Philosophical Psychology
Vol. 20, No. 2, April 2007, pp. 197–207

Faces and Brains: The Limitations of Brain Scanning in Cognitive Science

Christopher Mole, Corey Kubatzky, Jan Plate, Rawdon Waller, Marilee Dobbs, and Marc Nardone

The use of brain scanning now dominates the cognitive sciences, but important questions remain to be answered about what, exactly, scanning can tell us. One corner of cognitive science that has been transformed by the use of neuroimaging, and that a scanning enthusiast might point to as proof of scanning's importance, is the study of face perception. Against this view, we argue that the use of scanning has, in fact, told us rather little about the information processing underlying face perception and that it is not likely to tell us much more.

Keywords: Brain Scanning; Face Recognition; Fusiform Face Area; Neuroimaging; Perception; Representation

1. Introduction

Although the development of brain scanning techniques has dramatically changed the way in which cognitive science is conducted, it would be a mistake to suppose that every cognitive scientist has been convinced of brain-scanning's usefulness. A scanning-enthusiast, wanting to convince a scanning-skeptic that these techniques deserve their prominent place in the cognitive sciences, and looking for an example of an area in which scanning has proved its worth, might well look to the study of the cognitive mechanisms that enable us to perceive faces. Brain scanning has been used extensively in the study of face perception and data from scanning studies have been thought to enable a fresh approach to questions that, although of longstanding interest, have proved difficult to address by more traditional means.

Christopher Mole is a lecturer in philosophy at University College Dublin. Corey Kubatzky, Jan Plate, Rawdon Waller, Marilee Dobbs and Marc Nardone were students in a seminar taught by Christopher Mole at Washington University in Saint Louis.
Correspondence to: Christopher Mole, School of Philosophy, Newman Building, University College Dublin, Belfield, Dublin 4, Ireland. Email: christopher.mole@ucd.ie

ISSN 0951-5089 (print)/ISSN 1465-394X (online)/07/020197-11 © 2007 Taylor & Francis
DOI: 10.1080/09515080701209380

In this paper we shall explain why enthusiasm for the role of scanning in the investigation of face perception is misplaced. Scanning has, in fact, told us rather little about the information processing underlying face perception and if we are to advance our understanding of face perception it is not more scanning work that we need.

2. What We Knew Before Scanning

There are strong *prima facie* reasons to expect faces to be subjected to cognitive processing of a sort that is different from the processing involved in the perception of other stimuli. One can imagine good reasons why evolution might favor a cognitive architecture that treats faces in a special way. It might be, for example, that creatures with special resources allocated to the processing of faces are better able to exploit the useful information that faces carry about the age, sex, mood, direction of gaze and familiarity of their owners (Kanwisher & Moscovich, 2000, p. 1). We, and our ancestors, had good reason to care more about faces than other stimuli, and we had reason to care about them in early infancy. It is therefore plausible that we evolved special mechanisms for their processing. But this, like many evolutionary stories, is purely speculative. It is a question for cognitive psychology whether evolution has behaved as one might have expected in this case, and cognitive psychologists have long been in the business of finding out whether the processing of faces is indeed different from the processing of other stimuli, and of finding out, if faces *are* processed differently, what the difference is.

Several of the traditional methods of cognitive science provide lines of evidence that suggest answers to these questions. One source of evidence in support of the hypothesis that there is a difference in the cognitive treatment of faces comes from the behavior of very young infants, who, even in the first hours of life show a special sensitivity to, and seemingly a special understanding of, the expressions of faces. Children imitate faces, discriminate between them, and are interested by them (Fantz, 1963; Johnson & Morton, 1991, as cited in Nelson, 2001). This evidence from infant development does not establish the existence of a complex innate face processing system—there are too many controversies about the extent of the infant's abilities, and about the time at which these abilities are manifested (see Nelson, 2001)—but it does point towards the existence of *something* special about the way in which faces are processed.

Another line of evidence for this rather vague specialness claim comes from the fact that pictures of faces give rise to visual illusions not produced by other stimuli (Thompson, 1980), suggesting to some minds that the identification of faces is a more holistic process than the identification of other stimuli. Most stimuli are somewhat harder to recognize when they are presented upside down, but the difficulty of recognizing upside down faces is unusually great (Farah, Wilson, Drain, & Tanaka, 1998; Yin, 1969); and the parts of faces, unlike parts of other objects, are significantly easier to discriminate when they are presented in their usual context

(Tanaka & Farah, 1993). Fagan (1972) found an effect for upside down faces in children of four months, so it would seem that whatever specialness is revealed by these effects is probably not independent of the specialness suggested by the responses of very young children.

A third and richer source of evidence comes from the selective impairment of face recognition ability following brain damage. Bodamer (1947), Bruyer et al. (1983), Pallis (1955), and Humphreys and Riddoch (1987) have all described cases in which the recognition of faces is impaired although other visual processing remains intact. Care is needed, however, in inferring from the existence of such selective impairments to conclusions about the functional distinctness of the cognitive resources underpinning face recognition. The strikingly clear cases, such as the prosopagnosic sheep farmer, unable to recognize his wife's face but able to recognize the faces of his sheep (McNeil & Warrington, 1993), are extremely rare. "The condition is seldom encountered in a 'pure' form . . . there are usually additional types of visual recognition difficulty present" (Ellis & Young, 1997, p. 101). In cases such as this it is not clear how much double dissociations can tell us (Juola & Plunkett, 2000).

Taken together, these familiar and long-established forms of evidence provide reasonably strong support for the hypothesis (already made plausible by evolutionary considerations) that faces are subjected to processing different from, and somewhat independent of, the processing given to other visual stimuli. They don't, however, show anything very conclusively, and they don't tell us very much about what the difference is between face processing and other visual processing.

3. What Kanwisher Found

Since 1997, Nancy Kanwisher and her collaborators have conducted a series of brain-scanning experiments that they take to confirm the hypothesis that faces are subjected to processing that is different from, and somewhat independent of, the processing given to other visual stimuli. They take the results of these experiments to indicate that there is a particular area of the ventral pathway of the visual stream—an area they call the *fusiform face area*—that is the site of cognitive resources that are specialized for and dedicated to the processing of faces: They take this area of the brain to be "a module in human extrastriate cortex specialized for face perception" (Kanwisher, McDermott, & Chun, 1997, p. 4302).

The initial findings that Kanwisher and her collaborators took to be evidence for this claim were findings about the activation patterns of various brain regions (as revealed by fMRI) during the passive viewing of a sequence of several 30-second streams of grayscale photographs (shown at a rate of 45 photographs per 30-second stream). These findings point to the existence of a particular region of the fusiform gyrus that is more active when subjects are viewing pictures of faces than it is when the subjects are viewing other stimuli, including other stimuli representing living things. (We follow Kanwisher and her collaborators in referring to this area as the 'fusiform face area', or FFA, without thereby endorsing their account of what this area does.)

It was soon pointed out that there are several alternative hypotheses that could explain these findings. The hypothesis according to which the FFA is the site of a module "specialized for face perception" is not the only hypothesis, nor even the simplest hypothesis, that can explain the preferential activation of the FFA in response to pictures of faces. One alternative hypothesis (suggested by Wojciulik, Kanwisher, & Driver, 1998) is that the difference in activity in the FFA when subjects are shown faces, rather than other stimuli, is a reflection of the fact that faces are more *interesting* than other stimuli, and so are given more attention. Another alternative hypothesis (suggested by Gauthier et al., 2000) is that the FFA's differential response to faces is a reflection of the fact that our interest in faces is typically an interest in which *particular* face we are seeing, whereas our interest in other stimuli is typically an interest in which *sort* of stimulus we are seeing. A third hypothesis (suggested by Tarr & Gauthier, 2000, and researched extensively by Gauthier and collaborators) is that the FFA's differential response to faces is a reflection of the fact that subjects are typically *expert* in discriminating faces from one another, in a way that they are not expert judges of the other stimuli that Kanwisher used.

In order to test these various alternative hypotheses a great many more brain-scanning studies have been carried out, many of them ingeniously designed. Kanwisher and her collaborators, and the researchers who have engaged in debates with them, have examined the ways in which FFA activation is correlated with performance in a wide-variety of visual tasks involving faces and other stimuli. We now know that the FFA is activated when car or bird experts make judgments about the location of stimuli for which they are expert (even when the cars are seen side-on, rather than "face-forward") (Xu, 2005). We know that the degree to which the FFA is activated by stylized nonsense objects ("greebles") is correlated with the degree of expertise that a subject has acquired in recognizing these objects (Gauthier, Tarr, Andersen, Skudlarski, & Gore, 1999). And we know that FFA activation correlates with a subject's accuracy in judging the identity of the faces presented (Grill Spector, Knouf, & Kanwisher, 2004).

It is our view that all this scanning-based research on face perception is *not*, despite the enthusiasm with which it has been conducted and received, a fruitful and productive research project. Scanning does not confirm, or refute, our initial hypothesis about face processing being unlike the processing of other stimuli. Still less does it provide us with a more precise understanding of the ways in which face processing is unlike the processing of other stimuli.

We started out with a plausible but nonspecific hypothesis about the specialness of the cognitive processing to which faces are subjected. The null-hypothesis corresponding to this claim about face-processing's specialness is the claim that faces are handled in the same way as other visual stimuli. The specialness hypothesis was, we saw, made more plausible by evolutionary considerations and by the range of cognitive psychological evidence reviewed in §2. Kanwisher's work is widely believed to confirm a version of the specialness hypothesis, made more precise so as to be a claim about specialized resources *in the FFA*. This, as we shall show, is a mistake.

Scanning data has, in fact, told us nothing that could settle the debate between the original specialness hypothesis and the corresponding null-hypothesis according to which the processes by which faces are seen are not special, but are processes that are of just the same type as the processes by which other stimuli are seen and that are implemented by the very same cognitive resources.

We do not claim that scanning has been entirely unhelpful. It has enabled us to move away from the position of asking the rather broad question: "Are faces subjected to special cognitive processing, and if so, what's special about it?" and towards the position of asking a more specific question: "Does the FFA house a module specialized for face perception?" It has told us that *if* there are special resources for the processing of faces then the FFA is a likely site for them. But scanning can do nothing to answer the question of whether there *are* such resources.

4. What Haxby Found

A series of experiments by James Haxby and his collaborators help to illustrate why it is that, despite the findings of Kanwisher, the null-alternative nonspecialness hypothesis remains entirely unrefuted (Hanson, Matsukaa, & Haxby, 2004; Haxby et al., 2001). Haxby found, as Kanwisher had, that there is an area of ventral temporal cortex that responds maximally to pictures of faces when subjects perform a task requiring recognition (in this case, a one-back recognition task requiring subjects to recognize different views of the same face or object). In addition to Kanwisher's face area Haxby also found evidence for an area that responds maximally to pictures of chairs, an area that responds maximally to pictures of cats, an area that responds maximally to pictures of houses, an area that responds maximally to pictures of bottles, and other areas responding maximally to shoes and to scissors. And, most importantly, Haxby discovered that the patterns of activation in these other areas of cortex carry information about whether the stimulus presented was a face. He discovered that predictions about whether the stimulus presented was a face could be made with a high degree of accuracy (between 85 and 100%) on the basis of the patterns of activity in those parts of cortex other than the areas that respond maximally to faces. A great deal of information about faces was retained, even when data from the FFA was omitted from the analysis. Haxby also found, similarly (and again, importantly) that the patterns of firing in those areas of cortex other than those that respond maximally to cats carry information about whether the subject is looking at a cat; the patterns of firing in areas of cortex other than those that respond maximally to shoes carry information about whether the subject is looking at a shoe, and so on. In more recent work Haxby and his collaborators have used increasingly sophisticated ways of extracting the information carried by the patterns in their scanning data (e.g., by using neural networks and support vector machines) (Hanson et al., 2004). These studies confirm that there are a number of distributed and overlapping patterns of activation that carry information—often quite rich information—about the stimulus that a subject is presented with.

The challenge that Haxby's findings pose for Kanwisher is a methodological one. Visual processing, as Kanwisher understands it, is accomplished by several distinct processing streams (in addition, perhaps, to a general purpose processing stream) each of which is responsible for the handling of one sort of stimulus. Kanwisher takes the different patterns of maximal activation in different parts of cortex to reveal the distinct locations of these different streams, and, in particular, she takes the fact that the FFA is maximally responsive to faces to be evidence showing that the FFA is the location of processing resources that contribute only to the perception of faces. (She also makes a similar inference about the location of a para-hippocampal area specialized for place recognition.) But the inference from data showing the loci of maximal activation to a conclusion about the loci of modular processing resources is not a sound inference. There is every reason to suppose that the nonmaximal activation occurring outside the area that Kanwisher interprets as a face-recognition module does have an important cognitive role to play, as Haxby's work vividly shows.

The picture of visual processing that Haxby's work suggests is a version of the null-alternative to the specialness picture. In Haxby's picture it is the whole distributed pattern of activation that constitutes the processing by which different stimuli are perceived, whether or not these stimuli are pictures of faces, or places, or household objects. The suboptimal activation that occurs outside the FFA during the presentation of a face is, according to this picture, a part of the processing by which the subject comes to see the presented face, and the suboptimal activation that occurs inside the FFA when a shoe is presented is a part of the processing by which the subject comes to see the presented shoe.

The challenge that Haxby's results pose for Kanwisher is a challenge of a different and more radical sort from that posed by researchers such as Gauthier: It is not just that Haxby proposes another factor that was not controlled for in Kanwisher's original experiments. Gauthier and others accept that Kanwisher has located a specialized visual processing resource, but question whether she has done enough to show that it is specialized *for faces*. This leads Kanwisher to write as if the debate were all about whether it is faces or some other stimuli that elicit the strongest response from the FFA. She writes that:

> Strong claims of face selectivity entail the prediction that no non-face stimulus will ever produce a response as strong as a face; because the set of nonface stimuli is infinite, there is always some possibility that a future study will show that a putative face-selective cell or region actually responds more to some previously untested stimulus (say, armadillos) than to faces. (Kanwisher & Yovel, 2006, p. 2112)

Haxby's results pose a more radical challenge by showing that even if we knew the answer to the question of which stimuli elicit the maximal response from the FFA, the question of specialness or nonspecialness would remain. The nonspecialness hypothesis, which stands as the null-alternative to any hypothesis like Kanwisher's, is not refuted by the discovery that faces elicit maximal activation in the FFA. It may be faces that elicit the maximal activation from the FFA, but this would not show that the submaximal firing that other stimuli elicit from the FFA has no role in enabling the subject to perceive those stimuli. Nor would it show that the submaximal responses of areas other than FFA have no role in enabling the subject to perceive faces.

The null, nonspecialness hypothesis can accommodate Kanwisher's finding about the preferential activation of the FFA because this hypothesis says only that architecture responsible for the processing of faces is not distinguished from the general unspecialized processing architecture that enables visual perception in general. This null-hypothesis says nothing that would imply that stimuli of different sorts must produce uniform patterns of activation across the cortex. The presentation of some stimuli will, of course, lead to activation of some areas of cortex more than of others, but this does not mean that the maximally activated area is a specialized locus in which the processing of those stimuli takes place. The locus of maximal activation need not be the place where the processing is done. It might simply be an indication of the place at which the load is greatest. An analogy makes this clear. The fact that some manual tasks, such as hammering, lead to blisters on the palm of the hand does not show that it is the palm that performs the task of hammering. Other tasks, such as sanding, lead to blisters on the thumb, but this does not show that it is the thumb that performs the task of sanding. It is the whole hand that implements each of these processes. The blisters indicate where the load is greatest, but they do not tell us where the process is implemented, even if there is no other task that produces blisters in exactly those places. Why, then, should Kanwisher infer from the fact that faces lead to particularly strong activation in the FFA (and, perhaps, there are no other stimuli that activate exactly that area so strongly) to the conclusion that the FFA is the place in which the processes of face perception are done? A locus of preferential activation might just be the locus of the busiest part of general processing system. It need not be the locus of a module. There is no sound argument from data showing the loci of maximal activation to conclusions about the loci of specialized processing resources.

Haxby *et al.* (2001, p. 2428) make the same sort of point with an analogy to a less remote case. He suggests that the process by which objects are processed in the cortex might be analogous to the process by which colors are processed on the retina. Colors are encoded by the retina via the degree to which they activate the three kinds of cone. A Kanwisher-like experiment on color vision would conclude, on finding that there is one kind of cone that responds maximally to red light and that responds suboptimally to light of other wavelengths, that there is a specialized red-detecting system without a role in the detection of other stimuli. But this would clearly be mistaken. The suboptimal response of cones is just as much a part of the retina's processing of color as is the maximal responding. Encoding in the cortex might be similar. Even if we grant that faces are the stimuli to which the FFA responds most strongly, there is no reason to suppose that the suboptimal response of the FFA to nonface stimuli has no role in the perception of those stimuli.

5. What Haxby's Results Do Not Show

Haxby's results, by showing that the suboptimal firing patterns may be important, make vivid the way in which, by focusing only on maximal firing, Kanwisher's data fail to establish her hypothesis. But care is needed when we come to ask what positive conclusions can be drawn from Haxby's results. The results are presented as showing that the representation of faces and other objects in ventral temporal cortex is distributed and overlapping. (The title of Haxby et al.'s 2001 paper is "Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex.") But this is a case in which it matters rather a lot what is meant by 'representation'.

There is a sense of 'representation' for which Haxby's results clearly do show that there are distributed and overlapping representations of faces and objects; but according to this first sense of 'representation', the claim that there are distributed and overlapping representations of faces and objects does not contradict Kanwisher's hypothesis (although it does still pose the methodological challenge discussed above). If we strengthen the sense of 'representation' in such a way that Haxby's claim that there are distributed and overlapping representations of faces and objects *does* contradict Kanwisher's hypothesis, then the resulting sense of 'representation' is one for which Haxby's data do not show that representations of faces and objects are distributed and overlapping.

In the first of these senses of 'representation', a face is represented by a certain pattern of neural firing if and only if that pattern carries information about the face presented. Representations, in this sense, come cheaply. They need not have any role in *processing*, and they certainly need not have any role in a subject's awareness of the thing represented. In this sense of 'representation' there are representations of all sorts of things all over the place (not only all over the brain). Haxby's results certainly show that, in this sense of 'representation', there are distributed representations of faces, and that these overlap with representations of objects: They certainly show that distributed and overlapping patterns of firing *carry information* about faces and objects—they do so by showing that this information can (sometimes with the help of sophisticated computational mechanisms such as support vector machines) be used to make significantly above chance predictions about which sort of stimulus is being presented.

But if 'representation' is understood in this way then, in much the same way that the null-hypothesis can accommodate Kanwisher's findings, so Kanwisher's theory can accommodate Haxby's findings. Haxby shows that the suboptimal activation in the FFA when nonface stimuli are presented *carries information* about what those nonface stimuli are, but (since the carriers of information need not have any role in psychological processing) this does not show that the FFA has a role in the processing responsible for the perception of nonface stimuli. It might be, compatibly with Kanwisher's picture, that the information-carrying activation of the FFA in response to nonface stimuli is a sign of the epiphenomenal ticking over of the face-specific resources in the FFA, corresponding to the activity by which these modular resources look for a face, and fail to find one, in a picture of a shoe, or cat, or bottle.

A Kanwisher-style modular view says nothing that could imply that the face module must be silent when nonface stimuli are presented. Nor is it committed to the claim that the ways in which various nonface stimuli activate the face module must not carry any information about the categories to which those stimuli belong. Perhaps the process by which the FFA recognizes that a shoe is not a face is informatively different from the process by which a bottle is recognized as not a face. The modular view requires only that the activity displayed by the FFA has no role in the subject's seeing that the nonface stimulus belongs to a nonface category. Kanwisher's idea that the FFA and only the FFA is responsible for the subject's awareness of faces is compatible with Haxby's findings that the FFA is not the only structure whose activation carries information about faces and that the FFA's activation carries information about nonfaces.

'Representation' can, however, be understood in a stronger sense, so that the discovery of distributed and overlapping representations of faces and objects *would* contradict Kanwisher's modularity claim. Representations in the first sense come too cheaply: They need not play any role in psychological processing and so the discovery of face representations outside the FFA does not tell us whether any psychological processing of faces takes place outside the FFA. If we want the claim that there are distributed and overlapping representations to be a claim that does contradict Kanwisher's hypothesis then we must understand 'representation' in such a way that, in order to count as a representation, a state must play a role in psychological processing. On this more demanding notion of 'representation', a pattern of neural firing represents *x* for a subject if and only if the subject's cognitive architecture is such that it is in virtue of the presence of this pattern of neural firing that there is a mental state of the subject the content of which is '. . . *x* . . .'.

Understood in this more demanding way the claim that there are distributed and overlapping representations of faces and objects does contradict Kanwisher's hypothesis: The resources involved in processing faces are *not* special if the representations in virtue of which the subject sees faces are handled by the very same neural resources that are responsible for handling the representations of other stimuli.

But Haxby's results enable us to locate representations only in the first sense. They do not tell us what determines the content of the subjects' mental states and therefore tell us nothing about the location or distributedness of representations in the second sense of 'representation'. And yet, it is only when they concern representations, as understood in this second sense, that claims about distributed and overlapping representations have substantive consequences about the specialness or otherwise of the cognitive processing of faces.

6. Conclusion

The results of experiments from brain scanning experiments have not provided us with an answer to our long-standing question about whether the cognitive processing of faces is special, when compared with the processing of other visual stimuli. The null hypothesis, according to which the cognitive handling of faces employs the very same forms of processing as are involved in the handling of other stimuli is not refuted by the discovery that there is a region of the brain that works harder in response to faces than it does in response to other stimuli. Nor is the positive hypothesis, according to which there *is* a face-specific module, refuted by the finding that information about various stimuli is carried by the activation of distributed and overlapping brain regions. Claims about the usefulness of brain scanning for resolving the long-standing questions in cognitive psychology remain to be proven. They are not supported by the large number of scanning studies that have been devoted to the examination of face processing.

References

Bodamer, J. (1947). Die Prosop-Agnosie: Die Agnosie des Physiognomieerkennens. *European Archives of Psychiatry and Clinical Neuroscience*, *179*(1–2), 6–53.

Bruyer, R., Laterre, C., Seron, X., Feyereisen, P., Strypstein, E., Pierrand, E., & Rectem, D. (1983). A case of prosopagnosia with some preserved covert remembrance of familiar faces. *Brain and Cognition*, *2*(3), 257–284.

Ellis, A. W., & Young, A. W. (1997). *Human cognitive neuropsychology: A textbook with readings*. Hove, England: Psychology Press.

Fagan, J. F. (1972). Infant's recognition memory for faces. *Journal of Experimental Child Psychology*, *14*(3), 453–476.

Fantz, R. L. (1963). Pattern vision in newborn infants. *Science*, *140*, 296–297.

Farah, M. J., Wilson, K. D., Drain, M., & Tanaka, J. N. (1998). What is 'special' about face perception? *Psychological Review*, *105*, 482–498.

Gauthier, I., Tarr, M. J., Andersen, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform "face area" increases with expertise in recognizing novel objects. *Nature Neuroscience*, *2*, 568–573.

Gauthier, I., Tarr, M. J., Moylen, J., Andersen, A. W., Skudlarski, P., & Gore, J. C. (2000). Does visual subordinate-level categorization engage the functionally defined fusiform face area? *Cognitive Neuropsychology*, *17*(1/2/3), 143–163.

Grill Spector, K., Knouf, N., & Kanwisher, N. (2004). The fusiform face area subserves face perception not generic within category identification. *Nature Neuroscience*, *7*, 555–562.

Hanson, S. J., Matsukaa, T., & Haxby, J. V. (2004). Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: Is there a "face" area? *NeuroImage*, *23*, 156–166.

Haxby, J., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, *293*, 2424–2430.

Humphreys, G. W., & Riddoch, J. (1987). *To see but not to see: A case study of visual agnosia*. London: Erlbaum.

Johnson, M. H., & Morton, J. (1991). *Biology and cognitive development: The case of face recognition*. Cambridge, UK: Blackwell.

Juola, P., & Plunkett, K. (2000). Why double dissociations don't mean much. In G. Cohen, R. Johnston, & K. Plunkett (Eds.), *Exploring cognition: Damaged brains and neural networks: Readings in cognitive neuropsychology and connectionist modeling* (pp. 319–327). New York: Psychology Press.

Kanwisher, N., & Moscovitch, M. (2000). The cognitive neuroscience of face processing: An introduction. *Cognitive Neuropsychology*, *17*, 1–11.

Kanwisher, N., & Yovel, G. (2006). The fusiform face area: A cortical region specialized for the perception of faces. *Philosophical Transactions of the Royal Society of London B*, *361*, 2109–2128.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience*, *17*, 4302–4311.

McNeil, J. E., & Warrington, E. K. (1993). Prosopagnosia: A face-specific disorder. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, *46A*(1), 1–10.

Nelson, C. A. (2001). The development and neural bases of face recognition. *Infant and Child Development*, *10*, 3–18.

Pallis, C. A. (1955). Impaired identification of faces and places with agnosia for colours. Report of a case due to cerebral embolism. *Journal of Neurology, Neurosurgery & Psychiatry*, *18*, 218–224.

Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, *46A*(2), 225–245.

Tarr, M. J., & Gauthier, I. (2000). FFA: A flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nature Neuroscience*, *3*, 764–769.

Thompson, P. (1980). Margaret Thatcher: A new illusion. *Perception*, *9*(4), 483–484.

Wojcinlik, E., Kanwisher, N., & Driver, J. (1998). Covert visual attention modulates face-specific activity in the human fusiform gyrus: fMRI study. *Journal of Neurophysiology*, *79*, 1574–1578.

Xu, Y. (2005). Revisiting the role of the fusiform face area in visual expertise. *Cerebral Cortex*, *15*, 1234–1242.

Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, *81*, 141–145.