selective ROI was confirmed by comparison with a published atlas of “parcels” that identify the anatomical regions within which most subjects show activation for each contrast of interest (i.e., for face-, object-, and scene-selective ROIs) (Julian et al., 2012). Within each ROI, we then calculated the magnitude of response (percent signal change, or PSC) to each category of interest, using independent data from the Experimental runs. Percent signal change was calculated by normalizing the parameter estimates for each condition relative to the mean signal intensity in each voxel (Mumford, 2007). A 3 (region: OFA, FFA, pSTS) X 2 (hemisphere: Left, Right) X 4 (condition: arranged external features, rearranged external features, arranged internal features, rearranged internal features) repeated-measures ANOVA did not reveal a significant region x condition x hemisphere interaction ($F_{(10,120)} = 1.00$, $p = 0.45$,

$\eta_p^2 = 0.04$). Thus, both hemispheres were collapsed for all further analyses.

In addition to the standard ROI analysis described above, we also conducted a novel VSF analysis to explore the full function relating responses in each ROI to its volume (see also Norman-Haignere et al., 2016; Saygin et al., 2016). Specifically, we identified a unique anatomical search space for each face-selective ROI, including the ROI plus the surrounding cortex (henceforth referred to as “OFA,” “FFA,” and “pSTS”), using anatomical divisions from a standard atlas (i.e., the Harvard-Oxford anatomical atlas provided within FSL) (Desikan et al., 2006). The search space for each face-selective ROI was created based on the following combinations of anatomical parcels: “OFA” = lateral occipital cortex (inferior division) and occipital fusiform gyrus; “FFA” = temporal occipital fusiform cortex (posterior division), inferior temporal gyrus (temporooccipital part and posterior division); “pSTS” = superior temporal gyrus (posterior division), middle temporal gyrus (temporooccipital part and posterior division), supramarginal gyrus (posterior division), angular gyrus, and lateral occipital cortex (superior division). Next, each voxel within each search space was assigned to a 5-voxel bin based on the rank order of the statistical significance of that voxel’s response to faces > objects from the Localizer scans (here not smoothed). The response of each bin to the conditions of interest was then measured in these same voxels using independent data from the Experimental runs, and responses were averaged across each increasing number of bins (e.g., such that responses for “bin 2” equaled the average of voxels in bins 1 and 2, while responses in “bin 3” equaled the average of bins 1, 2, and 3, and so on down the entire extent of each ROI and beyond). The responses to each condition for each bin were then averaged across participants, ultimately resulting in a group average VSF for each condition. Notably, because this analysis uses only a rank ordering of significance of the voxels in each subject, not an absolute threshold for voxel inclusion, all participants can be included in the analysis – not only those who show the ROI significantly.

3. Results

3.1. The cortical face processing system responds strongly and selectively to both internal and external features

If a “face” includes both internal and external features, then face-selective regions should respond strongly and selectively to both internal and external features. To test this prediction, we compared the average response in OFA, FFA, and pSTS to four conditions: arranged internal features, arranged external features, objects, and scenes (Fig. 2). For all three regions, a four-level repeated-measures ANOVA revealed a significant main effect (all F ’s > 38.69, all p ’s < 0.001, all η_p^2 ’s > 0.67), with OFA, FFA, and pSTS each responding significantly more to internal features than to both objects (all p ’s < 0.001, Bonferroni corrected) and scenes (all p ’s < 0.001, Bonferroni corrected)—consistent with previous reports of internal feature representation in these regions (e.g., Yovel and Kanwisher, 2004; Pitcher et al., 2007; Liu et al., 2009). Further, both OFA and FFA responded significantly more to objects than scenes (all p ’s < 0.001, Bonferroni corrected), while pSTS responded marginally significantly more to objects than scenes ($p = 0.06$, Bonferroni corrected; $p = 0.01$, uncorrected) – confirming that objects are indeed a more stringent control category than scenes (Kanwisher et al., 1997; Downing et al., 2006; Pitcher et al., 2011; Deen et al., 2017). Critically, all three face-selective regions responded significantly more to external features than to both objects (all p ’s < 0.001, Bonferroni corrected) and scenes (all p ’s < 0.001, Bonferroni corrected)—revealing that these regions likewise represent external features. Finally, neither OFA, FFA, nor pSTS responded significantly differently to internal and external features (all p ’s > 0.14, Bonferroni corrected). Thus, these findings reveal striking selectivity for external features, with all three face-selective regions responding two to four times more to external features than to objects or scenes.

Given that all three regions showed a similar response profile, it is

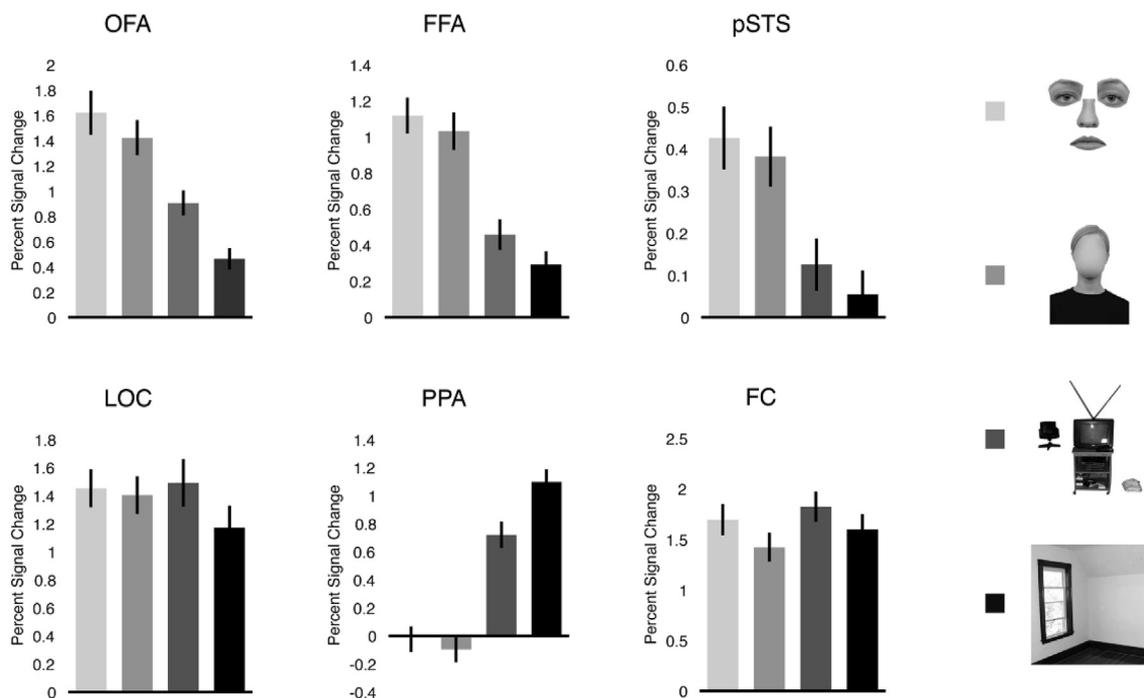


Fig. 2. Average percent signal change in each region of interest to the four conditions testing internal and external feature selectivity. OFA, FFA, and pSTS each responded significantly more to internal and external features than both objects and scenes (all p ’s < 0.001, Bonferroni corrected). None of these regions responded significantly differently to internal and external features (all p ’s > 0.14, Bonferroni corrected). By contrast, LOC responded to more to objects than scenes ($p < 0.001$, Bonferroni corrected), and similarly to internal features, external features, and objects (all p ’s = 1.00, Bonferroni corrected); PPA responded significantly more to scenes than any other condition (all p ’s < 0.001, Bonferroni corrected); and FC responded more to objects than all other conditions (all p ’s < 0.05, Bonferroni corrected), and weaker to external features than all other conditions (all p ’s < 0.01, Bonferroni corrected). Error bars represent the standard error of the mean.

possible that more general factors (e.g., differences in attention or lower-level visual features across the four conditions) could explain these results, rather than face-specific processing per se. To address this possibility, we compared responses in these regions with those in three control ROIs (i.e., LOC, PPA, and FC)—each of which would be expected to show a distinct response profile across the four conditions relative to face-selective cortex, if indeed our results cannot be explained by differences in attention or low-level visual features across the stimuli. For LOC, a four-level repeated-measures ANOVA revealed a significant main effect of condition ($F_{(3,57)} = 3.72, p < 0.05, \eta_p^2 = 0.16$), with a greater response to objects than scenes ($p < 0.001$, Bonferroni corrected), consistent with the well known object selectivity of this region. We found no difference in LOC response to either internal features and objects ($p = 1.00$, Bonferroni corrected) or external features and objects ($p = 1.00$, Bonferroni corrected), consistent with the hypothesis that LOC is sensitive to object shape across “object” domains (Kourtzi and Kanwisher, 2001; Grill-Spector, 2003). For PPA, a four-level repeated-measures ANOVA likewise revealed a significant main effect of condition ($F_{(3,57)} = 253.14, p < 0.001, \eta_p^2 = 0.93$), with a greater response to scenes than all other conditions (all p 's < 0.001 , Bonferroni corrected), reflecting the strong scene selectivity of this region. For FC, this same analysis revealed a significant main effect of condition ($F_{(3,57)} = 15.69, p < 0.001, \eta_p^2 = 0.45$), with a greater response to objects than all other conditions (all p 's < 0.05 , Bonferroni corrected), and a weaker response to external features than all other conditions (all p 's < 0.01 , Bonferroni corrected)—a distinct pattern from that in any other region, presumably reflecting differences in low-level visual information across the stimuli. Critically, a 6 (region: OFA, FFA, pSTS, LOC, PPA, FC) \times 4 (condition: internal features, external features, objects, scenes) repeated-measures ANOVA revealed a significant interaction ($F_{(15,240)} = 67.18, p < 0.001, \eta_p^2 = 0.81$), with a significantly different pattern of response between each face-selective region and all three control regions (interaction contrasts, all p 's < 0.05). Thus, the strong response of face-selective regions to external features cannot be explained by domain-general factors or low-level visual features, and rather reflects processing within the face processing system in particular.

But might the selectivity for external features in face-selective cortex result from the relatively high threshold we used to localize our ROIs, which included only the most face-selective voxels in each region? Indeed, a stronger test of face selectivity would ask whether a strong and selective response to external features (and internal features) is found at any threshold used to define the ROIs, not just the particular threshold we used. To address this question, we conducted a novel VSF analysis, which characterizes the complete function relating responses in each condition to ROI volume. Specifically, within each anatomical search space for each ROI (i.e., “OFA,” “FFA,” and “pSTS”), we calculated

responses to the four conditions above (i.e., internal features, external features, objects, scenes) in increasing numbers of “bins” of voxels (5 voxels per bin), ranked based on their response to faces $>$ objects from the independent Localizer runs (see Methods). The number of bins examined for each anatomical search space was chosen to be approximately twice the average size of each region in the ROI analysis above, allowing exploration of the pattern of responses across the entirety of the region itself, as well as the region of similar volume just beyond our ROI definition.

Complete VSFs for each ROI are shown in Fig. 3. As expected, face selectivity decreased in all three regions with increasing size, demonstrating the bin-by-bin within-subject replicability of face selectivity between the Localizer and Experimental runs. To statistically evaluate the pattern of responses across bins in each region, we performed statistical tests at five different sizes, ranging from the region including only the single most selective bin to the region including all bins, with even steps tested in between, calibrated to the average size of each region. At each size of “OFA,” “FFA,” and “pSTS,” we observed a significant main effect of condition (all F 's > 28.09 , all p 's < 0.001 , all η_p^2 's > 0.59), with i) a significantly greater response to both internal and external features compared to either objects or scenes (all p 's > 0.05 , Bonferroni corrected), ii) a significantly greater response to objects than scenes (all p 's < 0.01 , Bonferroni corrected), and iii) no significant difference in the response to internal and external features (all p 's > 0.34 , Bonferroni corrected). Qualitative inspection of the VSF in each region confirmed that this pattern of results is highly consistent across the entire volume of each face-selective ROI. These findings therefore reveal that selectivity for external features (and internal features) does not depend upon the threshold used for ROI definition, but rather extends across the entire volume of each face-selective region.

3.2. The cortical face processing system shows a similar division of labor across internal and external features

The cortical face processing system shows a similar division of labor across internal and external features

The findings above provide strong neural evidence that a face includes external features, not just internal features. But how is such external and internal feature information represented across the cortical face processing system? Previous work has found that OFA, FFA, and pSTS represent internal feature information differently, with OFA and pSTS representing the parts of internal features (e.g., eyes, nose, mouth), and FFA representing the coherent arrangement of internal features (e.g., two eyes above a nose above a mouth) (Pitcher et al., 2007; Liu et al., 2009). However, this division of labor has not been tested for external features. The finding that face-selective cortex exhibits a similar division

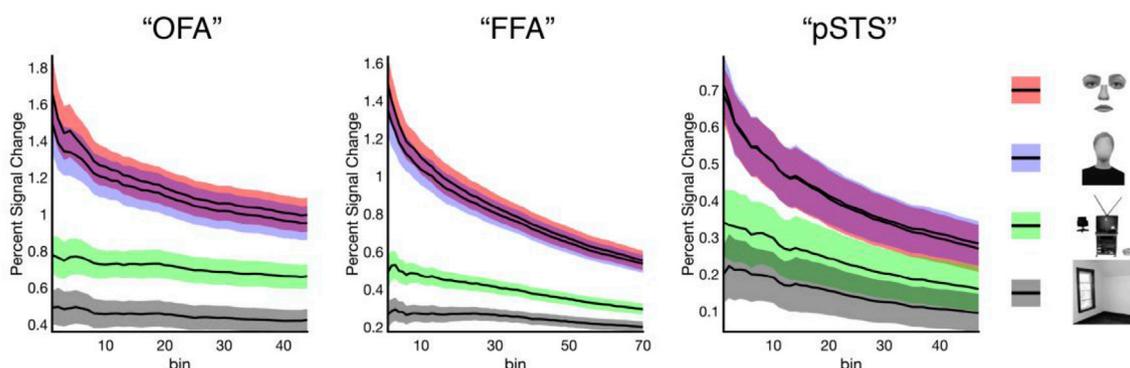


Fig. 3. Volume-selectivity function (VSF) analysis exploring external (and internal) feature selectivity in each face-selective region. The Localizer data were used to rank bins of voxels (5 voxels/bin) anywhere within a larger anatomical search space for each face-selective region (i.e., “OFA,” “FFA,” and “pSTS”) based on their response to faces $>$ objects. Experimental data were then used to extract the activation in each consecutive set of bins to internal features (arranged), external features (arranged), objects, and scenes, revealing that the strong selectivity for external features is found regardless of the threshold used for ROI definition. Colored bands around each line indicate the standard error of the mean.

of labor for external features as internal features would further support the hypothesis that external features, like internal features, are part of face representation.

Thus, to address this question, we compared responses in OFA, FFA, and pSTS to four conditions: i) arranged internal features, ii) rearranged internal features, iii) arranged external features, and iv) rearranged external features. If a region analyzes external or internal features at the level of parts, then it should respond similarly to such external or internal parts regardless of how they are arranged. By contrast, if a region is sensitive to the coherent arrangement of external or internal features, then it should respond more to external or internal features when arranged than when rearranged.

For OFA, a 2 (feature: internal, external) x 2 (arrangement: arranged, rearranged) repeated-measures ANOVA did not reveal a significant effect of arrangement ($F_{(1,19)} = 1.03$, $p = 0.32$, $\eta_p^2 = 0.05$) (Fig. 4), consistent with previous work showing that OFA represents the parts of faces (Pitcher et al., 2007; Liu et al., 2009). Critically, we further did not observe a significant feature x arrangement interaction ($F_{(1,19)} = 2.28$, $p = 0.15$, $\eta_p^2 = 0.11$), indicating that OFA analyzes the parts of internal and external features similarly.

Likewise, for pSTS, a 2 (feature: internal, external) x 2 (arrangement: arranged, rearranged) repeated-measures ANOVA did not reveal a significant effect of arrangement. ($F_{(1,19)} = 0.03$, $p = 0.87$, $\eta_p^2 = 0.001$) (Fig. 4), consistent with previous findings that pSTS represents the parts of faces (Liu et al., 2009). We did observe a significant feature x arrangement interaction ($F_{(1,19)} = 10.19$, $p < 0.01$, $\eta_p^2 = 0.35$). However, this interaction was driven by the difference in response to rearranged internal and rearranged external features ($p < 0.01$, Bonferroni corrected)—a contrast irrelevant to the question of whether pSTS represents the parts or wholes of internal and external features. By contrast, no other comparisons were significant (all p 's > 0.05 , Bonferroni corrected), including the critical comparisons between arranged and rearranged internal features, and arranged and rearranged external features. Therefore, this pattern of results suggests that pSTS extracts parts-based representations for both internal and external features.

Unlike OFA and pSTS, for FFA, a 2 (feature: internal, external) x 2 (arrangement: arranged, rearranged) repeated-measures ANOVA revealed a significant effect of arrangement ($F_{(1,16)} = 12.36$, $p < 0.01$, $\eta_p^2 = 0.44$) (Fig. 4), with FFA responding more to arranged than rearranged features, consistent with previous findings that FFA represents the coherent configuration of features (Liu et al., 2009). Crucially, we did not observe a significant feature x arrangement interaction ($F_{(1,16)} = 0.04$, $p = 0.85$, $\eta_p^2 = 0.002$), indicating that FFA shows similar sensitivity to the arrangement of internal and external features.

The analyses above suggest that the three face-selective regions represent face information differently, with OFA and pSTS representing internal and external parts, and FFA representing the coherent spatial

arrangement of internal and external parts. Testing this claim directly, a 3 (region: OFA, FFA, pSTS) x 2 (feature: internal, external) x 2 (arrangement: arranged, rearranged) repeated-measures ANOVA revealed a significant region x arrangement interaction ($F_{(2,32)} = 9.13$, $p < 0.001$, $\eta_p^2 = 0.36$), indicating that there were significant differences in sensitivity to arrangement information across the three regions. Interaction contrasts revealed that this interaction was driven by differences in responses to the arranged and rearranged conditions between regions that differ in the kind of face information represented (e.g., OFA and pSTS versus FFA, which represent parts and wholes of faces, respectively) (both p 's < 0.01), but not between regions that represent similar kinds of face information (e.g., OFA versus pSTS, which both represent face parts) ($p = 0.65$).

Further, a VSF analyses confirmed this pattern of parts- versus whole-based representation across the three regions (Fig. 5). At all region sizes tested, we found a significant region x arrangement interaction (all F 's > 14.56 , all p 's < 0.001 , all η_p^2 's > 0.43), with “FFA” responding significantly more to arranged than rearranged features relative to “OFA” and “pSTS” in every case (interaction contrasts, all p 's < 0.01). Taken together, these findings reveal that external and internal features show a similar pattern of parts vs. whole information processing across the face processing system, further supporting the hypothesis that a face includes both internal and external features.

3.3. Responses to external features do not reflect body-specific processing

An alternative account for our findings is that external features are not processed as part of the face per se, but rather are processed as part of the body. Indeed, our external feature stimuli included features that might more intuitively be thought of as part of the body than the face (i.e., the neck and shoulders). While this possibility is not likely the case, given that we removed body-selective voxels from our definition of the FFA, we nevertheless addressed this question directly by comparing the responses of the top face-selective voxels in the “FFA” search space (now including body-selective voxels), as well as the “OFA,” and “pSTS” search spaces, with those of the top body-selective voxels in these same search spaces. Specifically, we again conducted a VSF analysis, but now compared responses in bins of voxels ranked on their responses to faces $>$ objects (revealing face-selective regions) with those in bins of voxels ranked on their responses to bodies $>$ objects (revealing body-selective regions), using independent data from the Localizer runs. Critically, this approach allowed us to algorithmically assess body selectivity in and around the anatomical vicinity of each face-selective region, regardless of the extent to which these face- and body-selective ROIs overlapped, and regardless of whether a body-selective ROI could be defined in that vicinity using a standard threshold-based approach.

First, we tested the *selectivity* for external features in face- and body-

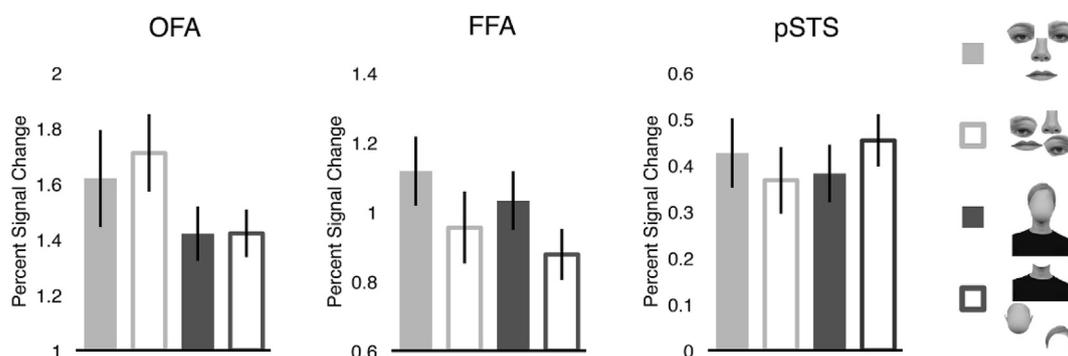


Fig. 4. Average percent signal change in each face-selective region to four conditions testing parts versus whole processing (i.e., arranged internal features, rearranged internal features, arranged external features, rearranged external features). OFA and pSTS responded similarly to arranged and rearranged face parts, both for internal and external features (both p 's > 0.32), whereas FFA was sensitive to the coherent arrangement of both internal and external features ($p < 0.01$), suggesting that all three regions extract similar kinds of representations across internal and external features. Error bars represent the standard error of the mean.

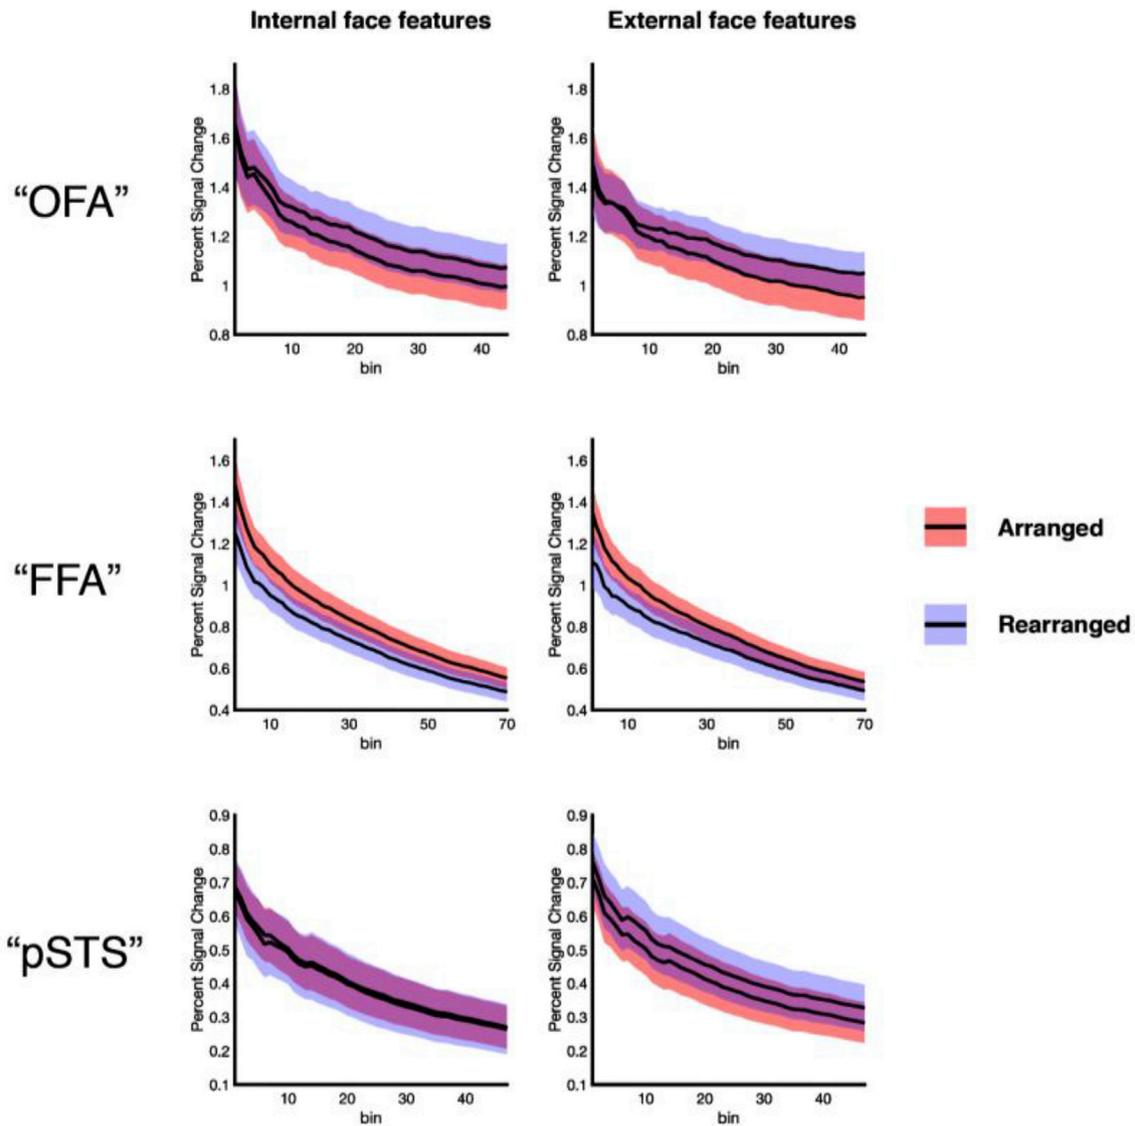


Fig. 5. VSF analysis exploring parts-based and whole-based representations in each face-selective region. The differential information processing performed across the three face-selective regions—with “OFA” and “pSTS” representing the parts of both internal and external of faces, and “FFA” representing the coherent arrangement of both internal and external features—is found at every threshold used to define the ROIs. Colored bands around each line indicate the standard error of the mean.

selective voxels within each search space (i.e., “OFA,” “FFA,” and “pSTS”). For all sizes in all three regions, a 2 (voxels: face-selective, body-selective) x 4 (condition: arranged internal features, arranged external features, objects and scenes) repeated-measures ANOVA revealed a significant interaction (all F 's > 7.97, all p 's < 0.001, all η_p^2 's > 0.29), with face-selective voxels responding significantly more to both internal, and crucially, external features than to objects, relative to body-selective voxels (interaction contrasts, all p 's < 0.05). These findings indicate that face-selective voxels are more selective for external features than nearby body-selective voxels in the same search space, supporting the hypothesis that external features are processed as a face (Fig. 6A; note that for clarity, the figure depicts only the pattern of responses found at the middle ROI size, representing the pattern found at all other sizes).

Second, we tested *how* internal and external feature information is represented in face- and body-selective voxels within each search space. For all sizes in all three regions, a 2 (domain: face-selective, body-selective) x 2 (feature: internal, external) x 2 (arrangement: arranged, rearranged) repeated-measures ANOVA revealed a significant domain x arrangement interaction (all F 's > 8.05, all p 's < 0.05, all η_p^2 's > 0.29), indicating that the pattern of responses in face-selective voxels across the internal and external conditions (reflecting the parts-

or whole-based representations found in each region) is distinct from that found in body-selective voxels in and around face-selective cortex. In particular, body-selective voxels within the “FFA” search space (which also contained the FBA) did not represent the coherent arrangement of external (or internal) features (indeed, a separate analysis confirmed that the functionally defined FBA showed the same pattern of results; Fig. 6B), unlike face-selective voxels in the same search space, while body-selective voxels within the “OFA” and “pSTS” search spaces (each containing portions of the EBA) responded more to rearranged than arranged internal and external features (indeed, a separate analysis confirmed that the functionally defined EBA showed this same pattern of results; Fig. 6B), unlike face-selective voxels in the same search space. Further, we observed no significant domain x arrangement x feature interactions in any region at any size tested (all F 's < 1.80, all p 's > 0.19, all η_p^2 's < 0.08; with the exception of one bin in FFA, and one bin in pSTS, each of which showed a marginally significant effect: both F 's < 3.75, both p 's > 0.06, both η_p^2 's = 0.16), indicating that the differential sensitivity to arrangement information between face- and body-voxels did not differ for internal and external features. Taken together, these results reveal qualitatively different representation of external features in face-selective cortex, relative to adjacent or overlapping body-selective

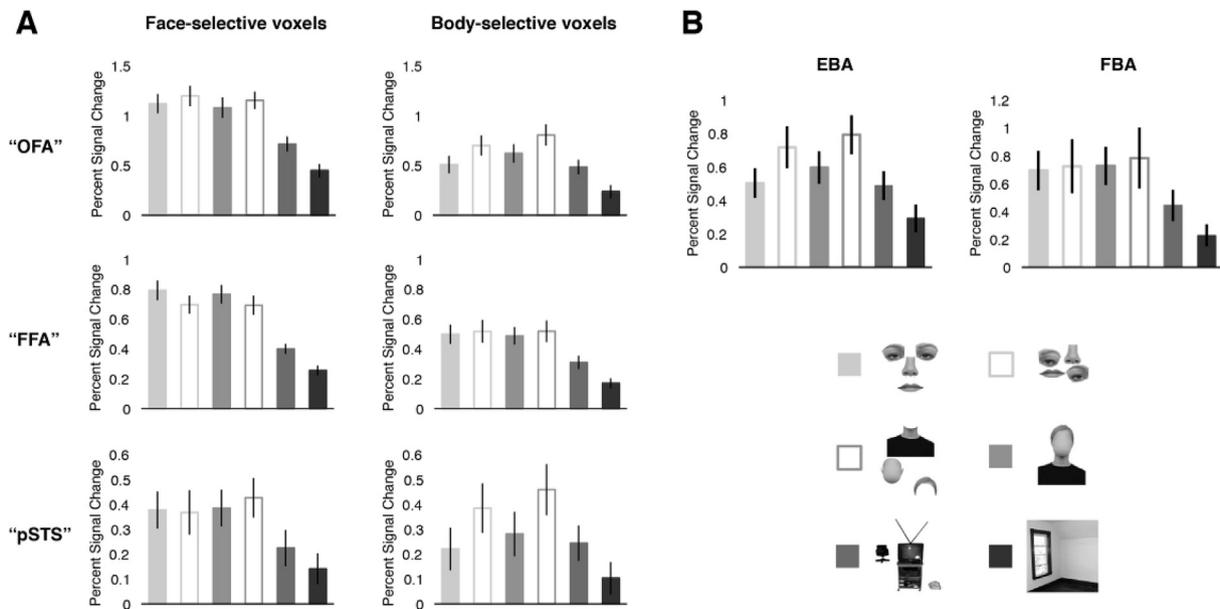


Fig. 6. A) VSF analysis of face- and body-selective voxels within anatomical search spaces for each region (i.e., "OFA," "FFA," and "pSTS"). For clarity, the figure depicts the pattern of responses found at a single ROI size only (i.e., the middle ROI size), rather than the entire VSF in each region; similar patterns were found at every other ROI size tested (see results). Voxels in each anatomical search space were ranked based on their response to faces > objects (revealing face-selective cortex, left column) or bodies > objects (revealing body-selective cortex, right column) in the independent Localizer dataset. Critically, in each search space, the pattern of responses in body-selective voxels significantly differed from that in the corresponding face-selective voxels, with stronger and qualitatively different representation of external features in face-selective voxels, relative to body-selective voxels. B) Average percent signal change to the six experimental conditions in the functionally defined EBA and FBA. The pattern of responses in EBA and FBA was similar to those found in body-selective voxels in the VSF analysis, and distinct from that in face-selective cortex. Error bars represent the standard error of the mean.

cortex, consistent with the hypothesis that external features are part of face representation.

4. Discussion

Despite the common intuition that a "face" includes internal features (e.g., eyes, nose, mouth) only, here we present strong neural evidence that a "face" is also composed of external features (e.g., hair, head outline, neck and shoulders). In particular, we found that regions of the cortical face processing system respond significantly more to external features than to both objects and scenes – rivaling the selectivity observed for internal features. We further explored the nature of external feature representation across the cortical face processing system, and found that face-selective regions show a similar division of labor across internal and external features, with OFA and pSTS extracting the parts of internal and external features, and FFA representing the coherent arrangement of internal and external features, further strengthening the claim that external features contribute to face representation, like internal features. These results help carve out a more precise understanding of exactly what features make up a "face", laying critical groundwork for theoretical accounts of face processing, which may benefit from incorporating external, as well as internal, features.

The present findings fit most strikingly with neuropsychological work on patient CK, who had impaired object recognition, but intact face recognition, allowing a clear view of the function of the face processing system without the 'contaminating' influence of the object processing system (Moscovitch et al., 1997; Moscovitch and Moscovitch, 2000). Counter to the hypothesis that face recognition depends on internal features only, CK was dramatically impaired at recognizing both internal and external features when these features were inverted, suggesting that CK's intact face processing system is highly sensitive to both kinds of features. Reinforcing these findings, here we directly tested external feature representation in the healthy cortical face processing system, revealing that external features are represented selectively within the face processing system. In so doing, our findings suggest that external

features, like internal features, are part of face representation.

Interestingly, however, individuals who suffer from damage to FFA (and are consequently prosopagnosic) often report greater reliance on external than internal features when recognizing faces (e.g., Rossion et al., 2003; Caldara et al., 2005). At first glance, this observation appears inconsistent with our finding that both internal and external features are processed by the same system, since damage to that system should impair recognition of both kinds of features. However, we propose that this observation does not represent a true dissociation of external versus internal feature information, but rather a dissociation between holistic face processing and feature-based face processing. For instance, it could be the case that in prosopagnosia, holistic face processing (as supported by FFA) is impaired, whereas feature-based face processing (as supported by OFA and/or pSTS) is nevertheless spared, allowing prosopagnosic individuals to utilize any salient face feature to recognize the face. Indeed, external features are often larger and more salient than internal features, and prosopagnosic individuals can even use particular internal features (e.g., the mouth) for face recognition (Caldara et al., 2005), suggesting that feature recognition is intact and that both kinds of feature (internal and external) can potentially be utilized.

Our results also shed light on developmental and adult psychophysical literatures investigating the role of external features in face recognition. These literatures suggest that humans use external features from infancy (Slater et al., 2000; Turati et al., 2006), and continue to do so across development, albeit with an increasing tendency to rely on internal over external features (Campbell and Tuck, 1995; Campbell et al., 1995; Want et al., 2003). Additional work in adults has shown a similar shift in the reliance on internal features as faces become more familiar (Ellis et al., 1979; O'Donnell and Bruce, 2001). While we tested only unfamiliar feature representation in adults, our findings are most consistent with the view that shifting biases for internal over external features across development and with increasing face familiarity do not reflect a total shift in the representation of what features comprise a face (i.e., since the adult face processing system is selective for unfamiliar internal and external features), but rather reflect a more subtle bias to use

particular features over others, while still employing both kinds of features.

Notably, while we found similar selectivity for internal and external features, other work found greater responses to internal than external features in face-selective cortex (Liu et al., 2009; Andrews et al., 2010; Betts and Wilson, 2010). Beyond the different methods used across these studies and ours (e.g., average activation versus fMRI adaptation in face-selective regions), the discrepancy across these findings most likely stems from the differing amounts of internal and external feature information presented in each study. For example, our internal feature stimuli included internal features presented on a white background, whereas other studies presented square cutouts of the internal features within the remainder of the face context (e.g., cheeks, forehead, etc.). Likewise, our external feature stimuli included the neck and shoulders, whereas other work included only the head outline and hair. In short then, the different ratios of internal versus external features across these studies would potentially lead to a greater bias for external features in our study relative to others. Importantly, under this account, it is still the case that both internal and external features are represented as “faces,” with the relative response to each depending on the amount of each kind of information presented. Indeed, the critical finding established in the present study is that both internal and external features are represented selectively within the face-processing system, providing evidence that a face is composed of both kinds of features.

But is it possible that regions in the cortical face processing system are not sensitive to external features per se, but rather respond to external features in isolation because these stimuli lead to participants to imagine internal features? Indeed, all three face-selective regions are known to respond during face imagery tasks, albeit more weakly than during face perception (Ishai et al., 2000; O’Craven and Kanwisher, 2000; Ishai et al., 2002; Cichy et al., 2012). There are three ways that such an imagery effect could have biased our results. First, it is possible that participants equally imagined internal features on the arranged and rearranged external features, which would lead to a similar response to both of these conditions in all regions. Counter to this prediction, FFA responded significantly more to the arranged than rearranged external features. Second, it is possible that participants imagined internal features more strongly on the arranged external features than the rearranged external features, which would lead to a greater response in all three regions to arranged external features (i.e., strong internal face imagery) than to rearranged external features (i.e., weak internal face imagery). Counter to this prediction, OFA and pSTS responded similarly to arranged and rearranged external features. Third, it is possible that the three face-selective regions show differential sensitivity to mental imagery. However, this account cannot explain the region by condition interaction observed here, and as such, we argue that such a possibility is not plausible. Thus, no potential mental imagery effect can readily explain the pattern of responses that we observed across face-selective regions to external features.

Related to the present study, one previous study found that FFA responds to “contextually defined” faces (i.e., images of blurred faces above “clear” torsos), and concluded that FFA not only represents intrinsic face features (i.e., the eyes, nose, and mouth), but also represents contextual information about faces (i.e., nonface information that commonly occurs with a face, such as a torso) (Cox et al., 2004). Thus, might FFA (along with OFA and pSTS) respond strongly to external features because these features serve as contextual cues for face processing, rather than because external features are intrinsically part of the face, as we have suggested? Although the context hypothesis is an interesting idea with much intuitive appeal, we do not think that it can explain our results for two reasons. First, the central evidence in favor of the context hypothesis is unreliable, with several subsequent studies failing to find evidence of context effects, and the original report failing to distinguish activity in FFA from that in the nearby and often overlapping FBA, leaving open the question of how much the findings can actually be attributed to FFA in particular (Peelen and Downing, 2005; Schwarzlose et al., 2005; Andrews et al., 2010; Axelrod and Yovel, 2011). Second, the theory itself

is underspecified, making few absolute predictions and leaving it almost impossible to disconfirm. For example, Cox and colleagues found a context effect for blurry faces *above* a body, but no context effect for blurry faces *below* a body. It is not clear precisely why the context hypothesis would predict this spatial arrangement effect; even if they had found a context effect for blurry faces below a body, the authors would have made the same claim about context representation in FFA. Furthermore, invoking the vague concept of “context” unnecessarily complicates our understanding of face processing when the same effects can already be explained by the intrinsic feature hypothesis (e.g., the intrinsic feature hypothesis readily predicts arrangement effects wherein the head goes above the neck and shoulders, just like the eyes go above the nose and mouth). For these reasons, we argue that the most parsimonious account of our findings is to simply treat external features as intrinsically part of the face (i.e., not as a more abstract or indirect contextual cue), as is widely accepted for internal features. Of course, while we found only similarities in internal and external feature representations, suggesting that both features are intrinsically part of the face, it could always be the case that another study might reveal qualitative differences in internal and external feature representation, potentially consistent with the hypothesis that responses to external features instead reflect face context representation.

Finally, the finding that external features—potentially even including the neck and shoulders—are part of the face raises the intriguing question of where and how the line gets drawn between a “face” and a “body.” Indeed, the findings that face-selective cortex i) is more selective for external features than nearby body-selective cortex, and ii) represents distinct kinds of information about external features, relative to body-selective cortex, indicate that external features are processed within face-selective cortex, but leave open questions about whether these features are processed as a body in body-selective cortex. Along these lines, one intriguing study found that both FFA and FBA show reduced neural competition for a head above a torso (i.e., the entire body extending from the shoulders to the waist) relative to a torso above a head, which was interpreted as evidence of an integrated representation of the face and body in both regions (Bernstein et al., 2014), potentially blurring the boundary between the “face” and the “body”. However, our findings suggest an alternative interpretation, wherein FFA represents the arrangement of the external features, including the neck and shoulders, leading to the arrangement effect observed by Bernstein and colleagues, while FBA represents the arrangement of the head on the body, but only when a sufficient amount of the whole body is shown (explaining why we did not observe an arrangement effect, since our external feature stimuli did not extend beyond the shoulders, and therefore may not have been processed as a body per se). This alternative account therefore raises the possibility that a “face” includes not only the front part of the head down to the chin, but rather extends to the neck and shoulders as well, whereas a “body” is only processed as such when a sufficient amount of the whole body is in view (e.g., when the entire torso is shown, not just the head and shoulders) (Taylor et al., 2007). Importantly, however, our study was not designed to address this question directly, and as such, future work will be required to establish precisely where and how the line gets drawn between face and body processing in the brain.

Taken together, the present results provide strong neural evidence that external features, along with internal features, make up a face. Understanding what features compose a face is important for developing accurate theoretical accounts of a variety of face recognition processes, from recognizing identities to emotions. Moreover, incorporating the analysis of external features in addition to internal features might facilitate the ability of computer vision systems reach human performance and/or improve their real-world applicability (Izzat and Jarudi, 2003). Finally, our results raise an important set of new questions. For example, how do these features interact to ultimately make a “face”? Analogously, how are these features coded at the neuronal level (e.g., Freiwald et al., 2009; Chang and Tsao, 2017)? In any case, the present finding that external features, like internal features, compose a face, highlights the

need to consider the role of external features in face recognition.

Acknowledgments

We would like to thank the Facility for Education and Research in Neuroscience (FERN) Imaging Center in the Department of Psychology, Emory University, Atlanta, GA. We would also like to thank Andrew Persichetti, Yaseen Jamal, Annie Cheng, and Bree Beal for insightful comments. The work was supported by Emory College, Emory University (DD), National Eye Institute grant T32EY7092 (FK), and an independent undergraduate research award from Emory College (EM). The authors declare no competing financial interests.

References

- Abudarham, N., Yovel, G., 2016. Reverse engineering the face space: discovering the critical features for face identification. *J. Vis.* 16.
- Andrews, T.J., Davies-Thompson, J., Kingstone, A., Young, A.W., 2010. Internal and external features of the face are represented holistically in face-selective regions of visual cortex. *J. Neurosci.* 30, 3544–3552.
- Arcurio, L.R., Gold, J.M., James, T.W., 2012. The response of face-selective cortex with single face parts and part combinations. *Neuropsychologia* 50, 2454–2459.
- Axelrod, V., Yovel, G., 2010. External facial features modify the representation of internal facial features in the fusiform face area. *Neuroimage* 52, 720–725.
- Axelrod, V., Yovel, G., 2011. Nonpreferred stimuli modify the representation of faces in the fusiform face area. *J. Cognit. Neurosci.* 23, 746–756.
- Bernstein, M., Oron, J., Sadeh, B., Yovel, G., 2014. An integrated face–body representation in the fusiform gyrus but not the lateral occipital cortex. *J. Cognit. Neurosci.* 26, 2469–2478.
- Betts, L.R., Wilson, H.R., 2010. Heterogeneous structure in face-selective human occipito-temporal cortex. *J. Cognit. Neurosci.* 22, 2276–2288.
- Brandman, T., Yovel, G., 2012. A face inversion effect without a face. *Cognition* 125, 365–372.
- Brunelli, R., Poggio, T., 1993. Face recognition - features versus templates. *Ieee T Pattern Anal.* 15, 1042–1052.
- Caldara, R., Schyns, P., Mayer, E., Smith, M.L., Gosselin, F., Rossion, B., 2005. Does prosopagnosia take the eyes out of face representations? Evidence for a defect in representing diagnostic facial information following brain damage. *J. Cognit. Neurosci.* 17, 1652–1666.
- Campbell, R., Tuck, M., 1995. Recognition of parts of famous-face photographs by children: an experimental note. *Perception* 24, 451–456.
- Campbell, R., Walker, J., Baroncohen, S., 1995. The development of differential use of inner and outer face features in familiar face identification. *J. Exp. Child Psychol.* 59, 196–210.
- Chang, L., Tsao, D.Y., 2017. The code for facial identity in the primate brain. *Cell* 169, 1013–1028 e1014.
- Cichy, R.M., Heinze, J., Haynes, J.D., 2012. Imagery and perception share cortical representations of content and location. *Cerebr. Cortex* 22, 372–380.
- Cox, D., Meyers, E., Sinha, P., 2004. Contextually evoked object-specific responses in human visual cortex. *Science* 304, 115–117.
- Cox, R.W., Jesmanowicz, A., 1999. Real-time 3D image registration for functional MRI. *Magn. Reson. Med.* 42, 1014–1018.
- de Haas, B., Schwarzkopf, D.S., Alvarez, I., Lawson, R.P., Henriksson, L., Kriegeskorte, N., Rees, G., 2016. Perception and processing of faces in the human brain is tuned to typical feature locations. *J. Neurosci.* 36, 9289–9302.
- Deen, B., Richardson, H., Dilks, D.D., Takahashi, A., Keil, B., Wald, L.L., Kanwisher, N., Saxe, R., 2017. Organization of high-level visual cortex in human infants. *Nat. Commun.* 8, 13995.
- Desikan, R.S., Segonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T., Albert, M.S., Killiany, R.J., 2006. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage* 31, 968–980.
- Downing, P.E., Jiang, Y., Shuman, M., Kanwisher, N., 2001. A cortical area selective for visual processing of the human body. *Science* 293, 2470–2473.
- Downing, P.E., Chan, A.W., Peelen, M.V., Dodds, C.M., Kanwisher, N., 2006. Domain specificity in visual cortex. *Cerebr. Cortex* 16, 1453–1461.
- Eimer, M., 2000. The face-specific N170 component reflects late stages in the structural encoding of faces. *Neuroreport* 11, 2319–2324.
- Ellis, H.D., Shepherd, J.W., Davies, G.M., 1979. Identification of familiar and unfamiliar faces from internal and external features - some implications for theories of face recognition. *Perception* 8, 431–439.
- Epstein, R.A., Kanwisher, N., 1998. A cortical representation of the local visual environment. *Nature* 392, 598–601.
- Farah, M.J., Wilson, K.D., Drain, M., Tanaka, J.N., 1998. What is "special" about face perception? *Psychol. Rev.* 105, 482–498.
- Freiwald, W.A., Tsao, D.Y., Livingstone, M.S., 2009. A face feature space in the macaque temporal lobe. *Nat. Neurosci.* 12, 1187–1196.
- Grill-Spector, K., 2003. The neural basis of object perception. *Curr. Opin. Neurobiol.* 13, 159–166.
- Grill-Spector, K., Kushnir, T., Hendler, T., Edelman, S., Itzhak, Y., Malach, R., 1998. A sequence of object-processing stages revealed by fMRI in the human occipital lobe. *Hum. Brain Mapp.* 6, 316–328.
- Harris, A., Aguirre, G.K., 2008. The representation of parts and wholes in face-selective cortex. *J. Cognit. Neurosci.* 20, 863–878.
- Harris, A., Aguirre, G.K., 2010. Neural tuning for face wholes and parts in human fusiform gyrus revealed by fMRI adaptation. *J. Neurophysiol.* 104, 336–345.
- Hu, Y., Jackson, K., Yates, A., White, D., Phillips, P.J., O'Toole, A.J., 2017. Person recognition: qualitative differences in how forensic face examiners and untrained people rely on the face versus the body for identification. *Vis. Cognit.* 25, 492–506.
- Ishai, A., Ungerleider, L.G., Haxby, J.V., 2000. Distributed neural systems for the generation of visual images. *Neuron* 28, 979–990.
- Ishai, A., Haxby, J.V., Ungerleider, L.G., 2002. Visual imagery of famous faces: effects of memory and attention revealed by fMRI. *Neuroimage* 17, 1729–1741.
- Izzat, N., Jarudi, P.S., 2003. Relative Contributions of Internal and External Features to Face Recognition. Massachusetts Institute of Technology.
- James, T.W., Arcurio, L.R., Gold, J.M., 2013. Inversion effects in face-selective cortex with combinations of face parts. *J. Cognit. Neurosci.* 25, 455–464.
- Julian, J.B., Fedorenko, E., Webster, J., Kanwisher, N., 2012. An algorithmic method for functionally defining regions of interest in the ventral visual pathway. *Neuroimage* 60, 2357–2364.
- Kamps, F.S., Julian, J.B., Kubilius, J., Kanwisher, N., Dilks, D.D., 2016. The occipital place area represents the local elements of scenes. *Neuroimage* 132, 417–424.
- Kanwisher, N., McDermott, J., Chun, M.M., 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J. Neurosci.* 17, 4302–4311.
- Kourtzi, Z., Kanwisher, N., 2001. Representation of perceived object shape by the human lateral occipital complex. *Science* 293, 1506–1509.
- Lai, J., Pancaroglu, R., Oruc, I., Barton, J.J.S., Davies-Thompson, J., 2014. Neuroanatomic correlates of the feature-salience hierarchy in face processing: an fMRI-adaptation study. *Neuropsychologia* 53, 274–283.
- Langner, O., Dotsch, R., Bijlstra, G., Wigboldus, D.H.J., Hawk, S.T., van Knippenberg, A., 2010. Presentation and validation of the Radboud faces Database. *Cognit. Emot.* 24, 1377–1388.
- Liu, J., Harris, A., Kanwisher, N., 2009. Perception of face parts and face configurations: an fMRI study. *J. Cognit. Neurosci.* 22, 203–211.
- Maurer, D., O'Craven, K.M., Le Grand, R., Mondloch, C.J., Springer, M.V., Lewis, T.L., Grady, C.L., 2007. Neural correlates of processing facial identity based on features versus their spacing. *Neuropsychologia* 45, 1438–1451.
- Moscovitch, M., Moscovitch, D.A., 2000. Super face-inversion effects for isolated internal or external features, and for fractured faces. *Cogn. Neuropsychol.* 17, 201–219.
- Moscovitch, M., Winocur, G., Behrmann, M., 1997. What is special about face recognition? Nineteen experiments on a person with visual object agnosia and dyslexia but normal face recognition. *J. Cognit. Neurosci.* 9, 555–604.
- Mumford, J., 2007. A Guide to Calculating Percent Change with Featquery (Unpublished Tech Report).
- Nestor, A., Plaut, D.C., Behrmann, M., 2016. Feature-based face representations and image reconstruction from behavioral and neural data. *P. Natl. Acad. Sci. USA* 113, 416–421.
- Norman-Haignere, S.V., Albouy, P., Caclin, A., McDermott, J.H., Kanwisher, N.G., Tillmann, B., 2016. Pitch-responsive cortical regions in congenital amusia. *J. Neurosci.* 36, 2986–2994.
- O'Craven, K.M., Kanwisher, N., 2000. Mental imagery of faces and places activates corresponding stimulus-specific brain regions. *J. Cognit. Neurosci.* 12, 1013–1023.
- O'Donnell, C., Bruce, V., 2001. Familiarisation with faces selectively enhances sensitivity to changes made to the eyes. *Perception* 30, 755–764.
- Peelen, M.V., Downing, P.E., 2005. Selectivity for the human body in the fusiform gyrus. *J. Neurophysiol.* 93, 603–608.
- Persichetti, A.S., Dilks, D.D., 2016. Perceived egocentric distance sensitivity and invariance across scene-selective cortex. In: *A Journal devoted to the Study of the Nervous System and Behavior*, 77, pp. 155–163.
- Pitcher, D., Walsh, V., Yovel, G., Duchaine, B., 2007. TMS evidence for the involvement of the right occipital face area in early face processing. *Curr. Biol.* 17, 1568–1573.
- Pitcher, D., Dilks, D.D., Saxe, R.R., Triantafyllou, C., Kanwisher, N., 2011. Differential selectivity for dynamic versus static information in face-selective cortical regions. *Neuroimage* 56, 2356–2363.
- Rice, A., Phillips, P.J., O'Toole, A., 2013. The role of the face and body in unfamiliar person identification. *Appl. Cognit. Psychol.* 27, 761–768.
- Rossion, B., Caldara, R., Seghier, M., Schuller, A.M., Lazeyras, F., Mayer, E., 2003. A network of occipito-temporal face-sensitive areas besides the right middle fusiform gyrus is necessary for normal face processing. *Brain* 126, 2381–2395.
- Saygin, Z.M., Osher, D.E., Norton, E.S., Youssoufian, D.A., Beach, S.D., Feather, J., Gaab, N., Gabrieli, J.D.E., Kanwisher, N., 2016. Connectivity precedes function in the development of the visual word form area. *Nat. Neurosci.* 19, 1250–1255.
- Schiltz, C., Rossion, B., 2006. Faces are represented holistically in the human occipito-temporal cortex. *Neuroimage* 32, 1385–1394.
- Schiltz, C., Dricot, L., Goebel, R., Rossion, B., 2010. Holistic perception of individual faces in the right middle fusiform gyrus as evidenced by the composite face illusion. *J. Vis.* 10.
- Schwarzlose, R.F., Baker, C.I., Kanwisher, N., 2005. Separate face and body selectivity on the fusiform gyrus. *J. Neurosci.* 25, 11055–11059.
- Sinha, P., Poggio, T., 1996. I think I know that face. *Nature* 384, 404–404.
- Slater, A., Bremner, G., Johnson, S.P., Sherwood, P., Hayes, R., Brown, E., 2000. Newborn infants' preference for attractive faces: the role of internal and external facial features. *Infancy* 1, 265–274.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., Niazy, R.K., Saunders, J., Vickers, J., Zhang, Y.Y., De Stefano, N., Brady, J.M., Matthews, P.M.,

2004. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* 23, S208–S219.
- Tanaka, J.W., Sengco, J.A., 1997. Features and their configuration in face recognition. *Mem. Cognit.* 25, 583–592.
- Taylor, J.C., Wiggett, A.J., Downing, P.E., 2007. Functional MRI analysis of body and body part representations in the extrastriate and fusiform body areas. *J. Neurophysiol.* 98, 1626–1633.
- Thompson, P., 1980. Thatcher, Margaret - a new illusion. *Perception* 9, 483–484.
- Tong, F., Nakayama, K., Moscovitch, M., Weinrib, O., Kanwisher, N., 2000. Response properties of the human fusiform face area. *Cogn. Neuropsychol.* 17, 257–279.
- Turati, C., Cassia, V.M., Simion, F., Leo, I., 2006. Newborns' face recognition: role of inner and outer facial features. *Child Dev.* 77, 297–311.
- Want, S.C., Pascalis, O., Coleman, M., Blades, M., 2003. Recognizing people from the inner or outer parts of their faces: developmental data concerning 'unfamiliar' faces. *Br. J. Dev. Psychol.* 21, 125–135.
- Young, A.W., Hellawell, D., Hay, D.C., 1987. Configurational information in face perception. *Perception* 16, 747–759.
- Yovel, G., Kanwisher, N., 2004. Face perception: domain specific, not process specific. *Neuron* 44, 889–898.
- Zhang, J., Li, X., Song, Y., Liu, J., 2012. The fusiform face area is engaged in holistic, not parts-based, representation of faces. *PLoS One* 7.
- Zhao, M.T., Cheung, S.H., Wong, A.C.N., Rhodes, G., Chan, E.K.S., Chan, W.W.L., Hayward, W.G., 2014. Processing of configural and componential information in face-selective cortical areas. *Cognit. Neurosci.* 5, 160–167.