

FFA: a flexible fusiform area for subordinate-level visual processing automatized by expertise

Michael J. Tarr and Isabel Gauthier

Much evidence suggests that the fusiform face area is involved in face processing. In contrast to the accompanying article by Kanwisher, we conclude that the apparent face selectivity of this area reflects a more generalized form of processing not intrinsically specific to faces.

How does the primate visual system process and identify objects? Central to this question is that visual recognition can occur at many levels, from coarse categories, for example “a car,” to highly specific individuals, “my dad’s 1960 TR3 convertible” (subordinate-level processing). The processing demands exerted by these two tasks seem to require separable visual recognition systems, one dedicated to determining category membership and one dedicated to individuation. The most conspicuous version of this argument is one in which there is a functional and neuroanatomical division between the mechanisms supporting object at the basic level and face recognition at the individual level. Here we evaluate this claim and ask whether putatively face-specific mechanisms in a region of the visual cortex known as the Fusiform Face Area (“FFA”) are selective specifically to the geometry of faces or whether, through the operation of other factors, they extend to non-face stimuli.

The comparison between faces and non-face objects is critical for evaluating three competing models. First, Kanwisher (this issue) suggests that face selectivity in the FFA as revealed by functional magnetic resonance imaging (fMRI) reflects a special-purpose, innate mechanism for detecting the geometry of faces, that is, a mechanism that determines only that a face is a face, but not the individual identity of the face^{1–3}. Evidence for this ‘detection model’ comes from studies in which selectivity for faces in FFA has been established over a wide vari-

ety of conditions, including line drawings, two-tone, gray-scale or colored images—stimulus manipulations that affect face-recognition performance but have little impact on activation in FFA^{1,4,5}.

For example, face inversion seems to only minimally influence activity in the FFA^{1,6,7}. This result appears at odds with the strong effect of inversion on face recognition performance⁸. However, inversion only impairs, but does not abolish, the perception of identity. Inverted faces are typically recognized well above chance level, which is consistent with the small but significant fMRI face inversion effects reported in some studies^{1–9}. Moreover, this same small inversion effect has been replicated with non-face objects in expert subjects⁹. The absence or relatively small inversion effect in the FFA may also mean that activity as measured by fMRI simply is not sufficiently fine-grained to pick up certain effects. As a case in point, fMRI cannot detect the robust and early effect of face inversion reported in ERPs: a ten-millisecond delay in the N170 face-sensitive potential^{10,11} that is recorded at temporo-occipital electrodes^{12,13} or a similar inversion effect in intracranial recordings¹⁴.

A second model holds that category selectivity reflects the presence of a geometrically defined feature map in ventral temporal cortex^{15,16}. The hypothesis is that small, localized brain regions are activated by objects, generally from a single category, with similar shape and image features. Evidence for feature maps comes from monkey neurophysiology¹⁷ suggesting a topography of features in inferior temporal cortex (IT) and from human fMRI studies indicating that across a single task, different stimuli selectively activate different regions of the ventral temporal cortex. For instance, distinct brain areas respond preferentially to letter strings, chairs, buildings or faces^{4,6,16,18}. However, geometric

similarity alone seems unlikely to account for category selectivity, in that task manipulations independent of image geometry can lead to activation maps indistinguishable from those produced by manipulating geometry.

Third, our preferred model, the flexible process map¹⁹, holds that an assortment of extrastriate areas support separable components of visual object processing and recognition. This model is distinguished from the detection and feature-map models in that specialization of large extrastriate areas (that is, those detectable with fMRI) is not *a priori* a function of stimulus geometry. Rather, through experience, observers associate particular geometries—which define object categories—with task-appropriate recognition strategies that automatically recruit components of the process map. In contrast to other accounts, the process-map model provides a general explanation for category selectivity, as well as why the FFA responds to a wide range of stimuli, including faces and non-face objects.

Factors affecting FFA processing

Given that face selectivity may not be intrinsically related to the specific geometry of faces, what else might account for the strong activation for faces in the FFA? Behavioral work^{20–22} suggests two critical factors associated with face recognition because of our particular experience with faces. First, the level of categorization at which objects are recognized is different—faces generally being recognized at the individual level, but objects such as chairs or birds being recognized at a more categorical level. Second, we have far greater expertise with faces as compared to almost any other visual category. Expertise with faces may automatize face processing at the individual level, thereby rendering task manipulations less effective for faces, for example, the rapid presentation rates used in some stud-

Michael Tarr is at the Department of Cognitive and Linguistic Sciences, Brown University, Box 1978, Providence, Rhode Island 02912, USA. Isabel Gauthier is at the Department of Psychology, Vanderbilt University, 301 Wilson Hall, Nashville, Tennessee 37240, USA. e-mail: isabel.gauthier@vanderbilt.edu or Michael_Tarr@brown.edu

commentary

ies (such as 300 ms with a 500 ms ISI⁵) may not affect face processing nearly as much as object processing.

Because extreme values for these two dimensions are typically associated with face recognition, any comparison between faces and non-face objects confounds these factors unless they are independently manipulated. When we do such manipulations, it is the level of categorization and the level of expertise, not the stimulus geometry, that are critical in determining selectivity in the FFA. Given such domain-general processes, we argue that the FFA may be better thought of as a 'flexible fusiform area'—a mechanism for visual recognition that exhibits a great deal of plasticity in response to both task demands and experience.

Evidence from fMRI

Functional MRI provides an ideal tool for evaluating these models because brain areas can be identified in a functional manner. In comparison to cytoarchitectonic criteria (like Brodmann areas) or gross anatomical landmarks (such as the Talairach system), a region performing the same function in different individuals can be defined, allowing for variability in the mapping of function to anatomy. Building on a long history of neuropsychological findings indicating specialization in the primate visual system for face processing²⁰, both PET and fMRI studies have localized face-selective regions in the ventral temporal lobe^{4,18,23,24}. Functional MRI has been used to identify a small number of face-selective voxels—the FFA—through a comparison of faces and non-face objects in a common task⁴. Importantly, this neural pattern replicates within the same study and generalizes across experimental conditions. That is, candidate face-selective voxels are recruited more for faces than for non-face objects even when the task and stimuli are varied. These results from imaging seem to support earlier accounts of face selectivity in the brain (and provide a method of localizing it). Different interpretations are possible, however, when categorization level and experience are taken into account.

To specifically address the roles of these two factors, several studies have manipulated the level of categorization and the level of expertise for non-face objects. These studies find that the response of FFA to non-face objects is affected by these manipulations independent of particular stimulus geometries^{25–27}. Logically then, these other factors can account for most of the strong face selectivity in FFA activity.

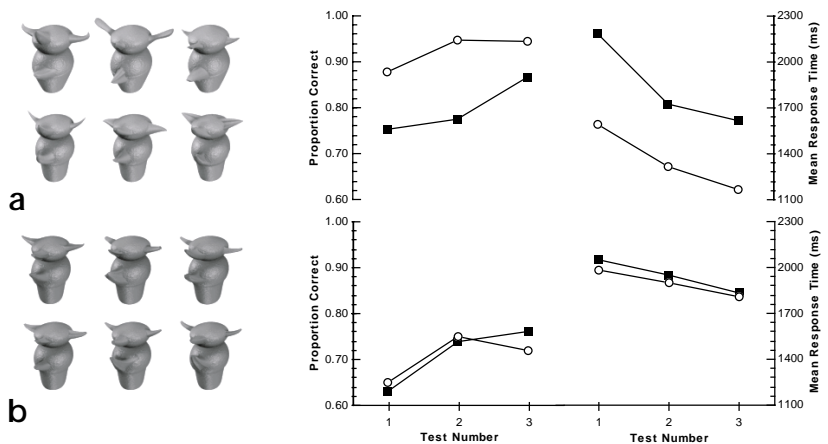


Fig. 1. Examples of Greebles. (a) Greebles from a set that Greeble experts could learn to recognize faster than novices. (b) Another set of Greebles, which the same experts could not learn faster than novices, presumably because they are more visually homogeneous than Greebles in the training set. Filled squares denote data from novices, and open circles denote data from experts⁴³.

Manipulating categorization-level

The FFA bilaterally shows an increased response when familiar artificial and natural kind objects are seen in the context of visually similar stimuli²⁵. The FFA shows a similar response for animals, even when their faces are obscured²⁸ (but see ref. 5). Moreover, when the level of categorization is manipulated by different labels that subjects match to the same image (for example, 'TR3' versus 'car' for the same picture of a TR3), the FFA is more active for judgments requiring classification at a subordinate level compared to more categorical judgments for a large variety of objects, living or artificial^{26,27}. Thus, living things, including animals and animals with hidden faces¹⁵, may recruit the FFA because of their visual homogeneity²⁹ and the level of visual discrimination that is most often applied to them (such as distinguishing a cow from a horse by default).

Manipulating level of expertise

The second factor influencing FFA activity for a wide range of objects is subject expertise. The FFA and a similarly selective occipital area ('OFA') in the right hemisphere^{4,7,30,31} are both recruited when observers become experts in discriminating objects from a visually homogeneous category (objects that share a common geometric structure). As detailed below, this occurs in subjects trained in the laboratory to be experts with novel objects called 'Greebles,' as well as bird and car experts with 20 years of experience^{9,25}. Expertise here means more than practice or priming with specific instances of a class because the FFA exhibits an equivalent response for new

Greebles never shown during training. Such expertise effects indicate a high degree of flexibility with regard to acceptable image geometries in the neural network mediating object recognition. Moreover, this flexibility does not have a critical period and can be harnessed for relatively novel learning experiences (such as learning to read brain scans or individuate nonsense objects such as Greebles).

Automatization by expertise

Object processing in the FFA can be automatized by expertise. Consider that Greeble experts show FFA activation in response to Greebles even under passive viewing instructions⁹. Moreover, whereas novices can easily focus on the top half of a composite Greeble (made out of the top of one and the bottom of another), Greeble experts cannot ignore the bottom distractor part. The behavioral changes measured by this 'composite effect' show a significant correlation across subjects and training sessions with activity in the right FFA of Greeble experts (I.G. and M.J.T., unpublished data). Thus, neural changes with expertise, at least in the right FFA, seem to be tied strongly to holistic processing strategies. This is consistent with a PET study³² reporting that the right FFA is more active when observers attend to whole faces than face parts, whereas the opposite pattern is found in the left FFA.

An fMRI study of bird and car experts further supports the idea that expertise automatizes processing in the right FFA²⁵. In novices, the homogeneous classes of birds and cars elicit more activity in the FFA than a set of various familiar objects, and this

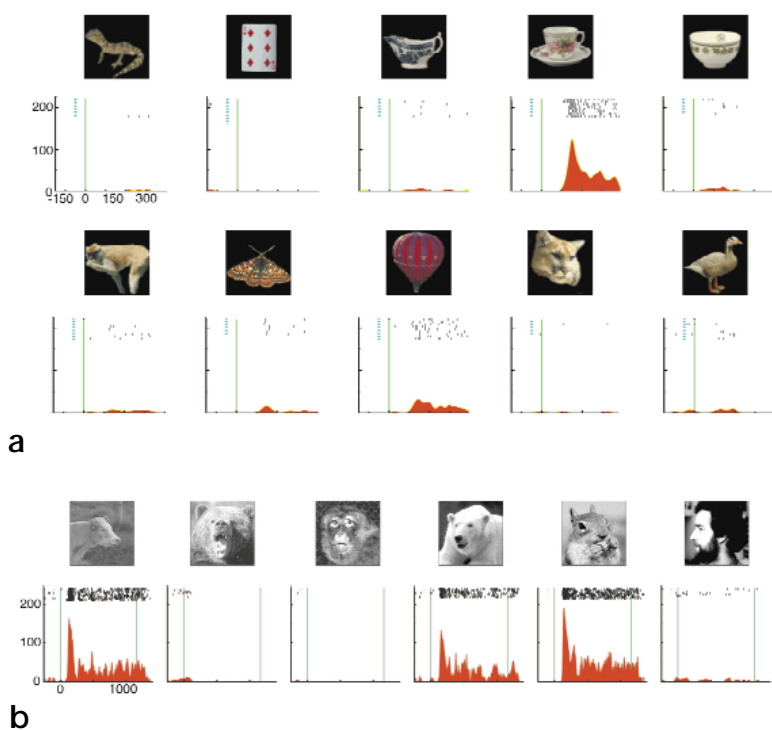


Fig. 2. Complex selectivity of IT cells. Adapted from ref. 49. **(a)** This cell responded most to the picture of a teacup, and also responded more to the picture of a balloon than all other images tested. Although the selectivity of this cell may be based on color, this set of responses illustrates that any hypothesis about the selectivity of a neuron is highly dependent upon the stimulus set used in an experiment. **(b)** A cell that responds to a subset of animals, with no obvious measure of similarity separating the effective from ineffective stimuli. The authors⁴⁹ argue that selectivity emerges from experience and that images that activate IT neurons need not be especially similar to one another.

effect is more pronounced when subjects attend to the identity of the objects than when they attend to their location. In contrast, experts show more activity in the right FFA for their category of expertise, birds or cars, regardless of whether they attend to the identity or location of the objects. In addition, a behavioral test of relative performance in matching birds or cars is highly correlated ($r \sim 0.8$) with activity in the right FFA for another set of birds or cars during the location judgments. Similarly, an ERP study of bird and dog experts found an enhanced N170—the negative potential typically associated with face-specific processing^{10,11}—for categorizing objects within their respective domains of expertise³³. Thus, both fMRI²⁵ and ERP³³ studies of experts provide evidence that the neural basis of face recognition extends to other domains of expertise.

Although these findings would already seem sufficient to reject the hypothesis of the FFA as a face detector^{1,2}, additional evidence is provided by fMRI habituation procedures assessing the nature of automatized processing for faces in the FFA. A task in

which subjects attended to the location of a face revealed that FFA and OFA activity habituates more to the repeated presentation of a single face than to different faces³¹. Thus, this activity seems to reflect more than generic face detection. Also supporting the hypothesis that the FFA is involved in individuation of faces, attention to face identity (rather than to eye gaze) increases activity in the vicinity of the FFA³⁴. It does not follow, however, that this area is not involved in individual face processing when subjects attend to dimensions other than identity. Indeed, several studies find that the FFA activity level in response to faces, as well as to other categories of expertise, is relatively immune to whether the task is recognition or not^{4,16,25}. Thus, it seems that the FFA automatically processes objects for which observers are experts at the subordinate level, but, as with most visual processes, the efficacy of the system may be further modulated by attention.

Evidence from neuropsychology

The most salient neuropsychological example of impaired visual object recognition

with brain injury is that of prosopagnosia. Patients with this disorder usually have suffered injury in the lingual and fusiform gyri, which typically encompass the FFA³⁵. The pattern of sparing and loss is striking—although they can recognize a face as a face, these patients are remarkably impaired at face recognition at the individual level, but sometimes seem to be able to recognize non-face objects. Although this deficit is often cited as evidence for a face-specific neural substrate, such as the FFA, it actually says little about the origins of face selectivity. Thus, prosopagnosia as a face-specific deficit does not distinguish among the detection, feature-map and process-map models—all three of which acknowledge the existence of and, indeed, attempt to explain category-selective neural substrates. To the extent that prosopagnosia speaks to these models at all, it is because of an intact ability. Even though the FFA is typically lesioned in prosopagnosia, face detection is spared when tested^{35,36}; thus, the detection model cannot account for face-selective activity in the FFA. Interestingly, autistic individuals do not show normal FFA activity when viewing faces and also have few problems with face detection³⁷.

At the same time, a patient (CK) with the complementary pattern of sparing and loss—preserved face recognition with impaired object recognition—is widely held as the other side of a double dissociation demonstrating distinct face and object recognition mechanisms³⁸. As with prosopagnosia, at issue is not whether there exist category-selective brain regions, but rather, whether these regions are domain-general or domain-specific. That is, all three models predict that a patient such as CK could occur; the question is simply why. Unfortunately, testing the domain-specificity of CK's spared recognition abilities would be difficult in that it is his holistic processing mechanisms that remain intact (given that it is generally agreed that the FFA is involved in holistic processing for faces and possibly other objects^{9,25,27,31}). Within the framework of the process-map model, the recruitment of these holistic mechanisms results from the interaction between subordinate-level recognition and expertise training. CK, however, seems to have lost the ability to recognize objects in a non-holistic manner and, in particular, the ability to move from a novice to an expert for any new object class. Thus, unlike normal subjects^{9,39}, CK cannot be bootstrapped into applying the intact components of his visual recognition system to new, non-face objects.

commentary

Beyond its specific inconsistency with the detection model, the syndrome of prosopagnosia reinforces the point that processing in the FFA is domain-general. In particular, patients with prosopagnosia are impaired with more than faces. Although this was raised as a possibility some time ago²⁰, it was rejected by some researchers⁴⁰ based on a study in which a prosopagnosic patient was differentially impaired at recognizing previously studied faces as compared to recognizing previously studied homogeneous non-face objects (such as eyeglasses). This experiment, however, did not consider potential response biases, for instance, the frequency with which the patient responded “old” to individual items actually shown during study relative to those items not shown during study, and the amount of time and effort expended by the patient for faces as compared to objects. Indeed, if the patient believed (perhaps incorrectly) that he was more impaired with faces relative to other objects, he may spend less time and allocate fewer attentional resources in encoding and recalling faces³⁶. These concerns were addressed in a study that measured recognition response times in conjunction with a bias-free (d') measure of recognition accuracy³⁶. Data from two patients with prosopagnosia were collected for their recognition of faces, visually similar non-face objects and Greebles. Although the patients did about as well as control subjects in recognizing objects and Greebles, they showed a speed-accuracy tradeoff to achieve this level of performance. That is, the patients spent far more time than the controls to achieve the same level of recognition accuracy. Moreover, when patients were forced to spend the same amount of time for each stimulus category, accuracy as measured by d' —a statistic capturing a subject's ability to discriminate or remember items independent of response bias—revealed that patients with prosopagnosia are most impaired whenever visually similar stimuli must be discriminated, regardless of category. Thus, prosopagnosia may actually be a domain-general deficit when assessed using a complete picture of performance.

Critiques of the process map

It is our contention that behavioral, neuroimaging, and neuropsychological results provide strong evidence in support of a domain-general model—the process-map model—for face processing and recognition. In particular, manipulations of the factors of level of categorization and level of expertise are able to account for nearly all of the apparent differences between face and object recognition. There are, however, skeptics.

Next we review and address some critiques of a domain-general account.

Several groups have attempted to refute the hypothesis that the reason that FFA seems face selective is because faces are recognized at the individual level, whereas other objects are recognized at a coarser level. One study compared passive viewing of faces and visually similar flowers appearing on a continuously changing background of nonsense patterns or non-face objects. When compared to a baseline of non-face objects, faces, but not flowers, produced activation in right fusiform gyrus⁴¹. Similarly, another study compared identity judgments for faces and hands and found higher activation for faces than hands as compared to a fixation baseline⁴.

How can these results be reconciled with other studies indicating a role for subordinate-level processing in the FFA? First, consider that the first study used passive viewing and did not require subjects to process flowers at the subordinate level. The second study showed faces and hands for 500 ms each. As reviewed above, subordinate-level judgments not automatized by expertise typically require additional processing time⁴². Second, these experiments only indicate that categorization level is not sufficient to account for all the specificity in the FFA. Note that this is already assumed by our model, in which the interaction of multiple factors (such as homogeneity, categorization level and expertise) constrain face-recognition performance.

In sum, experiments using non-face objects can provide some evidence that a given factor influences FFA response. However, the exact contribution of a factor to the combination that leads to the maximum FFA response cannot be inferred based on the same experiment (unless it can be precisely equated for faces and objects and we know how it interacts with other factors). This also means that we cannot refute the possibility that image geometry is one of several factors that contribute to maximum selectivity—at least not until maximum selectivity is obtained with non-face objects.

Even accepting that the FFA is involved in subordinate-level recognition, it remains possible that this area is preferentially selective for a specific geometry—that defining what sort of an image ‘counts’ as a face. Refuting this argument is complex—although we know that the FFA can become selective for other domains of expertise, for example, Greebles, birds and cars^{9,25}, it is often stated that “Greebles look like faces” and that experts can learn to apply face-selective mechanisms to other categories. These two claims seem reasonable—a

mechanism efficient for face processing might gain some of its efficiency by making assumptions regarding stimulus geometry. However, it is crucial that we be clear about what such claims imply, given that we have already established selectivity of the FFA in response to non-face stimuli.

As a specific and oft cited example, consider FFA selectivity for Greebles. Several findings indicate that Greebles are not treated as face-like in terms of their geometry. First, Greeble novices do not show Greeble-specific activation in FFA, but Greeble experts do show such activation⁹. Second, Greeble novices do not show face-like behavioral effects, but Greeble experts do show such effects^{39,43}. Third, CK, the agnosic patient who shows intact face recognition, but is dramatically impaired in non-face object recognition³⁸ cannot recognize Greebles any better than other objects—even when he is told to think of Greebles as faces or little people (M. Behrmann, I. G. & M. J. T., unpublished observation). At most, subjects can only learn to treat Greebles ‘like faces’ following an intensive training process that takes about 10 hours.

A second point is that subjects can see objects as face-like without recruiting those processes thought to be face-specific. For example, inverted faces are not recognized using the same configural processes as upright faces, yet anybody, including prosopagnosics, can identify inverted faces as faces. Likewise, schematic faces (made out of small geometric shapes) clearly look like faces but elicit about half as much activity as shaded faces do (equal to the activation obtained with the back of heads in a different experiment⁴⁴). Finally, perhaps the best example comes from patients with prosopagnosia, who can identify faces as such, but have severe problems perceiving and recognizing them as individuals³⁵. This is consistent with our experience training Greeble experts. Neither behavioral nor neural effects of expertise with Greebles are found even when subjects are well advanced in the training but before they meet our criterion for expertise (to recognize Greebles as quickly at the individual as at the family level).

Thus, seeing Greebles as ‘face-like’ (which any novice can do easily) does not predict expertise effects, including recruitment of the FFA. In contrast, a specific measure of holistic processing (the composite effect) with Greebles correlates with the recruitment of the FFA for Greebles in a different task (I.G. & M.J.T., unpublished data). This method, combining precise psychophysics and fMRI measurements of expertise acquisition, indicates that the

FFA is recruited by Greebles because experts acquire a more efficient strategy for discriminating Greebles, not because subjects learn to see Greebles as face-like. Reinforcing this point, even when Greeble experts successfully transfer their expert abilities to the recognition of new Greeble exemplars, their expertise may not apply to new Greebles that are more homogeneous than the original training set, even though the new Greebles clearly look equivalently face-like (Fig. 1).

The evidence for selectivity on a purely geometric level in the FFA is not strong, although it is possible that the FFA tends to respond to objects that are face-like because our perceptual expertise with faces makes this geometry a cue that the computations executed in the FFA are required. However, our results indicate experience recognizing Greebles, cars or birds at the individual level can make the geometries of these objects equally valid cues. Selectivity for particular image geometry is learned—a point driven home because objects of acceptable geometry do not activate the FFA by default. Thus, selectivity at a purely geometric level does not provide an adequate account of FFA specialization. In contrast, as already reviewed, we have established that level of categorization and expertise in a given object domain do affect the response of FFA independent of stimulus geometry. We argue that it is the manipulation of these factors, not stimulus geometry, that leads to recruitment of the FFA.

Could different interpretations for the function of the FFA depend on how we define this area? One method consists of defining a face-selective area in a group of subjects and measuring its behavior in other conditions (in the same or another group). This averaging method is not optimal, but it may be useful when individual localizers are not available. A study comparing group-averaged with individual functional definitions revealed that even though a FFA can be found in the large majority of subjects, it is small enough and variable enough in its location that the former approach may be inadequate²⁷. Specifically, if the defined region is large enough to encompass the FFA for all subjects, it also includes many non-face-selective voxels, and if it is small enough to contain only face-selective voxels, it may not include the FFA of many subjects.

Given these problems, many researchers have chosen to rely on individual localizers, leading to very small and highly selective FFAs^{4,9}. This approach sustains a certain illusion: that of a sharply defined area that is stable across conditions. Typi-

cally, a localizer is used to define a FFA at the outset of the study, and all further analyses reflect the response of only these voxels. Little consideration is given to the potential variability of the FFA. For instance, both the peaks of face selectivity and the extent of the FFA as defined using a faces-minus-objects comparison can differ substantially within subjects during passive viewing as compared to an active identification task²⁵. This issue is highly relevant to the nature of the selectivity in the FFA, as differences between faces and non-face objects must be weighed against the variability of face selectivity itself. To the extent the microstructure of this variability can be assessed with current techniques, there is no indication that the face-selective part of cortex is in any sense mutually exclusive with the rest of the object-recognition system. In a fMRI study of bird and car experts, several criteria for the definition of the FFA were compared (see http://www.nature.com/neuro/web_specials). The criterion most stringent in terms of size and face selectivity (face to object activity ratio, 12.8) led to a right FFA region of interest with an average size of 3 voxels (each voxel was 3.125 mm² in plane) but could be found in all 19 subjects. In comparison, the criterion used in many prior studies⁴ led to a larger area (average size of 6 voxels), which displayed more moderate face selectivity (face to object activity ratio, 5.2) but was less inclusive subject-wise (a right FFA only found in 7 of 19 subjects). Crucially, regardless of criterion, the FFA showed sensitivity to the level of categorization and expertise manipulations. This and several other analyses, such as center of mass comparisons and inspection of the peaks in unsmoothed activation maps, indicate that face expertise cannot be dissociated from object expertise at the spatial resolution of fMRI.

The existence of 'face cells' in IT⁴⁵ is often cited as evidence for a distinct recognition mechanism for faces. On the other hand, studies in which monkeys learn to recognize novel objects have revealed a remarkable similarity between properties of cells tuned to these trained objects and face cells^{46,47}. One interpretation is that "...it would be surprising if neurons exist that respond strongly to both faces and cars (but not other complex objects) in car experts (for example)"⁴⁸. The implication is that faces and objects would selectively activate independent populations of neurons, interspersed within the same fMRI voxel. Even if this were the case, the intermingling of at least two populations of neurons in a relatively small region of IT would suggest that

these neurons serve similar functional roles, but with somewhat different tuning (for instance, adjacent cortical columns are tuned to specific orientations, but the functional role of the entire area is a coding of orientation space—arguing that faces are special in this context is equivalent to suggesting that vertical lines are special, too).

Moreover, the finding that a given neuron responds most strongly to a specific stimulus does not imply that this same neuron is not functionally implicated in the processing of other stimuli. For example, neurophysiologists often rely on a reduction method¹⁷ in which, beginning with an image from a restricted basis set, they incrementally simplify the test stimulus while attempting to maintain the same level of cell response. Thus, the reduction method is bound to produce a single preferred feature set per cell. In contrast, IT cells can have responses that defy simple explanations in terms of singular preferred features (Fig. 2).

CONCLUSIONS

From grandmothers to scientists (and the two are not mutually exclusive), there is a certain fondness for the idea that the difficulty of face recognition combined with its social importance has exerted unique adaptive pressures. Thus, it is often argued, it would not be particularly surprising if there were a distinct neural and functional module for face recognition. Although the adaptive pressures may be real, this implies only that the primate visual system should be capable of efficiently performing this task, not that whatever perceptual mechanisms represent the present end-state of evolution must be exclusive to faces. Put another way, speculating on the evolutionary reasons for why a particular neural system might have arisen does not define the limits of this system. In an effort to explore these limits, we have collected evidence that FFA processing is domain-general. Specifically, it is involved in processing subordinate-level information for all objects, including faces. This processing is not restricted to objects of a certain geometry (that is, faces), as this factor cannot account for the majority of responses in the FFA to non-face objects. On the other hand, manipulations of the level of categorization and the level of expertise are able to predict when and if FFA activity will be found in a given task. Thus, even if stimulus geometry is somehow involved in FFA selectivity, it is likely to remain only one of several factors that contribute to the specialization of this brain area. That faces generally elicit the maximum response in the FFA is interest-

commentary

ing, but only as an effect that is inherently confounded with the processing biases that, through experience, are automatically engaged when we see a face.

ACKNOWLEDGEMENTS

This work was a collaborative effort, and the order of authorship is arbitrary. We thank D. Sheinberg for his careful reading of this review. M.J.T. thanks S. Pinker for a challenging discussion on some of the issues raised here; I.G. thanks N. Kanwisher for many stimulating discussions on our diverging opinions. This work was supported by NSF Award SBR-9615819.

RECEIVED 4 APRIL; ACCEPTED 20 JUNE, 2000

- Kanwisher, N., Tong, F. & Nakayama, K. The effect of face inversion on the human fusiform face area. *Cognition* **68**, B1–11 (1998).
- Bentin, S., Deouell, L. Y. & Soroker, N. Selective visual streaming in face recognition: evidence from developmental prosopagnosia. *Neuroreport* **10**, 823–827 (1999).
- Puce, A., Allison, T. & McCarthy, G. Electrophysiological studies of human face perception. III: Effects of top-down processing on face-specific potentials. *Cereb. Cortex* **9**, 445–458 (1999).
- Kanwisher, N., McDermott, J. & Chun, M. M. The fusiform face area: A module in human extrastriate cortex specialized for face perception. *J. Neurosci.* **17**, 4302–4311 (1997).
- Kanwisher, N., Stanley, D. & Harris, A. The fusiform face area is selective for faces not animals. *Neuroreport* **10**, 183–187 (1999).
- Aguirre, G. K., Zarahn, E. & D'Esposito, M. An area within human ventral cortex sensitive to "building" stimuli: Evidence and implications. *Neuron* **21**, 373–383 (1998).
- Haxby, J. V. *et al.* The effect of face inversion on activity in human neural systems for face and object perception. *Neuron* **22**, 189–199 (1999).
- Yin, R. K. Looking at upside-down faces. *J. Exp. Psychol. Gen.* **81**, 141–145 (1969).
- Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P. & Gore, J. C. Activation of the middle fusiform 'face area' increases with expertise in recognizing novel objects. *Nat. Neurosci.* **2**, 568–573 (1999).
- Bentin, S., Allison, T., Puce, A., Perez, E. & McCarthy, G. Electrophysiological studies of face perception in humans. *J. Cognit. Neurosci.* **8**, 551–565 (1996).
- Jeffreys, D. A. A face-responsive potential recorded from the human scalp. *Exp. Brain Res.* **78**, 193–202 (1989).
- Liu, J., Higuchi, M., Marantz, A. & Kanwisher, N. The selectivity of the occipitotemporal M170 for faces. *Neuroreport* **7**, 337–341 (2000).
- Rossion, B. *et al.* The N170 occipito-temporal component is delayed and enhanced to inverted faces but not to inverted objects: An electrophysiological account of face-specific processes in the human brain. *Neuroreport* **11**, 69–74 (2000).
- McCarthy, G., Puce, A., Belger, A. & Allison, T. Electrophysiological studies of human face perception. II: Response properties of face-specific potentials generated in occipitotemporal cortex. *Cereb. Cortex* **5**, 431–444 (1999).
- Chao, L. L., Haxby, J. V. & Martin, A. Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nat. Neurosci.* **2**, 913–919 (1999).
- Ishai, A., Ungerleider, L. G., Martin, A., Schouten, J. L. & Haxby, J. Distributed representation of objects in the human ventral visual pathway. *Proc. Natl. Acad. Sci. USA* **96**, 9379–9384 (1999).
- Tanaka, K. Inferotemporal cortex and object vision. *Annu. Rev. Neurosci.* **19**, 109–139 (1996).
- Puce, A., Allison, T., Asgari, M., Gore, J. C. & McCarthy, G. Differential sensitivity of human visual cortex to faces, letterstrings, and textures: A functional magnetic resonance imaging study. *J. Neurosci.* **16**, 5205–5215 (1996).
- Gauthier, I. What constrains the organization of the ventral temporal cortex? *Trends Cogn. Sci.* **4**, 1–2 (2000).
- Damasio, A. R., Damasio, H. & Van Hoesen, G. W. Prosopagnosia: Anatomical basis and behavioral mechanisms. *Neurology* **32**, 331–341 (1982).
- Diamond, R. & Carey, S. Why faces are and are not special: An effect of expertise. *J. Exp. Psychol. Gen.* **115**, 107–117 (1986).
- Tanaka, J. W. & Gauthier, I. in *Mechanisms of Perceptual Learning* (eds. Goldstone, R. L., Medin, D. L. & Schyns, P. G.) 83–125 (Academic, San Diego, California, 1997).
- Haxby, J. V. *et al.* Dissociation of object and spatial visual processing pathways in human extrastriate cortex. *Proc. Natl. Acad. Sci. USA* **88**, 1621–1625 (1991).
- Sergent, J., Ohta, S. & MacDonald, B. Functional neuroanatomy of face and object processing: A positron emission tomography study. *Brain* **115**, 15–36 (1992).
- Gauthier, I., Skudlarski, P., Gore, J. C. & Anderson, A. W. Expertise for cars and birds recruits brain areas involved in face recognition. *Nat. Neurosci.* **3**, 191–197 (2000).
- Gauthier, I., Anderson, A. W., Tarr, M. J., Skudlarski, P. & Gore, J. C. Levels of categorization in visual object studied with functional MRI. *Curr. Biol.* **7**, 645–651 (1997).
- Gauthier, I. *et al.* Does visual subordinate-level categorisation engage the functionally defined Fusiform Face Area? *Cognit. Neuropsychol.* **17**, 143–163 (2000).
- Chao, L. L., Martin, A. & Haxby, J. V. Are face-responsive regions selective only for faces? *Neuroreport* **10**, 2945–2950 (1999).
- Gaffan, D. & Heywood, C. A. A spurious category-specific agnosia for living things in normal human and nonhuman primates. *J. Cogn. Neurosci.* **5**, 118–128 (1994).
- Halgren, E. *et al.* Location of human face-selective cortex with respect to retinotopic areas. *Hum. Brain Mapp.* **7**, 29–37 (1999).
- Gauthier, I. *et al.* The fusiform "face area" is part of a network that processes faces at the individual level. *J. Cogn. Neurosci.* **12**, 495–504 (2000).
- Rossion, B. *et al.* Hemispheric asymmetries for whole-based and part-based face processing in the human brain. *J. Cogn. Neurosci.* (in press).
- Tanaka, J. W. & Curran, T. A neural basis for expert object recognition. *Psychol. Sci.* (in press).
- Hoffman, E. A. & Haxby, J. V. Distinct representations of eye gaze and identity in the distributed human neural system for face perception. *Nat. Neurosci.* **3**, 80–84 (2000).
- Tranel, D., Damasio, A. R. & Damasio, H. Intact recognition of facial expression, gender and age in patients with impaired recognition of face identity. *Neurology* **38**, 690–696 (1998).
- Gauthier, I., Behrmann, M. & Tarr, M. J. Can face recognition really be dissociated from object recognition? *J. Cogn. Neurosci.* **11**, 349–370 (1999).
- Schultz, R. T. *et al.* Abnormal ventral temporal cortical activity during face discrimination among individuals with autism and asperger syndrome. *Arch. Gen. Psychiatry* **57**, 331–340 (2000).
- Moscovitch, M., Winocur, G. & Behrmann, M. What is special about face recognition? Nineteen experiments on a person with visual object agnosia and dyslexia but normal face recognition. *J. Cogn. Neurosci.* **9**, 555–604 (1997).
- Gauthier, I. & Tarr, M. J. Becoming a "Greeble" expert: Exploring the face recognition mechanism. *Vision Res.* **37**, 1673–1682 (1997).
- Farah, M. J., Levinson, K. L. & Klein, K. L. Face perception and within-category discrimination in prosopagnosia. *Neuropsychologia* **33**, 661–674 (1995).
- McCarthy, G., Puce, A., Gore, J. C. & Allison, T. Face-specific processing in the human fusiform gyrus. *J. Cogn. Neurosci.* **9**, 605–610 (1997).
- Jolicoeur, P., Gluck, M. & Kosslyn, S. M. Pictures and names: Making the connection. *Cognit. Psychol.* **16**, 243–275 (1984).
- Gauthier, I., Williams, P., Tarr, M. J. & Tanaka, J. Training "Greeble" experts: A framework for studying expert object recognition processes. *Vision Res.* **38**, 2401–2428 (1998).
- Tong, F., Nakayama, K., Moscovitch, M., Weinrib, O. & Kanwisher, N. Response properties of the human fusiform face area. *Cognit. Neuropsychol.* **17**, 257–279 (2000).
- Perrett, D. I., Rolls, E. T. & Caan, W. Visual neurones responsive to faces in the monkey temporal cortex. *Exp. Brain Res.* **47**, 329–342 (1982).
- Kobatake, E., Wang, G. & Tanaka, K. Effects of shape-discrimination training on the selectivity of inferotemporal cells in adult monkeys. *J. Neurophysiol.* **80**, 324–330 (1998).
- Logothetis, N. K. & Pauls, J. Psychophysical and physiological evidence for viewer-centered object representation in the primate. *Cereb. Cortex* **3**, 270–288 (1995).
- Kanwisher, N., Downing, P., Epstein, R. & Kourtzi, Z. in *Handbook of Functional Neuroimaging of Cognition* (eds. Cabeza, R. & Kingstone, A.) (MIT Press, Cambridge, MA, in press).
- Sheinberg, D. L. & Logothetis, N. K. in *Perceptual Learning* (ed. Fahle, M.) (MIT Press, Cambridge, Massachusetts, 2000).