Are Faces Special?

Behavioral and Neural Support For and Against the Uniqueness of Faces in Object Recognition

Amrys O. Williams

11 May 2001

Are Faces Special?
Behavioral and Neural Evidence For and Against the Uniqueness of Faces in Object Recognition

---

Amrys O. Williams
Department of Brain and Cognitive Science
Massachusetts Institute of Technology
11 May 2001

Introduction

Are faces special, or does the brain treat them like any other object? This is a central question in the field of visual perception and, more specifically, object recognition. Studies to date suggest that humans and other primates respond differently to faces than they do to other visual cues in the environment; however, whether these differences have their bases in separate neural mechanisms is a matter of debate. Humans demonstrate an astonishing ability to recognize and differentiate hundreds of faces, despite the fact that all faces are similar. People pick up on subtle visual cues, but does this apprehension of detail extend to visual perception as a whole, or is there something peculiar about human face recognition?

This paper will first introduce the question of uniqueness with some background on object recognition. Evidence for and against the uniqueness of face perception will then be explored through overviews of behavioral and perceptual experiments regarding expertise, inversion, and innateness, as well as through studies of neuroanatomy, including lesion experiments, data from functional imaging, and case studies of individuals with brain damage.

Background: Thinking About Face Recognition

The Question

What do we mean when we ask whether faces are "special?" Toveé (1998) suggests that, in order for face processing to be considered unique, three criteria must be fulfilled. First, the procedure by which face recognition occurs must be different from the recognition procedure for non-face objects. Second, the neural apparatus responsible for face processing must be anatomically distinct from the general object recognition module. Third, faces must be processed differently at the neural level than are other objects.

It is distinctly possible that not all these assertions hold true. In fact, there is evidence that could dismantle the uniqueness argument altogether if it were built strictly on Toveé's foundation. It is useful, however, to keep his formulation of the question in mind, as it breaks down the rather arbitrary idea of "specialness" into terms that may be addressed by scientific research.

This article will cover a variety of experiments designed to address these and similar questions in an attempt to understand whether face recognition is indeed unique, or whether it is just a very practiced and refined type of object recognition.

Homogeneity and Within-Category Judgements

Though humans find face recognition to be an extremely fast and easy task, it actually requires incredibly detailed analysis. Faces are really quite similar: the differences from one face to another are often very subtle, and it is difficult to find another object category for which humans show the same recognition ability. Making discriminations among faces is an example of a within-category or within-class distinction, as it requires distinguishing between members of the same category (Liter & Bülthoff 1996). Within-class recognition is roughly akin to picking out a particular apple from a full bushel basket. It is a harder task than picking the apple out of a basket full of different kinds of fruit, which requires between-class discrimination. The difference is that all members of a category tend to look alike, making the identification task much more difficult.

However, face recognition differs from the first apple-in-a-bushel task in important ways. First of all, though faces constitute a homogenous class—like all apples, they share basic characteristics and features appear in a predictible configuration—they contain cues that facilitate identification. In the apple task, the only cues to differentiating between individual apples would be the length of the stem, the pattern of coloration on the skin, and possibly blemishes or bruises if they were present at all. Aside from stem length, the visual information is rather inconsistent and impoverished, as an apple looks pretty much the same from any view perpendicular to its core, and the spatial distribution of redness or greenness on the surface can vary indefinitely. In determining the differences between similar faces, however, the search space is greatly reduced to a finite set of charateristics—such as eye color, nose shape, cheekbone height, and lip fullness—all of which appear in a consistent relative spatial configuration (Diamond & Carey 1986). The eyes will never be found on the chin, nor will the nose be hidden behind the ears. When humans look for differences between faces, they are approaching a relatively predictable problem.

Recognition and Categorization

In order to speak further about face recognition, is it critical to define the term. In vision science, recognition usually refers to the successful categorization of an object as a member of a specific class of objects (Liter & Búlthoff 1996). Furthermore, the way in which humans categorize objects is fairly predictible. When shown a picture of a wingback chair, for example, a subject is likely to identify it not as furniture and not as a wingback chair, but simply as a chair. The categorization level at which people tend to classify an object is termed "basic level" (Rosch 1976) or "entry level" (Jolicoeur 1984). A superordinate category would be more general (*e.g.* furniture), a subordinate category would be more specific (*e.g.* wingback chair), and an individual exemplar would be the most specific (*e.g.* my father's wingback chair). All of these other classifications take longer than entry-level (Jolicoeur 1984). It is interesting, then, that face identification almost always occurs at the level of individual exemplar. For example, one is far more likely to classify the image in *Fig. 1* as "Jimi Hendrix" than as "the face of a black male" or simply as "a face." This demonstrates one of the central motivations for considering face recognition unique.

Fig. 1

The next section explores further behavioral arguments for and against uniqueness, including the relationship between knowledge and categorization, and studies exploring innateness.

Behavioral and Functional Evidence

The Expert Problem

Categorization, Experts, and Novices

Though the traditional categorization hierarchy makes sense, the term "basic level" seems a bit fuzzy. For example, the average person might identify a picture of a swallow simply as a bird, while an ornithololgist might call the same picture a swallow. What does this mean for the basic level?

A good way to begin is to think about what expertise means in terms of the categorization hierarchy. Expertise is detailed knowledge of a particular category—domain-specific knowledge—which implies the retention of more information at subordinate levels. It follows, then, that for experts the basic level in their category of expertise may be more subordinate than for novices.

In a study of dog and bird experts, Tanaka & Taylor (1991) found that, compared to novices, experts (a) organized their knowledge primarily at a more subordinate level of abstraction, (b) more often classified objects at a subordinate level, and (c) were faster at determining if images belonged to a subordinate category. These experiments established a connection between expertise and entry-level classification: individual differences in domain-specific knowledge affect the extent to which the traditional basic-level is central to categorization. More specifically, the more of an expert you are in a category, the lower in the hierarchy will be your entry-level classification of a category exemplar.

In another study, Gauthier *et al.* (1998) trained a set of subjects to be experts at identifying novel objects called "Greebles." The Greebles were designed to have a common spatial configuration, just as faces do. Gauthier and her colleagues then compared the recognition peformance of the experts to a group who had not been trained, and found the experts did much better than did the novices. This is not a surprising result, but it will have greater implications later. For now, it is enough to say that expertise effects occur with non-face stimuli.

The Inversion Effect and Expertise

It is well-established that people have difficulty recognizing upside-down faces, even if the faces are quite familiar (Diamond & Carey 1986). It has been hypothesized that the challenges presented by inversion are due to a mental rotation delay: humans recognize faces right-side-up, so in order to parse an inverted face, one must mentally right the image before attempting to identify it. During rotation, it is almost impossible to retain the subtle and detailed spatial interrelations that make face identification possible: therefore, important information is lost, and the face is not recognized (Rock 1974).

The Margaret Thatcher Illusion demonstrates the importance of spatial relationships in face processing, and how severely inversion handicaps the ability to discern them. The two adjacent photographs in *Fig. 2* appear, on first ispection, to be quite similar. However, if they are rotated 180 degrees, what seemed initially to be a minor difference between the two photographs produces a shocking effect. The marked perceptual difference between the upside-down and right-side-up images clearly illustrates how debilitating inversion is (Thompson 1980).

Fig. 2: The Margaret Thatcher Illusion (Thompson 1980)

In an attempt to investigate the relationship between expertise and the size of the inversion effect, Diamond & Carey (1986) conducted an experiment with subjects who were dog experts and subjects who were not. The two groups were shown pictures of faces and dogs, both right side up and inverted, and their recognition performace was measured. The results were clear: dog experts showed an inversion effect comparable to that observed with faces, while inversion had little effect on the novices. The experiment suggests that there is a connection between expertise and the magnitude of the inversion effect: the more expertise you have in a category, the greater will be the effect of inversion on your identification of category exemplars. If inversion affects other classes of objects, then faces may not be all that special: humans may just be face experts (Tarr & Gauthier 2000).

Evolutionary Bases and the Innateness Question

Earlier, it was explained that face recognition is a homogenous and relatively predictable problem. However, predictability does not always hold. People must be able to detect as well as account for differences in facial expression, which are not only transformations of the face which may make recognition more difficult, but are also important social cues. The detection and interpretation of facial expression is therefore of great ecological significance.

Experiments with Infants and Newborns

Since facial perception and interpretation seem to lend an evolutionary advantage, it begs the question as to whether or not face recognition is innate. Researchers have conducted several experiments with infants to explore this possibility, the results of which lend credence to the uniqueness camp.

In a study by Johnson *et al.* (1991), it was found that infants demonstrate the ability to track a face; in fact, that they prefer to track a face-like stimulus than another, similar stimulus (such as a conceptualized face or an inverted face). Bushnell *et al.* (1989) found that babies as young as four days show a strong preference for their mother's face over the face of a similar-looking female stranger. How can this recognition ability be present at such a young age? These researchers point to innateness as the most likely explanation.

The likelihood that these preferences are innate does indeed buttress the uniquness argument. But how well can we trust these infant studies? There is one problem in particular that makes them less watertight: the reduced sensitivity of infant vision, which calls into question the reliability of these results. Compared to adults' visual abilities, babies' vision is quite impaired. Infants lack the acuity, as well as the sensitivity to color and contrast possessed by adults. They also have great difficulty fixating objects. What this amounts to is that babies have a much more blurry and chaotic visual experience than adults do. How is it, then, that newborns are able to make such fine discriminations between faces? Despite the intuitive ecological significance, this is a question which remains unanswered. Therefore, the research points toward an innate preference for faces—and hence toward face perception as special—though the results are not sitting on a firm foundation.

The next section explores neuroanatomical and neurological studies which address the machinery aspect of the uniqueness question: is there an anatomically distinct neural apparatus for face recognition?

Neural Mechanisms and Machinery

Studies of neuroanatomy can be grouped into two categories: experiments using functional imaging to study the brain activity of normal subjects, and those investigating the behavior and neural activity of patients with brain damage. Both types of studies give researchers insight into which specific areas of the brain are involved in different tasks. The following sections describe functional imaging and brain damage case studies in more detail.

Functional Imaging and Cellular Recording

Fuctional imaging (particularly fMRI, Magnetic Resonance Imaging) and cellular recording are also used in conjuction with perceptual experiments to observe real-time neural activity in subjects carrying out a particular task. These techniques give scientists a clear picture of what is actually going on in the brain dynamically: which cortical areas are active during which tasks, and even which specific neurons are firing at any given time.

Cellular Recording

By observing and recording the activity of particular neurons, scientists have demonstrated the existence of a population of neurons in the visual areas of the temporal lobe which appears to be selective for faces (Rolls 2000). These face-selective neurons respond from 2 to 20 times more to face stimuli than to gratings, geometric patterns, or even complex non-face objects. The cells also respond differently to different faces, and to different facial features. Neurons were also found which appear to be sensitive to changes in facial expression. Together, these two sets of neurons demonstrate a mechanism for identifying individual faces and show that neural function is specialized in this area.

fMRI and the FFA: Fusiform Face Area or Flexible Fusiform Area?

Recent fMRI experiments demonstrate the existence of an area of inferior temporal cortex which responds to faces: its neurons fire whenever the subject is shown an image of a face. It has been dubbed the Fusiform Face Area (FFA), and has been among the most powerful and promising pieces of evidence for the uniqueness of face recognition. However, some studies have cast doubt on whether it is as face-specific as its discoverers (Kanwisher *et al.* 1997) originally believed.

Kanwisher's study (1997) identified an area in the fusiform gyrus that showed increased activation when subjects were shown faces as opposed to other common objects. Stronger responses in this Fusiform Face Area were found for intact over scrambled images of faces, frontal views of faces over frontal views of houses, and three-quarter views of faces (hair covered) over pictures of hands. These trials demonstrated that activity in the FFA does not reflect general image processing, visual attention, or subordinate-level classification, leaving Kanwisher and her colleagues to conclude that the FFA is selectively involved in face perception.

However, not all researchers were convinced. Gauthier *et al.* (2000) showed increased activation in the FFA when Greeble experts were shown Greebles. The FFA was also highly active in bird and car experts when they were shown pictures of birds and cars, respectively. The general result was that exemplars from homogenous categories stimulated the FFA more than familiar objects. It is therefore the individual's degree of expertise (and consequently the level of categorization)—not the attributes of the objects themselves—that determines the specialization of the FFA.

It would seem that the FFA is not inherently face-specific, but rather geared toward the recognition of objects in homogenous categories. It makes sense for such a neural mechanism to exist: a cortical area specialized for making fine distinctions among similar objects fits in with the evolutionary argument. It even strengthens it, as individuals who are equipped to, with practice, differentiate between members of a homogenous class are probably more fit than those who cannot. As a homogenous category, faces are included under this heading, so these results are not as controversial as they first appear. Though it comes down on the non-uniqueness side, describing the Fusiform Face Area instead as a Flexible Fusiform Area (Tarr & Gauthier 2000) accounts for all experimental results to date.

Agnosia and Beyond: Case Ctudies

Perhaps the most intriguing neurological research in face recognition consists of case studies of individuals with damage to regions of the brain that are important in visual processing. Scientists have studied patients with various types of visual agnosia, a range of disorders typically marked by an inability to name objects. One variety of visual agnosia is prosopagnosia; prosopagnosic patients exhibit a specific inability to recognize faces, but perform normally on other visual recognition tasks. The existence of so particular a deficiency lends credence to the idea that the brain does indeed treat faces differently from other objects.

Prosopagnosia

Although prosopagnosia is described as "face agnosia," controversy remains as to whether it is truly and uniquely face-specific. McNeil & Warrington (1993) tell of a prosopagnosic sheep farmer, known as W.J., who could recognize and distinguish between the faces of sheep but not of humans. To address the question of the face-specificity of prosopagnosia, W.J. was tested on identification of famous buildings, breeds of dogs, and types of flowers—with excellent performance. However, though these tasks do investigate within-category discrimination, they do not require the subject to differentiate between individual exemplars of a given class. W.J. was accordingly tested on sheep faces, as they constitute a class quite analogous to that of human faces. They found that W.J.'s sheep recognition performance far surpassed that of normal subjects tested on the same identification task. Patient M.T., also prosopagnosic, is able to differentiate makes and models of cars, and different kinds of fruits and vegetables (Kanwisher *et al.* 2000). These results clearly lend themselves to the notion that prosopagnosia is indeed face-specific.

Everything but the Face and Other Types of Agnosia

The case of patient C.K. is particularly interesting, as it bring expertise back into the discussion. C.K.'s brain damage resulted in severe visual agnosia—for everything except faces. Even though C.K. was an expert at identifying both airplanes and tin soldiers before the damage was sustained, he did not retain a more general ability to distinguish between members of homogenous categories, not even those at which he was expert (Kanwisher *et al.* 2000). His face recognition abilities were intact when the rest of his recognition machinery was destroyed; certainly this supports the idea that faces are neurologically unique.

Another interesting point about C.K. involves the inversion effect. Though comparable to normal subjects in upright face recognition tasks, C.K.'s ability is severely impaired relative to normal subjects when the faces are inverted. A patient known as L.H. suffers from a condition complementary to that of C.K.: he is impaired on upright face-recognition tasks, but much better at inverted faces. The nature of these two patients' impairments indicates that face processing is to some degree separate from other object recognition. Furthermore, these results indicate that inverted faces may not be treated as faces by the brain, but instead merely as objects (Toveé 1998).

Visual agnosia case studies largely support the uniqueness argument; however, another explanation has been proposed. Tarr & Gauthier (2000) posit that, after sustaining brain damage, patients such as C.K. become unable to utilize the general homogenous category machinery—what amounts to holistic processing machinery—for anything other than faces. They are unable to move from the novice to the expert domain in any other class of objects. While this argument still seems to set faces apart, and does not explain W.J.'s abilities, it is certainly plausible, and should remain in consideration.

Discussion

Summary

This article has examined a variety of evidence both supporting and calling into question the hypothesized uniqueness of facial processing. On the uniqueness side, there is the behavioral evidence from categorization and inversion effects, as well as the infant experiments providing evidence for innateness. Studies of prosopagnosics also strongly support uniqueness. Conversely, there is much evidence that faces are not completely unique. Gauthier's Greeble study demonstrates that many of the behavioral phenomena traditionally associated with face recognition alone can pop up in individuals with expertise in a category. Her functional imaging studies also suggest that areas of the brain once thought to be entirely face-specific can also be activated by other homogenous object classes, calling the uniqueness argument into question.

Early on, this article put forth Toveé's scientific formulation of "specialness." A good way to summarize the findings presented here is to ask how well they satisfy his criteria.

1. The procedure by which face recognition occurs must be different from the recognition procedure for non-face objects. The evidence from inversion effects and agnosic patients suggests that this is true. There seems to be a dissociation between the recognition of faces (and other object classes with which an individual is experienced) and that of ordinary objects.

2. The neural apparatus responsible for face processing must be anatomically distinct from the general object recognition module. There does seem to exist an anatomically distinct area for quick, holistic recognition of objects from categories in which one is expert, including faces. This is not face specific, but rather a more general module.

3. Faces must be processed differently at the neural level than are other objects. This ties in with the two previous criteria. Faces seem to belong to a more general class, the contents of which vary with each individual; it is this general class that might be processed differently than other objects.

There is a trend developing here, one in which faces are a subset of a broader category: homogenous object classes in which one has expertise. The literature is increasingly focused on this: the facial uniqueness argument has been reformulated as a question of domain-specificity (Kanwisher *et al.* 2000, Tarr & Gauthier 2000). Those arguing for uniqueness argue likewise for domain-specificity, particularly concerning the FFA; those who doubt that faces are truly special argue for domain-generality.

Conclusion

There seems to be an intuitive bent toward the uniqueness of facial processing: based on our own experience, it makes sense that faces should be treated differently from other objects, and the evidence on the surface seems to verify this hypothesis. However, delving deeper into the question—defining it scientifically, forcing the meaning of "uniqueness" to be phrased in specific terms, and exploring these criteria—casts doubt on the idea that faces are truly a special case in object recognition. If the trend towards domain-generality, in which faces belong to a more general class, continues to be supported by the research, it may be one on which both camps can agree.

References

Bushnell, I.W.R., Sai, F., & Mullin, J.T. (1989). Neonatal recognition of the mother's face. *British Journal of Developmental Psychology*: Vol. 7, 3-15.

Diamond, Rhea & Carey, Susan. (1986). Why Faces Are and Are Not Special: An Effect of Expertise. *Journal of Experimental Psychology*: Vol. 115, 107-117.

Gauthier, Isabel *et al.* (1998). Training "greeble" experts: A framework for studying expert object recognition processes. *Vision Research*: Vol 38, No. 15-16, 2401-2428.

Gauthier, Isabel *et al.* (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*: Vol 3, No. 2, 191-197.

Johnson, M.H. *et al.* (1991). Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition*: Vol. 40, 1-19.

Jolicoeur, P., Gluck, M. A., & Kosslyn, S. M. (1984). From pictures to words: Making the connection. *Cognitive Psychology*: Vol. 16, 243-275.

Kanwisher, Nancy *et al.* (1997). The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *The Journal of Neuroscience*: Vol. 17, No. 11, 4302-4311.

Kanwisher, Nancy *et al.* (2000). Domain specificity in face perception. *Nature Neuroscience*: Vol. 3, No. 8, 759-763.

Liter, Jeffrey C. & Bülthoff, Heinrich H. (1996). An Introduction to Object Recognition. In Object and Face Recognition: Selected papers from the Max-Planck Institut fur BiologischeKybernetik (pp 1-20). Tubingen.

McNeil, Jane E. & Warrington, Elizabeth K. (1993). Prosopagnosia: A Face-specific Disorder. *The Quarterly Journal of Experimental Psychology*: Vol. 46A, No. 1, 1-10.

Rock, I. (1974). The perception of disoriented figures. *Scientific American*: Vol. 230, 78-85.

Rolls, Edmund T. (2000). Functions of the Primate Temporal Lobe Cortical Visual Areas in Invariant Visual Object and Face Recognition. *Neuron*: Vol. 27, 205-218.

Rosch, E. *et al.* (1976). Basic Objects in Natural Categories. *Cognitive Psychology*: Vol. 8, 382-439.

Tanaka, James W. & Taylor, Marjorie (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology*: Vol. 23, No. 3, 457-482.

Tarr, Michael J. & Gauthier, Isabel (2000). FFA: a flexible fusiform area for subordinate-level visual processeing automatized by expertise. *Nature Neuroscience*: Vol. 3, No. 8, 764-769.

Thompson, P. (1980). Margaret Thatcher: a new illusion. *Perception*: Vol. 9, 483-484.

Toveé, Martin J. (1998). Is Face Processing Special? *Neuron*: Vol. 21, 1239-1242.