What is special about face recognition? Nineteen experiments on a person with visual object agnosia and dyslexia but normal face recognition. *Morris Moscovitch; Gordon Winocur; Marlene Behrmann.*

Author's Abstract: COPYRIGHT 1997 Massachusetts Institute of Technology

In order to study face recognition in relative isolation from visual processes that may also contribute to object recognition and reading, we investigated CK, a man with normal face recognition but with object agnosia and dyslexia caused by a closed-head injury. We administered recognition tests of upright faces, of family resemblance, of age-transformed faces, of caricatures, of cartoons, of inverted faces, and of face features, of disguised faces, of perceptually degraded faces, of fractured faces, of faces parts, and of faces whose parts were made of objects. We compared CK's performance with that of at least 12 control participants. We found that CK performed as well as controls as long as the face was upright and retained the configurational integrity among the internal facial features, the eyes, nose, and mouth. This held regardless of whether the face was disguised or degraded and whether the face was represented as a photo, a caricature, a cartoon, or a face composed of objects. In the last case, CK perceived the face but, unlike controls, was rarely aware that it was composed of objects. When the face, or just the internal features, were inverted or when the configurational gestalt was broken by fracturing the face or misaligning the top and bottom halves, CK's performance suffered far more than that of controls. We conclude that face recognition normally depends on two systems: (1) a holistic, face-specific system that is dependent on orientation-specific coding of second-order relational features (internal), which is intact in CK and (2) a part-based object-recognition system, which is damaged in CK and which contributes to face recognition when the face stimulus does not satisfy the domain-specific conditions needed to activate the face system.

Full Text: COPYRIGHT 1997 Massachusetts Institute of Technology

INTRODUCTION

Domain specificity of processing mechanisms is a central concern of cognitive neuroscience. This issue is especially relevant to research on the recognition of faces: Does face recognition involve processes and neural mechanisms that are different from those involved in object recognition, or is face recognition mediated by the same pattern-recognition mechanisms and processes as objects? This long-standing debate is not fully resolved, although converging evidence from a number of sources suggests that face recognition is special. In this paper we review briefly evidence and theories on both sides of the debate and present extensive data on face recognition from a single individual, CK, who has associative visual object agnosia and dyslexia (Behrmann, Winocur, & Moscovitch, 1992; Behrmann, Moscovitch, & Winocur, 1994) that is particularly circumscribed. His ability to recognize objects is severely impaired even though his acuity is normal and he can apprehend and draw the various features or components of objects he cannot identify. Among other things, what is remarkable about this individual is that his face recognition appears to be intact. He, therefore, is an ideal test case for

investigating face recognition in its relatively pure form, uncontaminated, for the most part, by processes that might be involved in recognizing objects. In other words, if face recognition is indeed special, as some have claimed, studying this individual provides us with a rare opportunity to determine in what way, if any, it is special. This investigation of face recognition will necessarily have some bearing on the debate concerning domain specificity and modularity of perception (Nachson, 1995). We will address some of the relevant issues as they arise in the individual experiments and in the concluding discussion.

We begin with a brief literature review on the issue of whether face recognition is special and follow that with an overview of current theories of face recognition before presenting CK. Then there are six sections each consisting of a number of experiments on different aspects of face recognition: (1) recognition of normal upright faces, (2) recognition of caricatures, (3) recognition of inverted faces, (4) recognition of perceptually degraded faces, (5) recognition of fractured faces and of face parts, and (6) recognition of faces made of objects. We will demonstrate that when the face is presented upright and intact, CK's recognition is perfectly normal and gives no indication of the profound deficits that are noted under other circumstances, such as when the face is inverted and when it is fractured. We conclude with a discussion of the implications of our findings for theories of face recognition and of modularity.

Are Faces Special?

Neuropsychological studies of patients with brain damage have demonstrated a double dissociation between recognition of faces and objects, indicating that the two processes are distinct (Newcombe, Mehta, & de Haan, 1994). At a neuroanatomical level, prosopagnosia (Bodamer, 1947; Hecaen & Agelergues, 1962), a severe deficit in face recognition, is usually associated with bilateral damage to the inferior aspect of the temporal cortex, in the region of the fusiform gyrus (Benton, 1980; Damasio, Damasio, & Van Hoesen, 1982; Damasio, Tranel, & Damasio 1990; Farah, 1990; Meadows, 1974; Sergent & Signoret, 1992a; Whitely & Warrington, 1977) though unilateral damage to the same region on the right (De Renzi, 1986a, 1986b; De Renzi, Perani, Cartesimo, Silveri, & Fazio, 1994; Landis, Cummings, Christen, Bogen, & Imhoff, 1986; Michel, Poncet, & Signoret, 1989; Tovee & Cohen-Tovee, 1993; Warrington & James, 1967) is sufficient to produce the deficit while sparing to a greater or lesser extent object recognition of equal difficulty. The opposite pattern of deficits, impaired object recognition but with relatively spared face recognition, can be obtained with damage on the inferotemporal cortex on the left, although as with prosopagnosia, visual object agnosia is more commonly associated with bilateral damage (Farah, 1990; Hecaen, Goldbloom, Masure, & Ramier, 1974; McCarthy & Warrington, 1990; Newcombe, Mehta, & de Haan, 1994).

Recent studies using functional neuroimaging and ERPs to faces and objects in normal people have corroborated the evidence based on studies of brain-damaged people by showing that faces and objects differ in both the pattern and location of electrophysiological and blood flow responses they elicit (Allison, Ginter, McCarthy,

Nobre, Puce, Luby, & Spencer, 1994a; Allison, McCarthy, Nobre, Puce, & Belger, 1994b; Bentin, Allison, Puce, Perez, & McCarthy, 1996; Bobes, Valdes-Sosa, & Olivares, 1994). For example, the amplitude of some ERP waveforms is greater and latency is shorter to faces than to objects. In ERP, PET, and MRI studies, faces and objects activate distinct loci in the same general location, usually in the right fusiform gyrus (Allison et al., 1994a, 1994b; Grady, McIntosh, Horwitz, Maisog, Ungerleider, Mentis, Pietrini, Schapiro, & Haxby, 1995; Haxby, Grady, Horwitz, Ungerleider, Mishkin, Carson, Herscovitch, Schapiro, & Rapoport, 1991; Haxby et al., 1994; Kanwisher, Chun, McDermott, & Ledden, 1996; McCarthy, Puce, Gore, & Allison, in press; Puce, Allison, Gore, & McCarthy, 1995; Sergent, Ohta, & MacDonald, 1992), although bilateral activation is observed in some individuals. Further evidence of selectivity comes from animal research. Cells that respond selectively to faces have been identified in monkeys in areas homologous to those that are activated by faces in humans (Bruce, Desimone, & Gross, 1981; Desimone, 1991; Gross, Rocha-Miranda, & Bender, 1972; Gross, Rodman, Gochin, & Colombo, 1993; Gross & Sergent, 1992; Heywood & Cowey, 1992; Logothetis & Sheinberg, 1996; Perrett, Hietanen, Oram, & Benson, 1992; Rolls, 1992; Yamane, Kaji, & Kawano, 1988; Young & Yamane, 1992).

At a functional level, the inversion effect has played an important part in distinguishing processes implicated in recognition of faces and objects: inverting stimuli from their canonical upright orientation impairs recognition of faces more than that of objects. This effect, first noted and exploited by Yin (1969; 1970) in studies of normal and brain-damaged people, is especially diagnostic of the differences between faces and objects. Once faces are inverted, the distinctiveness of face recognition at a perceptual and neural level is also lost in humans (Allison et al., 1994a; Bruce, 1988; Bruce & Humphreys, 1994; Jeffreys, 1993; Tanaka & Farah, 1993; but see Bentin et al., 1996; Wright & Roberts, 1996).

Despite this wealth of evidence, it has been argued that when factors such as task demands, discriminability, familiarity, and categorization are equated, behavioral and neurological differences associated with processing objects and faces are reduced or eliminated (see Bruce & Humphreys, 1994). For example, with respect to discriminability, it has been suggested that visual discrimination involved in face recognition is more demanding than that involved in object discrimination. Another version of the discriminability hypothesis states the opposite: Because we are so experienced at recognizing faces and because faces arguably are the most common complex visual stimulus we encounter, faces, not objects, are the easier of the two to discriminate. Neither version of the hypothesis can survive evidence of double dissociation from the neurological literature. If faces were simply more difficult to distinguish from one another than objects, there should not exist any patient who shows the reverse pattern, yet there is ample evidence of patients with object agnosia with relatively spared face recognition (Farah, 1990). Moreover, among prosopagnosics there should not be any who perform normally on object discrimination tasks that have been equated for difficulty with faces, but as already noted, such cases also exist. At the visual level, inversion does not alter the complexity of facial stimuli, yet not only does it impair face recognition, which is expected, but in doing so it eliminates many of the functional differences

between recognition of faces and objects (Bartlett & Searcy, 1993; Bruce & Humphreys, 1994; Rhodes, 1993; Rhodes, Brake, & Atkinson, 1993; Rhodes & Tremewan, 1994; Tanaka & Farah, 1993). By contrast, inverting objects typically has little effect on recognition, except that latencies may be slowed. If factors related to discriminability were at the root of the difference between faces and objects, one would expect that inversion would exaggerate the difference between them rather than reduce it.

Damasio and his colleagues (1982, 1990; Damasio, 1990) proposed a sophisticated version of the discriminability hypothesis to account for prosopagnosia: It is not a deficit of face recognition as such but a general deficit in discriminating among exemplars within, but not across, categories. We shall refer to this as the individuation hypothesis. Prosopagnosics simply are impaired at making fine discriminations among exemplars within a category, which is what face recognition requires. By contrast, object recognition typically involves discriminating one category of items (e.g., a chair) from another (a table). Consistent with the individuation hypothesis, Damasio et al.'s prosopagnosic patients were also impaired at distinguishing one object from another within categories. Contrary to the individuation hypothesis, however, there are prosopagnosic patients who do not have an associated deficit for within-category object discrimination. Though severely affected at recognizing faces, these people nonetheless can distinguish among different glasses (Farah, Levinson, & Klein, 1995; De Renzi, 1986a, 1986b), cars (Sergent & Signoret, 1992a) and sheep (McNeil & Warrington, 1993). Conversely, some patients have difficulty distinguishing among glasses (Farah, 1995) or cows (Assal, Favre, & Anderes, 1984) without being prosopagnosic.

The discriminability hypothesis has its counterpart in the behavioral domain. For example, a common explanation of the inversion effect, presented in many guises (see Valentine, 1988, 1991), is that identification processes for inverted faces are more complex or more difficult to execute than those for inverted objects, and, once executed, the target face may elude identification. An alternative explanation is that inverted faces do not initially engage mechanisms specialized for dealing with faces but instead are handled first by mechanisms used also to identify objects. As a result, faces that were processed holistically when upright are processed piecemeal when they are inverted (Rhodes, 1993; Tanaka & Farah, 1993; Yin, 1969). The latter process is ill-suited for identifying faces that are represented holistically. Because this controversy will be addressed directly by some of our experiments, we postpone further discussion of this point until we present our findings later in the paper. For the moment, we wish to draw attention to the point that the controversy revolves around the issue of whether face recognition is special or whether it depends on general-purpose visual processes that are also used for objects.

Theories of Face (and Object) Recognition

The question of whether face and object recognition are different is distinct from the question of what underlies that difference. We have already suggested that each may be mediated by different neural substrates. Less certain are the processes involved in face recognition that account for its special status. For those who believe that all stimuli are processed by a general-purpose visual processor, the difference between recognition of

faces and objects is one of degree rather than of kind. Typically, variations of the discriminability hypothesis are advanced to support this view, but as we have seen in the previous section, such hypotheses do not stand up well to the evidence.

Those who hold that face recognition is special have suggested a number of different explanations to account for its distinctiveness. One is that there is a face-specific processor whose domain is defined by facial stimuli. That is, the processor is tuned selectively to faces much in the way that neurons are in the monkeys' temporal lobes. We refer to this as the face-module hypothesis. An alternative explanation is that although a specialized processor exists, its domain is not restricted to faces but rather to all stimuli that can be processed and represented holistically. Among natural stimuli, faces are particularly well-suited to engage this holistic processor, although it is possible that other stimuli can do so too. Object recognition, on the other hand, is primarily analytic or part-based and depends on a processor with different characteristics from the one favored by faces. Insofar as aspects of face recognition are part-based, they, too, engage this processor. By the view entailed in this holistic hypothesis, face recognition is special because it has privileged, but not exclusive, access to a holistic processor. This view also holds that under some conditions, face recognition may depend on the operation of the analytic, part-based processor.

Versions of the Holistic Hypothesis

What precisely is meant by holistic and by part-based processes and representations is currently under discussion in the literature and vigorous investigation in the laboratory. The various proposals bear a family resemblance to one another. Although our experiments were not designed to distinguish among all of them, we will have occasion to refer to them throughout the paper. We, therefore, summarize them here and refer the reader to excellent reviews on face recognition for detailed expositions and critiques (Bruce, 1988; Bruce & Humphreys, 1994; Bruce & Young, 1986; Bruyer, 1986; Ellis, Jeeves, Newcombe, & Young, 1986; Farah, 1990, 1991; Nachson, 1995; Young, 1994).

All proposals that claim that face perception is holistic have in common the idea that the whole of the face, its global structure or gestalt as determined by the spatial relations among its components, is greater than the sum of its parts, the individual features that comprise the face. As applied to faces, this idea is linked to the neurologically based notion that the right hemisphere is holistic and, among other things, specializes in processing faces, and the left hemisphere is analytic and specializes in processing words and nameable objects (for an early version of this idea, see Jackson, 1874, 1915; which later was revived by Levy-Agresti & Sperry, 1968; see Bradshaw & Nettleton, 1981, and Moscovitch, 1979, for a review and critique of some of these ideas; and Rhodes, 1993 and Corballis, 1991, for some recent reviews, refinements, and developments). The difficulty, however, always has been to determine exactly what is considered a gestalt, which relations are crucial, and what constitutes a part. Some early views were that facial features that formed a gestalt would be processed in parallel whereas those that did not would be processed serially (Bradshaw & Wallace, 1971; Bradshaw & Nettleton, 1981). Another was that face recognition was dependent primarily on low but not high spatial

frequencies (Harmon, 1973; Sergem & Hellige, 1986). Both views have been repudiated on empirical and theoretical grounds (see Bruce, 1988; Moscovitch & Radzins, 1987; Moscovitch, Scullion, & Christie, 1976) and have been replaced by more computational theories (Beymer & Poggio, 1996; Hancock, Burton, & Bruce, 1996; Valentine, 1991; Yuille, 1991).

Configural Hypothesis

The more popular views contrast the configural properties of faces with the feature-based or less configural properties of objects. Rhodes (1988; Rhodes, Brake, & Atkinson, 1993) distinguished between first-order features, which are distinct, isolated entities (eyes, nose), and second-order features, which are the relations among first-order features and include their spatial relations and their location with respect to the contours of the face. For Sergent (1984; Takane & Sergent, 1983), configural relations are interactive in the sense that there is informational dependency and mutual influence among the parts of the face.

Second-Order Relational Hypothesis

Carey and Diamond (1994; Diamond & Carey, 1986), on the other hand, have distinguished between two types of relational features. First-order relational features are the spatial relations among parts or isolated features and are sufficient for identifying most objects or at least specifying their category membership. Second-order relational features are "distinctive variations of a shared configuration" (Carey & Diamond, 1994, p. 255), the spatial arrangements of the parts relative to some prototypical arrangement that exists for a class of items. Rhodes refers to stimuli that have this type of prototypical arrangement as homogeneous (Rhodes & McLean, 1990). Because the parts of a face always bear the same relation to each other, faces can be individuated only on the basis of their second-order relational features.

Norm-Based Coding Hypothesis

In a subsequent development of the basic idea of second-order relations, Rhodes, Brennan, & Carey (1987) suggested that the computations that are crucial are not the ones that involve features within a face but rather those that code relations between a face and a norm that is derived by averaging or superimposing a large number of faces (see also Valentine, 1991, for the advocacy of norm-based coding in face recognition). As Rhodes et al. (1993) correctly observed, each of these proposals assumes that representations of faces consist of parsed but interactive features and that the nature of their interactions determine whether the representations are holistic or not. In contrast to the norm-based hypothesis, the density alone (or noise) hypothesis states that individual faces are identified by their overall point-by-point representations in multidimensional space or by a principal component analysis (Valentine, 1991; Johnston, Milne, Williams, & Hosie, 1997).

Gestalt or Template Hypothesis

Another class of theories, however, exists that assumes that holistic representations are unparsed (Farah, 1990; Corballis, 1991; Tanaka & Farah, 1993; see Garner, 1974, Moscovitch, 1979, and Bradshaw & Nettleton, 1981, for review of some of the earlier ideas). Instead, they are represented as gestalts or templates in which the component parts, although separable in principle, are not processed or coded independently. Their identity depends on the gestalt of which they are a part. Insofar as faces are special, they are holistic in this sense. Part-based representations depend on decomposing items into their component parts and then integrating those parts in relation to each other and to the shape that binds them. Typically, object recognition is based on part-based representations, although some aspects of face recognition, such as recognizing inverted faces, can also be part-based.

Although Farah (1995; see critique in Carey & Diamond, 1994) attempts to distinguish this theory from some of the configural, relational, and norm-based theories, we do not think that the attempt is wholly successful. Instead, we believe that the two types of theories complement each other. The configural theories indicate the relation that parts forming a facial gestalt must have with each other, and with a norm, so that the face can be recognized. As we understand them, Farah's theory says nothing about the algorithm relating the parts to each other, and configural theories are neutral as to the ability of the parts to have an identity independent of the whole.

Each of these proposals captures some of the crucial differences between faces and other objects, and each has been useful in explaining different aspects of face perception. Nonetheless, many problems remain, which suggests either that the proposals are deficient or the ways they have been tested are not adequate. For example, in a number of studies, inversion either does not affect what is construed on principled grounds to be relational or norm-based processing or it paradoxically affects recognition of face parts. Thus, inversion does not affect recognition of caricatures, which is believed to involve norm-based coding more than recognition of veridical drawings (Rhodes et al., 1993). On the other hand, inversion does affect recognition of features such as eyes and mouth when presented in the context of a face but not when presented in isolation. As well, markers of what is taken to be a common underlying relational process seem to have different developmental time courses. Thus, the full effects of inversion are not observed until 10 years of age, whereas interference effects caused by combining seamlessly the upper and lower parts of two different faces (Young, Hellawell, & Hay, 1987) emerges as early as age 3 (Carey & Diamond, 1994). These results suggest one of the following: that there are different types of relational and norm-based processes each of which contribute to face recognition and make it special (Rhodes et at., 1993), that there is an inherent ambiguity in classifying aspects of the face as parts or relations (Rhodes et al., 1993), or that the theories themselves are inadequate. We hope that our study will provide some direction to resolving some of the recurring problems in the area and adjudicating among the various theories.

Rationale for Studying Face Recognition in People with Object Agnosia

One possible source of the controversy about the special status of face recognition, and the inherent difficulty of pinning it down, may be that in neurologically intact people, recognition of faces (and of objects) involves both holistic and part-based processes, making it difficult to identify the unique contribution of each. This is why investigating the face-recognition abilities of our patient is particularly valuable: Because his object-recognition system is damaged or absent, our investigation allows us to assess the limits of face recognition when processing relies primarily on the specialized face system. The results of the investigation can then be used to describe some of the properties of this processor and constrain our ideas about some others.

To our knowledge, a detailed examination of face recognition in a person who is agnosic for objects has been undertaken only once (McMullen, Fisk, & Phillips, submitted) and even in that study only a few tests of face recognition were administered. In studies of patients with object agnosia, face recognition is of peripheral interest: If it is examined at all, it is only to determine the specificity of the object agnosia (Feinberg, Schindler, Ochoa, Kwan, & Farah, 1994; see Farah, 1990, for review). Following the time-honored approach of understanding the normal by focusing on the abnormal, most neuropsychological investigations of face recognition have focused on prosopagnosia. This strategy is useful insofar as it provides information on how face recognition can be distorted when the mechanisms necessary for its normal function are damaged or absent (Moscovitch & Umilta, 1990). Such reports are especially informative if the individual's object recognition is intact because then we learn what the limits of face recognition are when it depends only on object recognition mechanisms. Often, however, the individual also is impaired at recognizing objects, and even when she or he is not, little is made of the theoretical significance of putting an object-recognition system at the service of face recognition.

To obtain a fuller understanding of processes involved in face recognition, studies of both types of agnosic patients are necessary. As noted, until now the focus has been predominantly on prosopagnosia. The experiments we present make a strong case for investigating face recognition in object agnosia and, by extension, in other types of patients with perceptual disorders whose face recognition, on the surface, appears to be intact. If this strategy proves useful for face recognition, a strong case can be made for applying it to other domains.

Mr. CK

CK was described extensively in a previous report (Behrmann et al., 1994), so only a brief summary will be provided here. CK, a right-handed man, was born in 1961, and emigrated from England to Canada in 1980 where he got married and was working toward a postgraduate degree in history. In January 1988, while jogging, he sustained a closed-head injury in a motor-vehicle accident. Initially in a coma, and having motor, sensory, cognitive, and emotional deficits shortly after he emerged from it, CK made a substantial recovery. A full neuropsychological investigation in 1991 revealed a verbal IQ of 96 and a performance IQ of 76, which likely underestimates his intelligence when considered in light of his severe object agnosia. In addition to the object agnosia and

dyslexia, which were well-documented, CK also had residual blindness in his upper left visual field and some mild left-sided weakness in his limbs that, despite his anti-convulsant medication (Tegretol), sometimes show clonus. Except for a hint of bilateral thinning in the occipitotemporal region, no damage was revealed on MRI or CT scans. His visual acuity is normal as are his language, memory, and reasoning. In 1991 he completed his MA degree with the aid of multitrack tape recorders and a voice-activated computer. CK is now in a manager-training program at a large organization. After an initial period of adjustment and despite his deficits, he has adapted quite well. This is testimony not only to his intelligence but also to his ambition and perseverance.

CK is introspective and insightful about his deficit. He willingly and spontaneously shares his insights with us and we provide him with feedback about his performance. We all believe that the more we know about his disorder, the easier it is for him to manage it and to explain his deficits to others who, he tells us, cannot comprehend why someone who is not blind and can see faces so well is so deficient visually in other domains.

It is worth providing some of our subjective impressions of CK's deficits to put on record what we have stated informally when answering questions about CK at colloquia and meetings. CK can navigate well in the world. He does not bump into objects and given proper contextual cues can infer what many objects in the environment are from their separate components. If he were to see a sheet of lined paper and a yellow pencil on a desk, it is likely that he would be able to identify them, by inference, rather than by immediate perception. When faced with objects that are not specified by the context, he may err in identifying them. Thus, on one occasion, he identified a pen placed in a holder fixed on a marble stand as a "trophy you must have won for your research" and had no idea that it was a pen until he touched it. On another occasion one of us brought him a cup of coffee that he had requested and placed it on the desk. Because some time passed without his drinking the coffee, CK was asked whether he no longer wanted it. He answered that he hadn't drunk it because he could not locate it since he had difficulty distinguishing the styrofoam coffee cup from other containers on the desk.

He identifies objects by noting their separate components, and because he knows we are interested in his reactions, he often provides a running commentary on how he pieces together these components using perception and reason to arrive at an answer. In this way, he told us that the object on the table was a "trophy" because he noted what appeared to be a stand with an object embedded in it. He also can copy the objects he cannot identify. His copies, however, are painstakingly piecemeal, focusing on separate parts without any appreciation of the whole. Having studied drawing as a young man, he can draw well from memory because the internal representation of objects is intact, as our studies of his visual imagery confirm (see Behrmann et al., 1992, 1994). Nonetheless, he claims that he has to cover up his own drawings so as to expose only a small portion at a time; otherwise the input he receives confuses him because, he says, it does not coincide with the image he has in his mind's eye. We believe, however, that he has some appreciation for the overall shape of the objects he copies or draws from memory because the parts are not randomly scattered or connected inappropriately but rather retain their proper relation with one another.

When he finds himself in an unusual or visually complex and unfamiliar environment, he often has a vacant look in his eyes, mixed with some concern or anxiety. For example, when we entered a cafeteria for lunch during a break in testing, his animated conversation ended abruptly as he looked about and realized that he did not have a clue as to the choice of foods available to him. Everything, he said, looked like different colored blobs, and he asked that they be identified for him so that he could choose his meal. He was able to eat appropriately (though we cannot report whether he actually knew which of the foods on his plate entered his mouth before he tasted it), and he regained his bright expression and engaging demeanor once he could focus on faces rather than on food.

His deficit in recognizing nonface objects also extends to those with which he is highly familiar, some since childhood. He was an airplane enthusiast and claims to have been able to recognize most of the planes in Jane's books on aircraft. He lamented that now he could not recognize any with certainty, a fact confirmed by us when we tested him formally - he scored at chance. He also has a large collection of thousands of toy soldiers that he had collected since childhood and that he now wants to sell because they no longer give him any pleasure. Where formerly he could distinguish an Assyrian foot soldier from a Roman one, and the latter from a Greek, he now could do so only by touch and often, because the soldiers were so small, the distinguishing marks among armies, let alone soldiers within the same army, were too small to sense easily.

In his personal life at least (we have no formal data on this), the deficit extends to parts of the human body. If a body part, such as a foot, protrudes from under a cover, he sometimes misidentifies it and may treat it as an inanimate object. Thankfully, the other senses can compensate for his peculiar, and sad, visual deficit.

The Nature of CK's Visual Object Agnosia

Because he can copy figures, draw objects from memory, and image objects well and because his visual acuity and face recognition appears to be intact, CK's condition is diagnosed as associative visual object agnosia - he cannot derive the associations or assign meaning related to the stimulus input that he can receive. Because he can identify the component parts of objects but not put them together into a coherent whole, his particular form of associative agnosia is called integrative agnosia (Riddoch & Humphreys, 1987). He seems to be unable to segment and group the elements that comprise objects and relate them to each other or to the overall shape of the object. He suffers from what Farah would call an impairment in part-decomposition (and, we would emphasize, synthesis). It is for this reason we thought that studying his ability to recognize faces would tell us what aspects of face recognition can proceed normally without part-decomposition and synthesis and what aspects cannot.

GENERAL METHOD AND TERMINOLOGY

Each of the experiments followed the same general procedure. CK was tested individually, as were the control participants. Unless there were published norms for the tests we administered, usually 12 participants, 6 males and 6 females, served as controls.

They were matched for age and education with CK. We took CK's performance to be deficient if it fell two or more standard deviations (SDs) below the mean of the controls. The procedure peculiar to each experiment will be described in the corresponding sections. Because there are so many experiments, a short conclusion is included at the end of each.

Because one of the issues this paper addresses is the modularity of face recognition, we did not want to prejudge the issue by referring to the mechanisms involved in face and object recognition as face and object modules, respectively, and the processes they mediate as face-specific and object-specific processes. Nonetheless, it is difficult to escape such terminology entirely given the nature of CK's preserved and impaired abilities. To use the general terms visual recognition mechanisms and processes is to err on the other side. We chose, instead, to use the terms face and object recognition systems and processes. These terms do not necessarily commit one to the view that the systems are modular or even that each system cannot be used to process information in the alternate domain if the stimuli and task demands permit. In short, the terms are as neutral as they can be while still allowing us to acknowledge a difference between those mechanisms typically used to process faces from those typically used to process objects. In the course of the paper, we specify which face stimuli and tasks engage each of these mechanisms and try to answer the question: What is special about face recognition?

How Good is CK at Recognizing, Matching, and Remembering Faces? Experiments 1 through 5

Experiment 1: Recognizing Photographs of Famous People

When we first examined CK, we showed him 17 photographs of famous people, many of them in nonprototypical poses, and he was able to identify them all (Behrmann et al., 1992, 1994). His performance surpassed the mean of our normal control group. To determine whether his face recognition is indeed normal or superior, we presented him and control participants with a set of 140 different photographs, most of them in color. A response was considered correct if the participant supplied the name or some identifying information, such as the President of France or the star of a particular television program.

Results and Discussion. Table 1 shows that CK's recognition of faces was equal to that of the normal control group. The slightly superior performance we sometimes observed probably was due to selection of faces that were especially familiar to CK and to the inclusion of one control who was an outlier.

The results from this study confirms that CK's recognition deficits, which are patently obvious and severe with regard to objects and words, do not extend to faces. His face recognition is not even mildly impaired as it often is in many patients with visual object agnosia (Farah, 1990). In point of fact, we do not know precisely how well CK compares with agnosic patients whose face recognition is reported to be intact because, as noted, such extensive tests of face recognition have rarely been undertaken. McMullen et al.'s

(submitted) patient has an apperceptive agnosia. As a result, his face recognition, well-preserved relative to object recognition, nonetheless may be worse than normal.

Conclusion. Recognition of famous faces is normal in CK despite his severe visual object agnosia.

Experiment 2: Recognizing Atypical Photos of Famous People

Normal face recognition often includes the ability to recognize individuals as they change with age, although admittedly our accuracy diminishes the more the person's face is altered from the typical image we retain. We wished to know whether CK's face recognition was normal only when confronted with fairly typical photos of famous people or whether he, too, could recognize faces that have been transformed by age and changes in style of dress and coiffure.

Method. To investigate this issue, we obtained a book of photographs of famous people at different ages. Prior to testing, we had other individuals order the faces in terms of how closely they resembled the target face that was judged to be the most typical or well known. CK and his controls then were shown the photos in sequence from the least similar to the target to the most similar, with the target being shown last ([ILLUSTRATION FOR FIGURE 1 OMITTED] for an example).

Table 1. Mean Number of Famous People Recognized from Photos.

		Mean	SD	Range
Set A (n = 12)	Controls	54	12	34-69
	CK	66		
Set B (n = 12)	Controls	54	13	33-68
	CK	53		

Results and Comment, Table 2 shows that CK performed as well as the controls on this test. He not only was able to identify target faces as well or better than controls, confirming the results of Experiment 1, but he was as good as they were at extrapolating from the typical image of that face to photos in which the face was transformed by age.

Conclusion. Recognition of age-transformed faces of famous people is normal in CK.

Table 2. Mean Scores on Recognition Test of Atypical Photos of Famous People Across Their Lifespan (Max. = 50).

	Mean	SD	Range
Controls (n = 13)	30	10	14-45
CK	32		

Experiment 3: Judging Family Resemblance

Judging family resemblance is said to be possible, although the evidence on this issue is weak. We were not so much interested in knowing whether CK could judge true family

resemblance among biologically related people, but whether he could choose faces that were created by combining evenly the features of two other individuals, one male, one female. Put another way, could he pick the "parents" of computer-created "offspring" faces? This question was of interest to us because it would provide information about CK's discrimination abilities as well as his ability to focus on facial features, which is what seems to be required to perform this task adequately [ILLUSTRATION FOR FIGURE 2 OMITTED]. There were 49 offspring created from 14 parents. One point was awarded for each correctly identified parent. This experiment is also of interest because good performance depends on the ability to compare facial features embedded in different configurations. As such, the results of this study could tell us something about the nature of holistic and part-based processes in face recognition.

Results and Comment. CK performed as well as controls on this test (see Table 3), indicating that his ability to make this difficult discrimination has not suffered. It also suggests that he can identify facial features or components quite well. Whatever holistic process is involved in face recognition, it does not preclude identification of facial components even when more seems to be involved than simply matching one identical face with another. An alternate interpretation is that the task requires only simple pattern matching between a feature, such as the lips of the offspring, and the corresponding feature on the parent's face. Even if the second alternative were true, it would still argue against the strong form of the holistic theory that states that perception of the part is influenced by the whole (Farah, 1995). From that perspective, because CK has an impaired part-based system, it should have been more difficult for him than for controls to escape the configural influences and perceive the parts well enough to achieve reasonably high levels of accuracy on this test. That he was no worse than controls suggests that an intact part-based system is not needed for this task. We will have more to say about the effects of facial configuration or the perception of facial features after presenting the results of Experiment 17.

```
Table 3. Mean Number of Correct Responses on a Seven-Item Forced
Choice Recognition Test for Identifying Each of the "Parent" Photos
from the Photo of a "Child."

                        Mean     SD      Range

Controls (n = 16)        79       8       62-93
CK                       78
```

Conclusion. Appreciation of family resemblance is normal in CK and suggests that he can recognize parts of faces even when they are embedded in different facial configurations.

Experiment 4: Matching Faces from Different Views and Under Different Lighting Conditions (from Behrmann et al., 1994)

In each of the previous experiments, CK had either to identify familiar faces or he had to judge facial similarity when the facial components were clearly visible and could be matched by superposition, in principle, if not in practice. We wished to know whether

CK could match unfamiliar faces even when he could not rely on a simple pattern-matching strategy or on previously stored representations of faces. To this end, we administered The Face Recognition test, a standardized test developed by Benton and Van Allen (1973; Benton, Hamsher, Varney, & Spreen, 1978) for examining face recognition in brain-damaged individuals. On each of 48 trials of this test, a participant chooses which of six sample faces matches a target face, even though on most of the trials the samples are oriented or lit differently from the target.

Results and Comment. As Table 4 shows, there was no significant difference between CK's performance and that of neurologically intact men. CK's performance on this test reveals that he can discriminate and match faces across different views and lighting without relying on previous experience with the face. Simple pattern or feature matching would not support this performance. As his perception in real life would indicate, CK indeed seems to represent faces normally and uses that information in discriminating one face from another.

Table 4. Scaled Score on the Face Recognition Test (Benton et al., 1978): Matching Faces Photographed from Different Views and under Different Lighting Conditions (Max. = 54).

	Scaled Score
Mean standardized score for men	45.6
CK	49

Conclusion. CK is normal at matching unfamiliar faces from different viewpoints and lighting conditions, indicating that he does not rely on memory or sensory feature matching in face perception.

Experiment 5: Memory for Unfamiliar Faces

Although CK could identify familiar faces as well, or better, than controls, we had only anecdotal evidence, and some weak experimental support, that he could learn to identify unfamiliar faces normally. He learned to identify our faces well and on one occasion, in a shower at the university gym, he recognized one of us despite the change in setting and apparel that he was accustomed to seeing, and before that individual recognized CK. He never complained about any difficulties in learning people's faces and their names. In examining the famous faces we asked him to identify, there was no difference in his ability to recognize people who became famous before the accident as compared to those who became famous afterward, although admittedly the former comprised the large majority of our sample.

To determine whether he could learn new faces, we simply exposed him to 16 photographs of faces taken from a yearbook and tested his recognition for them immediately and after a 20-minute delay. His performance was compared with that of 18 young and 18 old adults, evenly divided between males and females.

Results and Comment. Table 5 shows that CK performed no differently from either young or old controls on this task. His ability to recognize faces is not confined to faces that were familiar to him before his accident but extends to new faces, a result that we expected in view of how easily he learned to identify us. What this experiment shows beyond what we know anecdotally is that his acquisition and retention are normal within the delay period sampled by our study, and we have no reason to doubt that it is normal for more extensive delays.

Conclusion. Immediate and delayed memory for unfamiliar faces is normal in CK.

```
Table 5. Mean Number of Unfamiliar Faces Recognized When Tested
Immediately after Study and after a 20-Min Delay (Max. = 16).
```

	Immediate		Delay	
	Mean	S.D.	Mean	S.D.
Young controls (n = 18)	13.2	1.5	11.1	1.6
Old controls (n = 18)	11.6	1.8	10.2	1.8
CK	12.3		11.7	

Discussion of Experiments 1 through 5

These experiments confirm beyond a doubt that CK's recognition of faces is normal under a wide range of conditions and testing procedures. He can recognize famous people as well as controls do even from atypical photographs that depict the individual from epochs of life in which they were not famous. CK can detect family resemblance among individuals and can identify the same individual from photographs that alter viewpoint and lighting. He also can learn to identify new faces and remember them normally.

From these data it would appear that the visual impairment that is responsible for his object agnosia does not extend to faces. We chose these particular experiments to document his preserved abilities because we believe they represent a fair sample of the variety of face-recognition abilities typically demanded of us in real life. Although in all our experiments we used photographs rather than real faces, from our experience with CK and from his own and his wife's reports, we do not think that our findings misrepresent his real-life abilities.

In the next set of experiments we examined whether CK's face-recognition abilities extend to caricatures, stimuli that resemble faces and, when well-executed, clearly depict individuals we recognize. Caricatures are at once impoverished representations of faces, lacking the detail of photographs, and "deep" in the sense that they capture the essence of the individual face by exaggerating its distinguishing features. Would such impoverished stimuli be sufficient to engage CK's face-processing system and yield an output specific to the individual being depicted or would caricature identification be dependent on feature-identification processes that typically are associated with the object-recognition system that is impaired in CK?

Recognizing Caricatures: Experiments 6 through 8

Recognizing familiar faces seems to depend on encoding distinctive characteristics that distinguish one face from another. Thus, distinctive faces are remembered better than faces judged to be less distinct (Bartlett, Hurry, & Thorley, 1984; Benson & Perrett, 1994; Cohen & Carr, 1975; Going & Read, 1974; Light, Kayra-Stuart, & Hollander, 1979; Valentine & Bruce, 1986a, 1986b). These distinctive properties may even be exaggerated in memory, effectively making our memory of a face not unlike that of a caricature. Early studies comparing learning and memory of photographs with caricatures found no benefit for one over the other as this caricature hypothesis would suggest (e.g., Tversky & Baratz, 1985). Such results, however, are difficult to interpret. As Rhodes, Brennan, and Carey (1987) suggested, the more detailed information about a face that a photo contains in comparison to a caricature may compensate for whatever the photo loses in distinctiveness. A better test, Rhodes et al. proposed, would be to vary distinctiveness in a single representational medium. They did just that by using Brennan's (1985) program for generating line drawings of individual faces in which distinctive features can be identified and varied continuously. They created caricatures of people with whom their participants were acquainted that exaggerated distinctive features and anti-caricatures that distorted the distinctive features to an equal extent but in a direction that brought them closer to the norm, i.e., made them less distinctive. In a later study, Carey, Rhodes, Diamond, and Hamilton (submitted) did the same thing with photos of famous people. As predicted, participants identified the caricatures as well or better than the veridical line drawings of the face and much better than the anti-caricatures. Benson and Perret (1991) obtained the same result when comparing photographic-quality caricatures to veridical photos.

The question we wished to ask was whether caricature recognition, i.e., identifying, encoding, and retrieving these salient features, involves the same processes as does recognition of faces in real life or at least from photos. It may be the case that object-recognition mechanisms that are involved in isolating and identifying parts (part-decomposition and synthesis) are necessary for recognizing caricatures, but another type of process is involved in recognizing real faces. That some procedures, such as inversion, reduce face recognition but not the caricature effect (Rhodes & Tremewan, 1994), indicates that the hypothesis is worth considering. In neurologically intact people it is difficult to tease the two apart. As a test of the hypothesis, we wished to know whether CK, whose object-recognition system is impaired, would recognize caricatures normally. In Experiment 6, we had him try to recognize caricatures of famous people taken from magazines and newspapers, in Experiment 7 we had him try to recognize cartoons, and in Experiment 8 we tested the caricature hypothesis directly by using the stimulus set of caricatures of famous people that Carey et al. (submitted) created and kindly lent us. The latter experiment also provided an opportunity to test Rhodes et al.'s norm-based hypothesis of caricature and face recognition against the density-alone or noise hypothesis.

Experiment 6: Recognizing Detailed and Impoverished Caricatures of Famous People

We chose two sets of caricatures: one that preserved a lot of detail and included shading (such as Levine's creations) and another set that relied more on the line (such as

Hirschfeld's) [ILLUSTRATION FOR FIGURE 3 OMITTED]. We wished to know whether CK would find one more difficult than the other, on the assumption that the detailed caricature resembled a face more and would be more likely to engage the face mechanisms that were intact.

Results and Comment. Table 6 shows that CK was able to identify caricatures as well as controls. There was no difference between the detailed and minimal sample. Clearly, the face recognition system does not need a real face or a very good facsimile of one, such as a photograph, to be activated. Even an impoverished caricature, which no one could mistake for a real face, will do.

```
Table 6. Mean Recognition and Percentage of Correct Identification
of Caricatures of Famous People Whom the Participants Recognized
(Max. = 29).
```

	Mean Recognition	Percent Identified	S.D.	Range
Controls (n = 12)	28.3	80.4	13.3	50-93
CK	29	97		

Experiment 7: Recognizing Cartoon Faces of Imaginary Characters

Although caricatures may never pass for real human faces, they nonetheless represent them. It is reasonable to suppose that if face recognition mechanisms exist, it is human faces they are coding. If CK can identify caricatures, it is because caricatures reactivate internal representation of veridical human faces that he once experienced.

It is of interest, therefore, to know whether CK could also recognize cartoon faces of characters he never could have experienced as having a real human face, characters such as Mickey Mouse, Goofy, and Pluto, or pictures of faces of characters such as Grover and Big Bird from Sesame Street or Kermit from The Muppets. If CK can identify them, it suggests that caricatures, as well as cartoons, are processed in their own right by face mechanisms rather than only through their association with a veridical face [ILLUSTRATION FOR FIGURE 6 OMITTED].

The other reason we were interested in this question is that it would provide some idea of the minimal information that the face recognition system needs in order to be engaged. Typically cartoons present very limited information to be identified as faces; eyes, nose, and mouth are essentially the only properties that are humanlike. Yet we all see these cartoons as faces. We do not know, however, whether we rely on processes involved in object recognition to identify these faces as compared to photos and caricatures of real faces where such processes are not needed. Testing CK should provide a solution to this problem.

We asked CK and controls to identify 31 cartoon faces.

Results and Comment. CK performed somewhat better than controls on this task (see Table 7). Although the cartoons could never be taken for real faces, CK recognized them

almost flawlessly. Because his object recognition is severely impaired, it is unlikely that object recognition processes contributed to his performance. The most likely explanation is that he recognized cartoon faces by recourse to his face-recognition system, which is intact and fully engaged by the cartoons.

These results aid in the interpretation of the previous experiment; they indicate that normal processing of caricatures of human faces does not rely on activation of a previously stored representation of the veridical face that the caricature depicts - although clearly such an association is necessary if the caricature is to be recognized. Rather, despite their impoverished depiction and exaggeration of features often beyond what can be expected of normal faces, caricatures can activate face recognition mechanisms. That the mechanisms can also be engaged by cartoons whose faces bear an even poorer resemblance to real faces than do caricatures indicates that the type of information needed to activate the mechanisms is minimal indeed. At this point we can hazard the guess that the face recognition system will be activated by any stimulus configuration in which eyes, nose, and mouth are identifiable and in proper relation to one another. Because some of the cartoons lacked a proper nose or mouth (e.g., Big Bird), it may be necessary that only two of the features be present in the proper gestalt. Later experiments will elaborate on this point. A recent study by Suzuki and Cavanagh (1995) suggests that simply having two curved lines at the location of the eyes and one at the location of the mouth is sufficient to form a gestalt. They found that detection of the curvature of only a single line is more difficult when it appears in the context of the face than when it appears within a nonface aggregate of lines.

```
Table 7. Mean Number of Upright Cartoon Faces Identified Correctly
(Max. = 31).

                      Mean     S.D.     Range

Controls (n = 12)     28.0     2.6      23-31
CK                    30
```

Conclusion. Recognition of cartoons is normal in CK, indicating that even rudimentary faces of imaginary creatures can activate a face-recognition system.

Experiment 8: Norm-Based Coding in Recognition of Veridical, Caricature, and Anti-Caricature Drawings of Faces

According to the norm-based coding hypothesis, faces are recognized by a process that codes the deviation of facial features from a prototype or norm. A face is coded according to the deviation of its features along a vector that originates in the norm and passes through the region of deviation. What makes some features of a face salient is the extent of their deviation along these "psychologically privileged" vectors from the norm. Caricatures produce their effect by further increasing the features' distance from the norm in the same direction, i.e., along the privileged vector, and in the same proportion in relation to other features [ILLUSTRATION FOR FIGURE 4A AND B OMITTED]. It is for this reason that successful caricatures appear both grotesque and yet startlingly recognizable. If features are altered along the same vectors by an equal amount as in

caricatures but in the opposite direction, i.e., toward the norm, anti-caricatures are produced [ILLUSTRATION FOR FIGURE 4C OMITTED]. As a result, faces are less distinguishable from the norm and consequently should become less recognizable. As predicted, Rhodes et al. (1987), Carey (1992), and Carey et al. (submitted) found that for personal acquaintances and for famous people, caricatures were recognized as well as veridical line drawings or better than them, whereas anti-caricatures were recognized far less well.

These results were also consistent with a density-alone hypothesis (or noise hypothesis), (see Bartlett & Searcy, 1993). According to this hypothesis, individual faces are not represented at encoding on the basis of comparison with the norm. Rather they are identified by their overall point-by-point representation in multidimensional space (Valentine, 1988, 1991) or by a principal component analysis that does not involve comparison with the norm at encoding. If the density-alone hypothesis is correct, comparable deviation from the veridical face as in the caricatures but along a direction orthogonal to the privileged vector (lateral caricatures) should produce a face that is as recognizable as that of the caricature [ILLUSTRATION FOR FIGURE 4D OMITTED]. Contrary to the density-alone hypothesis, Rhodes, Carey, and their collaborators found such faces to be virtually unrecognizable because they bore almost no resemblance to the particular faces from which they were derived (Carey, 1992, p. 99; Carey et al., submitted).

If CK showed the same pattern of performance as that of controls, it would provide added evidence that face recognition in CK involves the same processes as in neurologically intact individuals. It would also indicate that processes involved in norm-based coding do not rely on part-based object-recognition systems to specify and isolate the features that enter into the calculation. Norm-based coding can be implemented by the face-recognition system alone.

To test these hypotheses, we asked CK to identify 60 line drawings consisting of veridical representation and a caricature, anti-caricature, and lateral caricature of each of 15 famous people.(1)

Results and Comment. Table 8 shows that the pattern of CK's performance resembled that of the controls, although he performed at a level that was about 50% higher. For both CK and controls, identification of caricatures and veridical drawings was about equivalent and twice as good as that of anti-caricatures. Lateral caricatures were barely recognized by all participants.(1)

These results indicate that processes involved in norm-based coding are intact in CK and suggest strongly that in controls, as in CK, they are implemented by the face-recognition system. Indeed, CK's superior performance overall may imply that norm-based coding may be hampered by interference from competing object-recognition systems in controls. Although provocative and tantalizing, this conclusion is not warranted by the data. CK showed no greater proportional advantage of caricatures over veridical drawings or over anti-caricatures, as that hypothesis would suggest. The most likely explanation is that

many of the famous people were older individuals (e.g., Rodney Dangerfield and George Shultz) who were more likely to be familiar to CK, who was in his mid-30s at the time of testing, than to the young undergraduates in Carey et al.'s (submitted) study.

Despite CK's superior performance overall, he performed as poorly as controls in identifying lateral caricatures. These results reinforce Carey et al.'s (submitted) conclusion that the norm-based hypothesis provides a far better account of face recognition than the density-alone hypothesis. Because versions of the density-alone hypothesis are used in many computer models of face recognition (Beymer & Poggio, 1996; Hancock, Burton, & Bruce, 1996; Turk & Pentland, 1991; see Valentine, 1991), these models, as successful as they are, ultimately [TABULAR DATA FOR TABLE 8 OMITTED] will need to be modified to explain face recognition in people.

Conclusion. CK displays the normal pattern of recognizing veridical line drawings and caricatures better than anti-caricatures and lateral caricatures, thereby indicating that norm-based coding best describes the operation of his isolated face-recognition mechanisms.

Discussion of Experiments 6 through 8

Experiments 6 through 8 demonstrated that processes involved in the recognition of caricatures and cartoons can be normal in an individual who is profoundly agnosic for objects. These results are noteworthy because it is conceivable that identification of cartoons and caricatures is dependent on part-based processes typically used in recognizing objects. Caricatures, and to some extent cartoons, exaggerate particular facial features. As well, cartoons use minimal facial features to make a character immediately recognizable. One could argue that identification of cartoons and caricatures is achieved simply by noting the exaggerated features without regard to the configuration in which they are embedded. That CK recognized caricatures and cartoons normally argues against the interpretation that part-based processes involved in object recognition are needed for identification of these "stylized" faces. Also, like normal controls, CK recognized caricatures formed by exaggerating features that are salient with respect to a norm, but not caricatures in which other features were exaggerated. The latter result indicates that norm-based coding is used to recognize faces even by an individual who has little recourse to object-recognition processes.

Insofar as CK's performance can be taken to reveal the operation of a face-recognition system operating in relative isolation from mechanisms involved in object recognition, the results of Experiments 6 through 8 indicate that the face-recognition system has the following characteristics:

1. It does not need real faces or even photos of real faces for its input; line drawings containing minimal identifying information, which presumably is configural, is sufficient. The precise nature of that minimal information is an issue that will be addressed in later experiments.

2. It does not need input derived from possible real human faces; imaginary cartoon faces will do.

3. It uses a norm-based coding process to identify faces.

Introduction to Experiments 9 through 19

It was important to establish how well-preserved CKs face-recognition abilities are in a variety of circumstances for the reader to appreciate the deficits we will demonstrate in subsequent experiments. First, however, let's consider the significance of our findings to date. We have provided what appears to be the best-documented demonstration of intact face recognition in a patient with visual object agnosia. Had we stopped our experiments here, we would have concluded along with other investigators that face and object recognition are mediated by distinct neural mechanisms. It is not a conclusion we wish to dispute, although we have not provided conclusive evidence to support it. Experiments 1 through 8 also have not helped distinguish among the various theories of face recognition, though Experiment 8 did support the norm-based coding hypothesis over its competitors.

The experiments that follow will demonstrate that under some conditions, CK's face recognition is impaired, often profoundly, indicating that sometimes processes involved in object recognition do contribute to normal face recognition. Despite these findings, we shall argue that they do not invalidate the conclusion that face recognition is distinct from object recognition. Instead, as we shall see, the experiments will help specify the processes involved in recognizing faces and the stimulus characteristics that faces must have to activate these processes; in doing so, the experiments help provide evidence to distinguish among theories of face recognition. Once we appreciate what is special about face recognition, we can then begin to understand the contribution made by processes typically employed in object recognition.

Recognizing Inverted Faces: Experiments 9 through 12

We noted earlier that recognition of faces is believed to be hampered by inversion much more than the recognition of objects. This inversion effect has served as a diagnostic marker for distinguishing between processes involved in face and object recognition since it was first noted by Yin in 1969. Despite almost three decades of research on the topic, it is still being debated whether upright and inverted faces are processed in the same way or whether fundamentally different processes are used in each case (e.g., Bruce & Humphreys, 1994). We refer to the latter hypothesis as the dual mechanism hypothesis. Proponents of the first hypothesis, which simply is a version of the discriminability hypothesis (see "Introduction"), assert that inverted faces are harder to recognize because they are unfamiliar or atypical (Valentine, 1991). As a result, it is difficult for face-recognition processes to assign proper values to the features that norm-based or density coding requires, or, if they do so, the values assigned do not match those of stored representations of upright faces. Those who favor the dual mechanism hypothesis argue that inverted faces cannot engage face-processing mechanisms, and as a result, recognition of inverted faces must rely on processes typically used to recognize objects.

A common version of this hypothesis is that upright faces are recognized holistically whereas inverted faces are recognized piecemeal, by part-based, feature-dependent processes that are also used to recognize objects. Thus, in prosopagnosic patients, recognition of inverted faces is no worse, and sometimes can be even better, than recognition of upright faces (Farah, Wilson, Drain, & Tanaka, 1995).

Because he can recognize upright faces normally but is severely impaired at recognizing objects, CK is an ideal person on whom to test these two hypotheses. According to the discriminability hypothesis, CK should perform no worse than a neurologically intact person in recognizing inverted faces because the processes used to achieve normal recognition of upright faces are the same as those used to recognize inverted faces. If, on the other hand, inverted faces are recognized by recourse to mechanisms that are typically used to recognize objects, CK would be expected to suffer from inversion much more than do normal controls. That is, CK would be disproportionately impaired by inversion.

We begin by reporting CK's ability to identify inverted faces of famous people and well-known cartoon characters. We then report his performance on a perceptual matching task consisting of unfamiliar faces. In each of the three experiments we also tried to control for the possibility that the drop in performance when the faces were inverted was related to the difficulty of the task rather than to inversion per se. In the final experiment, we tried to determine which aspects of the face are most affected by inversion.

Experiment 9: Identification of Inverted and Disguised Faces of Famous People

Despite his normal, and sometimes above average, performance in identifying pictures of upright faces, we were interested to know whether CK would perform similarly on identifying these faces when they were inverted. To control for the expected difficulty that even neurologically intact people would have in identifying inverted faces, we included a condition of disguised faces (with mustaches, glasses, and wigs) that were as difficult for controls to recognize as inverted faces [ILLUSTRATION FOR FIGURE 5 OMITTED]. If CK, unlike controls, finds inverted faces significantly more difficult to recognize than disguised faces, it would constitute evidence against the discriminability hypothesis. Instead, such an outcome would support the dual mechanism hypothesis that inverted faces are recognized by different mechanisms and processes than are used for upright faces, presumably ones that also contribute to object recognition, which is deficient in CK.

CK and control participants attempted to identify faces of 140 famous people in photos, half of which were presented as inverted, and half as disguised.

Results and Comment. These results are presented in Table 9. Because one control participant was much worse than others in identifying faces, we present the results both with that individual's score included and excluded.

Table 9 shows clearly that CK was severely impaired in identifying inverted faces, even though he was as good as or better than normal in identifying disguised and undisguised

upright faces. The results argue against the discriminability hypothesis of the inversion effect and strongly support the dual mechanism hypothesis. Once a face is inverted, normal people rely on mechanisms involved in object recognition to identify the face. Because those mechanisms are damaged in CK, he finds it extremely difficult to recognize inverted faces. When he is successful, he seems to rely on a single, salient identifying feature. When faces are upright, they engage face-recognition mechanisms even when the faces are disguised. Although we do not have a good theory as to why disguised faces are difficult to recognize, the discriminability hypothesis provides a plausible account.

The inversion effect was especially pronounced in CK if the inverted faces were viewed first, suggesting that seeing the upright face may have helped in subsequent identification of the inverted face. That type of facilitation, or priming, is also seen in CK's ability to perform better than controls on identifying disguised faces when they were first seen as undisguised.

Despite our attempt to control for difficulty, it may still be argued that although identifying inverted and disguised faces are equally difficult for controls, the underlying process used to overcome these difficulties may be different. Seeing inverted faces may be so unusual that the procedures used to overcome that distortion are different from those used for seeing through a disguise. After all, we all have encountered people with changed hair styles, new glasses, new headdresses, and altered facial hair and have had some practice dealing with these changes (see Experiment 2), whereas it is much less likely that we would have had much occasion for identifying inverted faces. To address these concerns, we conducted an experiment involving cartoon characters.

Conclusion. CK's severe impairment at recognizing inverted faces despite his normal recognition of upright disguised faces indicates that recognition of inverted faces is mediated by different mechanisms from those used to recognize upright faces; presumably those mechanisms that are necessary for recognizing inverted faces are damaged in CK and contribute to object recognition.

Table 9. Percentage of Correct Recognition of Inverted and Disguised Faces That Were Identified Correctly (see Table 1) When Upright and Undisguised.

	Inverted			Disguised		
	Mean	SD	Range	Mean	SD	Range
Set A						
Controls						
(n = 12)	66	15	31-87	69	13	45-88
(n = 11)	70	11	55-87			
CK	42			85		
Set B						
Controls						
(n = 12)	71	11	42-82	66	16	35-88
(n = 11)	73	6	62-82			
CK	14			68		

Experiment 10: Identification of Inverted Cartoon Faces

We chose to examine the inversion effect in cartoon faces for a number of reasons. Cartoon faces often preserve only the minimal amount of facial information. Often one of the three primary features of the face (nose, eyes, and mouth) is replaced by a part not found in humans such as a beak or a snout. Also, the contour of the face and its peripheral features, such as ears and hair, are far different from those seen in people. No individual, however grotesque, has ears growing on the top of the head like Bugs Bunny or has a head shaped like the Roadrunner's [ILLUSTRATION FOR FIGURE 6 OMITTED]. For this reason, cartoon faces may be identified by part-based processes that focus on these features rather than by holistic processes that are concerned with configurations among features. However, the fact that we perceive cartoon faces as faces and that CK recognizes them normally, suggests that the same processes that are used to recognize human faces are also used to recognize cartoon faces. Our intuitions notwithstanding, it may still be the case that the separate features are so striking that CK would not have to integrate them, as he might for common objects, in order to recognize the cartoon character. Detection of one or two salient features may be sufficient. If this is correct, inverting cartoon faces should have a much less detrimental effect on CK's recognition than inverting human faces. If, on the other hand, CK is as impaired in recognizing inverted cartoon faces as inverted human faces, it would suggest that even cartoon faces must activate his normal face-recognition system. That outcome would provide us with further knowledge about the minimal stimulus attributes necessary for activating that system.

Another reason for testing the inversion effect with cartoon faces is that such faces are trivially easy for normal people to recognize. In fact, performance is virtually identical for the upright and inverted cartoons that we used. As well, unlike humans, cartoon characters are placed in such fantastic situations that their faces are seen as inverted considerably more often than human faces. For these reasons, the discriminability hypothesis predicts that inverting cartoon faces should have a much smaller detrimental effect on CK's recognition than inverting human faces. A contrary outcome would argue strongly against the discriminability hypothesis and for the dual mechanism hypothesis.

Participants attempted to identify 31 inverted cartoon faces.

Results and Comment. CK was profoundly impaired in identifying inverted cartoon faces despite performing flawlessly when identifying the same faces upright (see Tables 7 and 10). Contrary to the discriminability hypothesis, even though upright cartoon faces were easier to identify than faces of famous people, CK was as poor at identifying inverted cartoon faces as he was at inverted human faces and even poorer if we consider the number of standard deviations he fell from the mean of the control group in the two conditions.

These results also indicate that although facial features in cartoons are primitive and sometimes absent, replaced, or inappropriately located, and although the contour of the face does not always conform to that of humans, it is nonetheless treated like a human

face and is capable of engaging the face-recognition system so long as it is presented upright. Once it is inverted, it is no longer recognized by CK, even though the cartoon's salient features provide distinctive cues.

Inverted faces, even ones as rudimentary as cartoons, cannot engage face-recognition mechanisms. That controls, but not CK, recognize inverted faces accurately suggests that they rely on part-based processes used in object recognition, a recourse that is unavailable to CK. Indeed, CK can detect isolated features but only rarely can he use them to identify the character. For example, when viewing the inverted Bugs Bunny, he commented on what he mistook to be long legs attached to the chin, and in examining the inverted Bart Simpson, he wondered why a saw was positioned at the bottom of the face. He did, however, identify Mickey Mouse by his ears. We postpone a more-detailed discussion of the part-based processes necessary for recognition of inverted faces until the end of this section.

```
Table 10. Mean Number of Correctly Identified Inverted Cartoon Faces
(Max. = 31).

                         Mean      SD       Range

Controls (n = 12)       27.0      3.0       22-31
CK 5
```

Conclusion. CK is as impaired at recognizing inverted cartoon faces as inverted faces of people, indicating that recognition of even these easily recognized, rudimentary, inverted faces with salient features is dependent on nonface-recognition processes that are damaged in CK.

Experiment 11: Perceptual Matching of Upright and Inverted Faces of People, Dogs, and Cats

The previous experiments on inversion all required the participants to identify a familiar face from semantic memory. These results may be interpreted as indicating that the inversion effect depends crucially on gaining access to memory representations of faces. A number of recent studies, however, have shown that the inversion effect can be perceptual (Searcy & Bartlett, 1996; Farah et al., submitted). To determine whether the source of CK's difficulty in identifying inverted faces is related to perception or memory, we tested his ability to match unfamiliar faces perceptually. In the same study, we also used heads of real dogs and cats to see if the inversion effect also applies to individuals whose physiognomy bears some resemblance to the human face.

In both the human and animal conditions, participants were required to select the 16 target faces from a set of 32 that appeared either in the same orientation as the targets or opposite from them.

Results and Comment. CK was as profoundly impaired at perceptually matching inverted to upright unfamiliar human faces as he was at identifying inverted familiar faces from memory (see Table 11). Indeed, his score in the inverted-upright condition was slightly

lower than chance. However, he performed nearly flawlessly in both of the other conditions, as did controls in all three conditions.

CK's impaired performance in the upright-inverted condition is strong evidence that the inversion effect observed in the previous experiments is perceptually based. He cannot identify inverted familiar faces for the same reason he cannot match inverted to upright faces: The mechanisms needed to apprehend faces cannot be engaged by inverted faces. As a result, an accurate representation of the inverted face cannot be delivered to memory for matches with stored representations of familiar faces, which this study suggests are also represented in the upright orientation. These results suggest that the face-recognition system includes processes involved in perceptual apprehension of faces as well as those needed to recover stored representations. The controls' performance suggests that perceptual matching of inverted to upright faces depends on processes typically used in object recognition but which can be applied to faces when need be. Bartlett and Searcy (1993; Searcy & Bartlett, 1996) reach the same conclusion based on their perceptual matching studies in normal people. Farah et al. (submitted) showed that interference effects from [TABULAR DATA FOR TABLE 11 OMITTED] face and letter masks were much greater for upright than for inverted faces, leading them to conclude that the inversion effect is perceptually based. Although we believe their conclusion is correct, we think their experiment is flawed because they did not use a counterbalanced design to show that an inverted face mask has little effect on an upright face.

Two other findings are worth noting. CK could match inverted to inverted faces well even though he cannot match inverted to upright faces. This result suggests that he can use simple sensory features to match one picture with another, an interpretation that may also apply to the results of parent-offspring matches in Experiment 2. One reason this strategy fails in the upright-inverted condition is that the faces have first to be normalized to the same orientation for feature matching to take place. Normalization to the same orientation may be an operation associated with object recognition (Jolicoeur, 1985; Hamm & McMullen, 1996) that may not be applicable to faces (Harman & Moscovitch, 1996). Another possibility is that the face-recognition system mandatorily codes faces when they are presented upright (see Farah, Wilson, et al., 1995) but cannot do so when they are inverted. As a result, CK is forced to match two radically different representations: a structural representation of a face when upright with a conglomeration of separate features or parts when inverted. Controls can identify and integrate these parts laboriously so that they can then match inverted to upright faces with a moderate degree of success if they are given enough time. Our prediction is that if they had to match inverted to upright faces quickly, say under masked, tachistoscopic conditions, the failure rate would come to resemble CK's.

Both interpretations receive some support from the second noteworthy finding: CK's performance when matching heads of dogs and cats. He does no worse in the upright-inverted condition than in the upright-upright condition; in both cases he scores lower than controls but much better than he did in the upright-inverted condition for human faces. Although animal heads have some facelike quality, they do not engage the face-recognition system fully as cartoon faces do. As a result, CK probably relies primarily on

sensory feature matches in both cases. His performance is less than perfect, however, because animal heads engage face-recognition mechanisms partially and may sometimes lead him to rely on imperfect information derived from their output. Evidence from subsequent experiments (see Experiment 18) supports this interpretation.

The results of this experiment indicate clearly that the inversion effect can have a perceptual basis. They also indicate that successful matches between inverted and upright faces depend on comparing and coordinating the results of perceptual analyses from mechanisms mediating object recognition with those mediating face recognition. Some of the experiments we report below help identify those processes more precisely.

Conclusion. CK is impaired at matching an inverted human face to a simultaneously presented upright face, indicating that the defect occurs at a perceptual level. Matching of animal heads, although somewhat impaired, is no worse in the upright-inverted condition than in the upright-upright one, and matching inverted-inverted faces is normal, suggesting that matching by simple sensory features is relatively preserved in CK.

Experiment 12: The Effects of Inverting External versus Internal Facial Features on Recognition

A lingering problem for theories of face recognition is to specify which facial components are crucial. Is it just the eyes, nose, and mouth or are the hair and the contour of the face also important? Are each of these components important in isolation or is the spatial relation of one component to the other the crucial factor, as the various relational or holistic theories would predict? If the latter, is the relation of internal features to external contour important or just the relation of internal features to each other? Among the various ways these problems have been approached are by deleting one or more facial features to see how recognition is affected (Ross & Turkewitz, 1981, 1982); by inverting internal facial features, yet still allowing them to occupy their canonical location in the face (experiments on the Thatcher illusion: Thompson, 1980; Bartlett & Searcy, 1993; Rhodes et al., 1993); by altering the spatial relations among the features (Bartlett & Searcy, 1993; Rhodes, 1988; Rhodes et al., 1993; Searcy & Bartlett, 1996; Tanaka & Sengco, in press); by having participants match whole faces that share either internal or external features (Nachson, Moscovitch, & Umilta, 1995; Sergent, 1984); by having participants identify faces when internal and external features are presented in isolation (see references in Nachson et al., 1995); or by multidimensional scaling derived from comparing faces differing in internal and external features (Sergent, 1984; Takane & Sergent, 1983). Although there are some inconsistencies among the results of these studies, in general they have confirmed that both the identity of internal and external features themselves and their relation to one another are important determinants of face recognition. The emphasis placed on the relative importance of one type of information over another will vary with other factors. When matching two simultaneously presented faces, feature identity may be more important than configuration (Searcy & Bartlett, 1996), with external features being favored as familiarity increases (Ellis, Shepherd, & Davies, 1979; Young, Hay, McWeeny, Flude, & Ellis, 1985, but see de Haan & Hay, 1986). When identifying faces or matching them from memory, configurational or

relational information gains in importance, especially as familiarity increases (Nachson et al., 1995; Rhodes, 1988; Rhodes et al., 1993; Ross & Turkewitz, 1982; Searcy & Bartlett, 1996).

Despite these advances, Rhodes et al. (1993) observe correctly that the isolated/relational distinction carries with it an inherent ambiguity that is highlighted by attempts to manipulate one variable independently of the other. When varying isolated features, the relation between them necessarily changes. As well, it is not known whether the spatial distance between features is itself a feature or exists only as configurational information. A second issue, acknowledged by some investigators (Rhodes et al., 1993; Bartlett & Searcy, 1993; Searcy & Bartlett, 1996), is that there may be two mechanisms involved in processing faces, one part-based and the other holistic, each of which contributes differently to performance on the different tests. Thus, finding that both internal and external features can contribute to performance on a sequential face-matching task (Nachson et al., 1995) may simply reflect the separate contribution of each of the two processes and tells us little about which type of information engages processes peculiar to face recognition.

One way in which Rhodes and her colleagues and Bartlett and Searcy address these problems is by studying how other factors, such as inversion, interact with manipulation of isolated features or configurations. Rhodes et al. (1993), Bartlett and Searcy (1993), Searcy and Bartlett (1996), and Tanaka and Sengco (in press) found that inversion has a much greater effect on discriminating changes in spatial relations than in isolated features and, by implication, on identifying faces on the basis of relational information than on the basis of isolated components (but see Rhodes et al., 1993, Experiment 3 for contrary evidence). In these studies, no attempt was made to vary the spatial relations of internal and external features independently of each other. Consequently, we still do not know whether inversion has a differential effect on information derived from these two sources or whether the spatial relations between internal features and external contour are as important for identification as the spatial relations among internal features.

Studying CK provides an excellent opportunity to resolve some of these issues. Because we know that inversion is much more disruptive of CK's performance than that of controls, we can determine whether inversion of internal or external features will cause him more difficulty [ILLUSTRATION FOR FIGURE 7 OMITTED]. If he is unaffected much by either, it would suggest than he can use one source of information independently of the other for identification. If he is equally and severely affected by both, then both sources of information, presumably in relation to one another, are needed for identifying faces. Because CK cannot use his object system to compensate effectively for information that his face mechanisms cannot handle, differences between his performance and that of controls would tell us what type of information engages the face-processing system and what type is handled by other mechanisms that are also used for objects.

In this experiment, participants attempted to identify famous people from photos in which the internal features (eyes, nose, and mouth) were either inverted as a unit or the external features were inverted [ILLUSTRATION FOR FIGURE 7 OMITTED].

Results and Comment. Inverting the internal features impaired everyone's performance but especially CK's (see Table 12). Whereas recognition in controls dropped by about 20%, it dropped by over 60% in CK. Interestingly, inverting only the external features had no noticeable effect on recognition either for CK or for controls. The results confirm what most thieves and highwaymen know: Internal features carry the burden of information in face recognition. For psychologists too lost in theory, our results help explain why the bandanna is worn across the face rather than as a kerchief around the face.

Our results also indicate that for identification, external contour is not nearly as important as internal features, that the spatial relations between internal features and contour are not of vital importance, and that presenting eyes, nose, and mouth in the proper spatial relation is probably sufficient for good recognition. Put succinctly, the face recognition system is sensitive to internal features presented in an upright orientation and in the proper spatial relation to each other. Experiments 16 and 17 examine the effects that altering these spatial relations have on recognition. The far better recognition achieved by controls relative to CK when internal features are inverted suggests that they use other mechanisms, presumably those that can also be used for objects, either to identify the inverted internal portion or to identify the upright contour and the external facial features, such as hair style, that may define a person.

Our results are consistent with previous findings that indicate that people are recognized better on the basis of internal facial features and the spatial relation they bear to each other than on the basis of external features (for references, see "Introduction" to this section). They can also be reconciled with findings showing that on some tests both internal and external features influence recognition because both part-based and holistic processes likely contribute to performance on those tests (e.g., Nachson et al., 1995; Young et al., 1987, Experiment 4), especially if the individual had much practice identifying [TABULAR DATA FOR TABLE 12 OMITTED] the face and its component parts (see also Experiments 18 and 19, below, for demonstrations of competition and interference between holistic encoding of faces and part-based encoding of objects).

Conclusion. Only inversion of internal features (eyes, nose, and mouth) leads to poor recognition in normal people and a profound loss in CK, indicating that face-recognition mechanisms are much more sensitive to internal features (and probably the spatial relations among them) than to external features, facial contour, or the spatial relations between internal and external features.

Discussion of Experiments 9 through 12

The major finding in this set of experiments is that CK was much more than normally impaired at recognizing inverted human faces, even when task or discrimination

difficulty was controlled. Because CK's recognition of all manner of upright faces under a variety of viewing and test conditions is normal, the results argue strongly against the hypothesis that recognition of inverted and upright faces are processed by the same system (Valentine, 1988, 1991). Instead, the results support the hypothesis that inverted faces are processed by separate mechanisms, presumably those mediating part-based processes that are crucial for object recognition and which are damaged in CK.

Each of the experiments in this set was designed not only to compare the discriminability hypothesis with the dual mechanisms hypothesis but also to provide information about the attributes of the stimulus that can engage a face-recognition system. Based on our experiments, we can add the following to points 1,2, and 3 as characteristics of the stimulus properties necessary for activating the face-recognition system:

4. The face cannot deviate too much from the upright orientation (Experiment 9).(2) Photos of real animal heads may weakly engage the mechanism while simultaneously engaging object-recognition mechanisms (Experiment 11). Thus, the inversion effect is less pronounced for them.

5. Within a human face, it is the internal features of the face and the spatial relations among them that carry the most information (Experiment 12). Face-recognition mechanisms are far less sensitive to external features, facial contour, or the spatial relations between internal and external features. Inversion only of internal features leads to poor performance.

6. The face-recognition system includes processes necessary for representing unfamiliar faces perceptually as well as in memory (Experiment 11).

Introduction to Experiments 13 through 17: Parts and Wholes

Throughout the previous set of experiments we speculated that in neurologically intact people, recognition of inverted faces is handled by separate mechanisms employing part-based processes that may also be involved in object recognition (Farah, 1990; Bartlett & Searcy, 1993; Searcy & Bartlett, 1996; Valentine, 1988, 1991). If our speculation is correct, decomposing an upright face into its component parts, altering them, or manipulating the spatial relations among them should have the same deleterious effect on CK's recognition as inversion. As well, by seeing what sort of decomposition or alteration CK can tolerate and what sort he cannot, we can gain a better understanding both of how part-based processes contribute to normal face recognition and of the stimulus attributes necessary for engaging face-recognition mechanisms. On a more theoretical level, these studies may be useful in distinguishing among the various theories of holistic and part-based processes in recognition of faces and objects.

Contour Assignment and Depth Segregation of Perceptually Degraded Faces and Objects: Experiments 13 and 14

CK's visual object agnosia traditionally is considered to be a peripheral agnosia rather than a central agnosia in part because he has preserved internal visual representations of the objects that he cannot perceive.(3) His problem is that he cannot access those representations from vision. A hallmark of peripheral visual object agnosia is that perceptual degradation markedly impairs even the limited performance that such people can attain (De Renzi, Faglioni, Grossi, & Nichelli, 1991; McCarthy & Warrington, 1990). This certainly is true of CK. Requiring him to identify objects in noise, overlapping line drawings of objects, or shaded pictures of objects in unusual lighting has catastrophic effects on his performance (see below and Behrmann et al., 1994).

A possible reason for the devastating effect of these procedures is that they introduce discontinuities in a stimulus, effectively creating disjointed parts where once there was a unified whole. Whereas we can easily recover the whole, CK cannot. The question we wished to ask was whether such procedures would have similar effects on CK's face recognition. If recovering a whole from component parts created by overlap or shadowing depends on the application of a part-based, object-recognition system, CK should perform as poorly with faces as with objects. On the other hand, if identification under such degraded conditions depends on the interaction of top-down processes with bottom-up ones that are domain-specific, face recognition should be spared since both types of processes seem to be intact in CK insofar as faces are concerned.

Two other related questions are also addressed by this study. One concerns the locus of the deficit in integrative visual object agnosia and the other pertains to theories of depth segregation and contour assignment. Because CK's agnosia is considered to be peripheral, one possibility is that the locus of the deficit is at early stages of processing concerned with extraction of information about shape and contour for all stimuli. Consistent with this hypothesis is one class of theories that posits that depth segregation is a low-level perceptual process whose function is the proper assignment of contours to figure and ground (see Nakayama, Shimojo, & Silverman, 1989; Palmer & Rock, 1994a, 1994b). According to this hypothesis, depth segregation is crucial for object identification and must be completed before object identification can occur. If the theory is correct, CK should be as deficient at identifying faces as objects. If, however, CK can identify faces but not objects under conditions in which good performance depends on the proper assignment of foreground and background, the depth-segregation-first hypothesis is challenged. Instead, the results would favor an alternative theory that states that object recognition occurs in parallel with depth segregation and contributes to it (Peterson, 1994a, 1994b; Peterson & Gibson, 1994a, 1994b). The latter results would also suggest that the deficit in integrative agnosia arises beyond regions concerned with depth or figure-ground segregation, extraction of simple features, and shape or contour.

Experiment 13: Depth Segregation and Recognition of Overlapping Faces

A possible way of splitting a face or object into component parts is to introduce discontinuities in the stimulus. This can be accomplished by creating overlapping figures of line drawings [ILLUSTRATION FOR FIGURE 8A OMITTED]. In our experiments on object recognition (Behrmann et al., 1994), we found that although CK was very

impaired at recognizing line drawings of common objects, he nonetheless could identify about 50% of them when they were not overlapping. When presented with overlapping drawings, he could not identify any of them, nor could he choose the targets from among the lures at a level greater than chance. He could not even trace the outline of a single drawing correctly because he could not tell which contour belonged to which object. Reflecting on his own performance, he said that he did not know how to assign parts to each object and so could not use that information to deduce what each object was.

We were interested in knowing whether overlapping line drawings of faces would have a similarly disruptive effect on his performance. If overcoming discontinuities produced by overlap requires integrative, part-based processes used in object recognition, CK should be as impaired in recognizing overlapping faces as objects.

A subsidiary reason for conducting this experiment is that it may serve as a test of some theories of depth segregation and provide information about the locus of CK's deficit. As we noted, one class of theories posits that depth segregation precedes object recognition (Palmer & Rock, 1994a, 1994b) whereas another posits that the two processes occur in parallel (see Peterson, 1994a, 1994b).

Distinguishing one overlapping figure from another requires depth or figure-ground segregation: The target item serves as the figure and the remaining items form the ground from which it must be segregated. If depth segregation precedes object recognition, CK should perform as poorly on the overlapping faces test as on the overlapping objects test. If, however, CK performs differently on the two versions of the test, it would provide evidence for the parallel depth-segregation and object-recognition theory.

Participants were presented with three overlapping caricatures of famous people, whom they first had to identify. If they failed, they then had to select the targets from five, free-standing caricatures [ILLUSTRATION FOR FIGURE 8 OMITTED].

Results and Comment. CK performed as well or better than controls in identifying and recognizing overlapping faces, even though controls found this test to be far more difficult than the overlapping object test on which CK was profoundly impaired (see Table 13). The discontinuities created by overlapping figures and the difficulty of appropriately assigning contours and parts seem to depend, in large measure, on the integrity of specialized perceptual mechanisms. Because the mechanisms involved in identifying objects were damaged in CK, the degradation caused by overlapping figures exacerbated an already severe deficit. The face-recognition mechanisms, however, were intact and able to support normal performance on this task.

There is no denying that laying one figure over another introduces discontinuities in contours and divides into parts what formerly was a seamless whole. One could make the case that some type of part-based process is needed to integrate the resulting discontinuous figure into a unified whole. Whatever the process, it does not seem to be impaired in CK since he can use it to recognize overlapping faces. It therefore cannot be the part-based process believed to be so crucial for object recognition. Alternatively, it

could be argued that overlapping one figure with another does not truly break the figure into discontinuous parts but rather introduces background visual noise that interferes disproportionately with processing in the damaged object-recognition mechanisms as compared to the intact face-recognition mechanisms.

The results argue against the sequential depth-segregation-first theory (Palmer & Rock, 1994a) and in favor of the parallel depth-segregation and object-recognition theory (Peterson, 1994a; 1994b). Our finding that CK's performance on this type of depth-segregation task differs markedly depending on whether faces or objects are used indicates that depth segregation cannot take precedence over object recognition because if it did, performance should be equivalent on the two tests. Instead, the results indicate that object recognition occurs in parallel and can contribute in a top-down fashion to depth segregation. The alternative interpretation that the differences occur at a later identification phase that follows depth segregation does not hold for two reasons: CK's object identification is far worse on the overlapping figure test than on identification of isolated objects; and he cannot even trace the contour of different overlapping objects, something he could do perfectly for isolated objects even when he could not identify them.

We delay the discussion of the locus of the deficit in integrative agnosia to the general discussion that follows Experiment 14.

Conclusion. CK shows normal recognition of overlapping faces but not objects, indicating (1) that if part-based processes are necessary for overcoming figural discontinuities caused by overlap, they are intact in CK and (2) that object/face recognition occurs in parallel with depth segregation and contributes to it.

```
Table 13. Mean Correct Identification and Forced Choice Recognition
of Overlapping Caricatures of Famous People (Max = 27).

                      Identification     Forced Choice Recognition

                  Mean   SD   Range       Mean    SD    Range

Controls (n = 12) 10.0   8.2  2-21        26.0    1.9   25-27
CK                14.0                    26.0
```

Experiment 14: Recognizing Mooney Faces and Objects

In the previous experiment, the discontinuities produced by overlap did not produce a noticeable break in the figure as much as a seam across which the contours continued to be joined. Would CK's face recognition be preserved if performance depended on closure, the filling in of a perceptual gap between component parts? To test this hypothesis, we used Mooney (1956) figures, black and white pictures of objects and faces constructed of patches of intense flat light and shadow [ILLUSTRATION FOR FIGURE 9 OMITTED]. To identify these pictures, the "empty" spaces, whether of shadow or light, need to be filled in to complete the figure. Patients with parietal lesions who have visual, constructive problems perform poorly on this and other tests of closure whether objects or faces are involved (Milner, 1980; Warrington & Rabin, 1970). If figure completion is

an integrative part-based process similar to one that occurs in normal object recognition, CK should be as impaired at recognizing Mooney faces as objects.

Like the previous experiment, this one has implications for theories of object recognition and figure-ground or depth segregation. Figure recognition is accompanied by segregation of the figure from the ground. If the depth-segregation-first hypothesis is correct, CK should perform equivalently on faces and objects. If, however, performance on faces exceeds that on objects, it would argue in favor of the parallel object-recognition segregation theory.

Participants viewed a set of Mooney figures consisting of faces and objects that they had to describe.

Results and Comment. As Table 14 shows, CK identified all seven faces correctly but recognized only one object, whereas controls recognized all the objects but missed one face on average. Perhaps more than on any other test, CK's performance on this one places the specificity of his preserved abilities in stark contrast to his impaired ones. After having been frustrated at identifying the first couple of objects ([ILLUSTRATION FOR FIGURE 9A AND B OMITTED]), CK said, upon being presented with his first face [ILLUSTRATION FOR FIGURE 9E OMITTED] "I told you I can't identify any of these; they all look like blobs to me." Encouraged to persist and allow the figure to emerge, he exclaimed, "My God! I can see it. It's a face." He then described it correctly as that of a child looking to the right and down but then he spontaneously asked, "What's he looking at? It's white and has something spread out that looks like wings. I guess he's watching a bird fly by." When shown Figure 9c, he described the person correctly as an old man staring to the left but then he was puzzled, "What is that in front of his mouth? It's not a beard. It looks like something he's eating." One can fairly draw a line separating the stimuli that the face-recognition system will accept and those it cannot. The line demarcates the border where the face ends and other objects, including body parts, begin. Interestingly, the only object he was able to identify was a horse's head [ILLUSTRATION FOR FIGURE 9D OMITTED]. "It looks like a face," CK said on seeing it, "but it's too long, could it be a cow or a horse?"

It is clear that the closure needed to bring the parts of the stimuli together is beyond CK's abilities if the stimuli are objects but not if they are faces. Thus, it is not the case that the part-based process needed for closure, if indeed it is part-based, is damaged in CK. The deficit is not in the process but in the specific domain or representation on which the process operates. Requiring closure makes it even more difficult for CK to identify objects, just as it does to identify faces. Because he is already so poor at one and so good at the other, the difference between the two is exaggerated.

Like the findings on depth segregation and identification of overlapping figures, the results from this study indicate that closure is driven as much by top-down processes derived from object identification as from bottom-up processes driven by the stimulus input. The type of closure demanded by the Mooney figures test is a special case of depth segregation since an important aspect of the task is to assign the regions of shadow and

light to figure and ground. The results clearly favor the theory that object recognition and depth segregation occur in parallel over the theory that depth segregation precedes object recognition (Peterson, 1994a, 1994b).

Conclusion. CK's recognition of Mooney faces is intact but that of objects is very impaired, indicating that if closure is an integrative, part-based process, it is functioning normally in CK; it is just the domain or representations to which it is applied that is impaired.

Discussion of Experiments 13 and 14

These experiments were designed to test the idea that if face recognition were made more dependent on part-based processes than holistic ones, CK would be as impaired in recognizing faces as objects. Because there is little in current theories to specify which facial parts enter into the part-based process, or what the integrative process combining these parts might be, we used our intuition to guide us. Even more than testing ideas derived from these theories, our experiments provide information that can constrain them.

We began first by seeing whether face recognition is impaired by introducing discontinuities in facial stimuli, caused either by overlapping one face with another or by separating one part of a face from another by regions of light or shadow. We reasoned that good performance requires the assignment of parts to the appropriate figure so that closure and a good form is created that can be distinguished from ground (Peterson, 1994a, 1994b). Other investigators (Nakayama et al., 1989; Palmer & Rock, 1994a, 1994b, and references therein) assume that figure ground or depth segregation takes precedence, but in either case, assignment of parts to figure or ground is involved.

```
Table 14. Correct Identification of Mooney Faces and Objects.

                     Faces (Max. = 7)       Objects (Max. = 9)

                   Mean    SD    Range     Mean    SD    Range

Controls (n = 12)  6.6     0.7    5-7      7.2     1.7    3-9
CK                 7.0                     1.0
```

The results indicate that insofar as these processes can be construed to be part-based, they are intact in CK and do not account for the dissociation between face and object recognition. Although the same process is involved regardless of the class of stimuli to which it is applied, it is only object recognition that is impaired. Moreover, it is not only impaired in the sense that the object cannot be recognized, but unless the object is clearly isolated (as in the "bird" example), CK cannot even determine what the shape or contour of the object is. The objects appear as fragmented "blobs." Part-based processes involved in depth segregation and closure, therefore, are not what distinguish perception of faces from that of objects.

Our findings also are directly relevant for theories about the relationship between object recognition and depth or figure-ground segregation. CK's performance on tests involving

depth segregation is determined by the type of stimulus rather than by the process. Thus, his performance provides strong evidence for the view that top-down processes associated with object recognition occur in parallel with depth segregation and contribute to it, and it speaks against that the view that depth segregation must precede object recognition.

CK's integrative visual object agnosia has traditionally been classified as a peripheral, rather than a central, agnosia for two reasons: (1) The agnosia is easily exacerbated by stimulus degradation that is achieved by adding visual noise, shadows, or overlapping figures, and (2) his internal visual representations of objects, as revealed by tests of imagery, are intact, which means that the agnosia does not originate centrally. Given this classification, one might have assumed that the early processes involved in feature extraction and depth or figure-ground segregation were impaired. The results of Experiments 13 and 14, however, indicate that the locus of the deficit cannot be at these early peripheral stages because his performance on faces, in contrast to that on objects, is normal.

Instead, the results suggest that the locus of the deficit in integrative agnosia is at the region at which peripheral processes common to both faces and objects diverge to deliver their input to domain-specific regions involved in identification of complex visual stimuli. The most likely locus is the infero-temporal cortex, a region rich in areas dedicated to domain-specific perception that lies anterior to striate and extrastriate areas concerned with identification of basic visual features and contour (Allison et al., 1994b; Gross, 1973, 1992; Gross et al., 1993; Ungerleider, 1995; Van Essen, Anderson, & Felleman, 1992). This region also coincides with areas that are damaged in visual object agnosia (McCarthy & Warrington, 1990) and that are activated on tests of object, word, and face recognition in humans (see "Introduction" for references on neuroimaging).

Our results also suggest that the agnosic person's sensitivity to visual degradation may not necessarily arise only because peripheral mechanisms are damaged but also (or primarily) because perception under degraded conditions does not benefit from the contribution of top-down processes. When such top-down information is available, as it is for faces in CK's case, he easily overcomes the effects of visual degradation.

Even though discontinuities between parts were introduced in Experiments 13 and 14, the parts still retained the same spatial relation to one another. Put another way, although there were discontinuities, there was no distortion or violation of the facial gestalt. In the next set of experiments, we examine the effects that such distortions produce on face recognition.

The Effects of Altering the Spatial Relations among Parts (Spatial Distortion) on Face Recognition: Experiments 15 through 17

Almost all theories of face recognition emphasize the importance of configural or holistic processing (see "Introduction"). It is the configuration of the face, the spatial relation of the parts to each other, that is as crucial for face recognition as the particular form of the

individual parts. Indeed, some theorists have asserted that the individual parts are identifiable only in relation to other components (Farah, 1990; Tanaka & Farah, 1993). When isolated, they are not recognized as belonging to a particular individual. Thus, a nose is recognized as belonging to a particular individual only when the nose is embedded in the face and not if it is presented in isolation. By these accounts, inversion alters the perceived spatial relations that parts bear to each other, making them behave more like isolated units than integral units (Garner, 1974) of a configuration or gestalt. If the face-recognition mechanisms that are intact in CK rely on configuration for their input, altering the spatial relations among facial components should impair his recognition as much as inversion. As well, if configuration is important, removing any single component should have less of a deleterious effect on recognition than changing spatial relations, so long as sufficient components remain for maintaining the configuration. These hypotheses were tested in this set of experiments.

Experiment 15: Recognizing Fractured Faces in Which Spatial Relations among Internal Features Are Altered

In Experiment 12, we altered the configuration of the face by inverting the internal and external features independently of each other. Although the main finding from that study was that face recognition was impaired only if the internal features were inverted, the study also indicated that altering the spatial relations between internal and external features had little effect on recognition. In this study we altered the spatial relations among components of internal features (as well as external ones) without inversion. We accomplished this by cutting photos of faces of famous people into five or six parts and spreading them apart while retaining a semblance of the first-order relations among them. That is, the arrangement, from top to bottom, of forehead, eyes, nose, mouth, and chin was preserved although the spatial distance between them was altered [ILLUSTRATION FOR FIGURE 10 OMITTED]. Also, all of the main internal components of the face, the eyes, nose, and mouth, were left intact, thereby enabling subjects to use them as aids to recognition. All participants attempted to identify 40 such fractured faces of famous people. If identification of the parts themselves, without regard to the configuration they form, plays a substantial role in face recognition, performance should not suffer.

Results and Comment. Altering the spatial relations among components had as deleterious an effect on recognition as did inversion. Of the faces they recognized intact, controls identified about 80% correctly whereas CK could identify only about 40%, a score that fell 6 SDs away from the mean (see Table 15).

Table 15. The Number of Correctly Recognized Intact Faces and the Percentage of Them That Had Been Recognized When Fractured (Max. = 40).

	Intact (Number Correct)			Fractured (% Correct)		
	Mean	SD	Range	Mean	SD	Range
Controls (n = 12)	33.1	6.8	21–39	81	7	68–92
CK	34			38		

The result confirms the hypothesis that the configuration of component parts needs to be maintained if face-recognition mechanisms are to be engaged. Because it was primarily the second-order relations that were altered, the results are consistent with hypothesis that these relations are crucial for identification (Carey & Diamond, 1994; Rhodes, 1988).

We know from Experiment 12 that inverting external features, which also alters the spatial relation between internal and external features, has little effect on recognition. The drop in performance in this experiment, therefore, is most likely attributable to the alteration in spatial relations among internal features. Because first-order relations were preserved for the most part, as were the identity of the parts themselves, we do not know how much they contributed to recognition.

One interpretation of these results is that fracturing the faces destroyed the facial gestalt, thereby forcing participants to resort to part-based, rather than holistic, processing in order to identify the face. Because CK's part-based processing mechanisms presumably are damaged, he performs much worse than controls. Although the interpretation is consistent with the data, the terms holistic and part-based are not formulated precisely enough to describe the algorithm or transformations used to arrive at a correct identification. For example, what constitutes a holistic representation or a holistic process? Does "holistic processing of faces" imply a norm-based process on which face recognition depends? If so, don't the "parts" of each face need to be compared with the norm? And aren't some parts more crucial for identification than others? How is this different from part-based processing? More crucial, part-based processing of fractured faces would seem to be either impossible or paradoxical if we accept Farah and her colleagues' view that parts of faces lose their "identity" when they are isolated from the facial configuration. How then can controls perform as well as they do if they rely primarily on identifying isolated parts?

We do not wish to abandon the holistic/part-based distinction. We raise these concerns to address them better in the following experiments. Our approach is to try to specify the types of distortion that CK and controls can tolerate in the faces they view and still recognize them. We return to the questions we raised in the discussion at the end of this section and in the final discussion at the end of the paper.

Conclusion. Recognition of fractured faces is very impaired in CK, and to a lesser extent in controls, suggesting that alterations of spatial relations among internal features of faces is as detrimental to recognition as inversion.

Experiment 16: Recognizing Faces That Are Misaligned Along the Horizontal or Vertical Midline

The parts in fractured faces were displaced spatially both in the horizontal and vertical direction. As well, many parts were displaced from their original location. We wished to determine whether recognition would be impaired if displacements involved whole units, such as the top or bottom of the face, and if it mattered whether the displacement was horizontal or vertical [ILLUSTRATION FOR FIGURE 11 OMITTED]. Consequently,

we had participants attempt to identify faces from photos that were misaligned along the vertical or horizontal midline. Their performance would provide information about the spatial relations that are crucial for face recognition and, perhaps, constrain what the term holistic means when it is applied to faces.

Results and Comment. CK was impaired at recognizing faces that were misaligned along the horizontal but not the vertical, whereas controls' performance was not affected as severely (see Table 16). The drop in performance for CK and controls from the intact condition was comparable to that observed for fractured faces, even though in this experiment most of the components retained their relation to one another. Indeed, in the vertical condition, no loss of recognition was observed.

The results indicate that distortion along the horizontal plane is sufficient to break the facial gestalt so that the face mechanisms can no longer process the information holistically. As a result, subjects need to rely on other mechanisms for recognition.

Table 16. Mean Correct Recognition of Photos of Famous People That Are Intact or Misaligned Along the Vertical or Horizontal Midline (Max = 10).

	Misaligned Vertical			Intact		
	Mean	SD	Range	Mean	SD	Range
Controls (n = 12)	7.6	1.5	5-10	8.0	1.2	7-10
CK	9			9		

	Misaligned Horizontal			Intact		
	Mean	SD	Range	Mean	SD	Range
Controls (n = 12)	7.5	1.5	5-10	8.0	1.4	5-10
CK	4			8		

An alternative but complementary interpretation is that in norm-based coding, values for spatial relations along the horizontal are more crucial than those along the vertical for determining the identity of the face. Another possibility, however, is that as long as the relations among at least one full set of internal parts (left or right) are unaltered, norm-based coding can proceed and arrive at a correct "solution." Vertical misalignment satisfied that condition but horizontal misalignment did not, presumably because faces are sufficiently symmetrical along the vertical midline so that there is near total redundancy in the two vertical halves but not the two horizontal ones. It would be interesting in light of these hypotheses to vary vertical misalignment so that it violates this condition, say by "stretching" the face or some components of it away from the others. Whichever of these latter interpretations is correct, they will add precision to the notion of what is meant by a facial gestalt and, perhaps, clarify what is meant by holistic processing.

Conclusion. Horizontal misalignment of the top and bottom half of the face leads to far greater impairment in face recognition than vertical misalignment of the left and right

half of the face, suggesting that spatial ordering of at least one full complement of internal features needs to be preserved as proper input to face-recognition mechanisms.

Experiment 17: Face Recognition with One Part Missing and Recognition of the Isolated Part

The results of the previous experiments indicate how important retaining the configuration of the face is for recognition. One distinction between configurational and part-based coding is that configurations can tolerate the loss of a single feature so long as other features are present. For part-based coding, however, the loss of even one feature is likely to lead to impaired performance, especially if that part is crucial (Ross & Turkewitz, 1981, 1982). Additionally, within certain boundary conditions, altering the spatial relations of parts to each other should be more detrimental for recognition than eliminating parts.

In this experiment, we tested face recognition when a single component, either the eyes, nose, or mouth, was removed [ILLUSTRATION FOR FIGURE 12 OMITTED]. We also tested recognition of the part that was removed from the face by presenting it and a lure in isolation alongside the face with the missing part. If, as some holistic theories claim (Farah, 1990; Tanaka & Farah, 1993; see discussion in Carey & Diamond, 1994), a part has no identity, or separate representation, independent of the whole configuration when it is no longer embedded in its proper configuration, it should be very difficult to distinguish the part belonging to the face from the lure.

Results and Comment. CK performed at least as well as controls both in identifying faces with a single part missing and in recognizing the missing part (Table 17). Participants served as their own controls to confirm that the missing part was identified because it belonged to the face rather than simply because there were sensory cues that indicated which stimulus fit better in the allocated space. CK scored perfectly, and controls somewhat worse, in choosing the correct target when they could identify the face and thus draw on their memory of the missing part. Both CK and controls, however, scored at or near chance when they could not identify the face.

Consistent with the configurational hypothesis, all participants were no worse at identifying faces with a single part removed than they were at identifying the intact face. Also, CK himself scored much better in this experiment in which one part was removed but the configuration was retained than in Experiment 15 in which all parts were visible but the configuration was altered. Experiments on single-unit recording from face-sensitive cells in monkeys found similar effects: Removing a part did not eliminate the response but altering the configuration did (Desimone et al., 1984; Perrett, Rolls, & Caan, 1982).

Table 17. Mean Correct Forced-Choice Recognition of Single Isolated Face Parts (Eyes, Nose, or Mouth) for Identified and Non-identified Faces of Famous People (Max. = 15).

	Identified	Recognized Part

	Mean	SD	Range	Mean	SD	Range
Controls (n = 12)	10.0	2.9	4–13	7.0	3.3	1–12
CK	10			10		

	Nonidentified			Recognized Part		
	Mean	SD	Range	Mean	SD	Range
Controls (n = 12)	5.0	2.9	2–11	3.0	2.6	0.8
CK	5			2		

Removing one part still allows norm-based coding to proceed based on the preserved spatial relations among the remaining parts. Misaligning the bottom half of the face with the top alters the spatial relations among parts on which norm-based coding depends.

Although the results provide additional evidence in favor of the hypothesis that part-based information is not as crucial for face recognition as configural information, they are problematic for the strict holistic hypothesis (see also Experiment 3 on parents and offspring). As stated by Farah and her colleagues (Farah, 1995; Tanaka & Farah, 1993), one implication of that hypothesis is that parts are not recognized when isolated from the configuration to which they belong. In a sense, they do not have an independent identity. Our results indicate otherwise. Both CK and controls generally succeeded in correctly attributing the isolated missing part to the face. This finding, as well as that of Experiment 3, indicates that face mechanisms, even in the absence of intact object-recognition mechanisms, are capable of representing facial components explicitly and separately from the whole face. Thus both the identity of the parts and their configuration are represented by face mechanisms. We will speculate in the final discussion on how the two types of information are used and how they are related to the operation of part-based mechanisms typically used to identify objects and inverted faces.

Our results do not contradict the weaker version of the holistic hypothesis, which states that parts are recognized better in a face context than when they appear in isolation. Thus, Tanaka and Farah (1993) showed that participants are better at recognizing facial features belonging to particular individuals when the part is embedded in the face than when it is presented in isolation. The same was not true of houses or of inverted faces. Although we did not test CK on this experiment directly, we think it very likely that he would perform as normal people do on faces, but much more poorly with objects. Indeed, Harman and Moscovitch (unpublished observations) tried to teach him to identify four different cars and four different exemplars of the same breed of dog, and he failed at both tasks after dozens of attempts. He was, however, able to learn the names of four new faces within two trials. We then had him view the face through a movable, round aperture whose size varied. Identification of the individual when only a single component was visible at a time was poor both for CK and for controls, although both improved when two components were visible simultaneously.

Taken together, the results indicate that although representation of faces contains information about the identity of single parts, as well as the spatial relations between them, the identity of the part by itself is a poor cue to recognition of the face. For that

matter so may be information about spatial relations if such information could be presented independently of the parts. The implication of this conclusion for a theory of face recognition will be left to the final discussion.

Conclusion. CK can recognize faces when only single features are removed and the general configuration retained, and CK can recognize isolated features as belonging to particular faces, indicating that face mechanisms represent both the identity of individual components as well as the spatial relations among them.

Discussion of Experiments 15 through 17

When the spatial relations among the parts of the face were altered, face recognition suffered in normal people and was severely impaired in CK. For both CK and normal controls, the effect of altering spatial relations was as detrimental to face recognition as inversion. Because previous experiments indicated that altering the spatial relations between internal and external parts of the face did not affect recognition (Experiment 12), we can conclude from Experiments 15 and 16 that what is crucial for face recognition are the relations among internal parts. Misalignment of parts only along the horizontal plane certainly makes recognition difficult, and except for shifts along the vertical midline, it is likely that distortion in the vertical plane, such as "stretching," would have a similar effect. Because the face recognition system relies on configural information, it can tolerate the loss of a single feature without recognition being affected (Experiment 17; and Desimone et al., 1984; Perrett et al., 1982; Yamane et al., 1988). Together, the results from this set of experiments suggest that an upright left or right half-face in which the parts retain their second-order spatial relations but from which one part may be removed is sufficient to support good face recognition.

In light of some holistic theories of face recognition (Tanaka & Farah, 1993; Farah, 1995), an unexpected finding was that both controls and CK were able to choose correctly the isolated face part that belonged to a particular face. Contrary to other reports, this result indicates that individual face parts can be recognized in isolation. Differences in task requirements may account for the discrepancy between our findings and those previously reported in the literature. In most experiments in which recognition of isolated parts was tested, participants effectively were required to identify the face to which the part belonged. As we noted earlier, this finding means simply that an isolated part is not a good cue to face recognition; it does not mean that a configural face-recognition system does not have access to the identity of the individual parts, nor does it mean that individual parts lose their identity when not embedded in the configuration. On the contrary, the results of Experiment 17 suggest strongly that when a representation of a face is recovered from the face-recognition system, it contains information about both the identity of the parts and the spatial relations between them. The parts can then be identified in isolation if what is required is identifying the part that fits the incomplete face rather than using the part to recover information about the face.

To the points noted previously regarding the characteristics of a face that activate the face-recognition system, we add the following:

7. The spatial relations between internal features of the face need to be preserved. Distortion of spatial relations leads to loss of identification that is dependent on the face-recognition system.

8. Loss of any single component does not affect face recognition.

9. Information about specific facial features, as well as the spatial relations among them, is represented in the face-recognition system.

Experiments 18 and 19: Modularity of Face Recognition: Shallow Output and Informational Encapsulation/Cognitive Impenetrability

Fodor (1983; 1985) argued that much of high-level perception is mediated by domain-specific perceptual-input modules that satisfy a number of criteria that distinguish them from general-purpose central systems. In evaluating Fodor's proposals from a neuropsychological perspective, Moscovitch and Umilta (1990) proposed that to be considered a module, a system needs to satisfy only three major criteria: domain specificity, informational encapsulation/cognitive impenetrability, and shallow output. Thus a module (1) is activated only by stimuli in a specific domain, (2) is informationally encapsulated so that higher-order cognitive information cannot penetrate it to gain access to its intermediate-level representations or influence its output, and (3) its output is shallow in that it is not semantically interpreted, it only provides information pertinent to its domain, and it does not contain any information as to the source from which its output is derived.

The neuropsychological evidence reported in the literature and gathered in this study makes face recognition a good candidate for modularity in Fodor's strong sense. Most of the experiments reported in this study attest to the specificity of the domain over which the face-recognition system operates. The evidence for the other two criteria is less compelling. On the one hand, Rhodes and Tremewan (1993) found that preceding a picture of a face with its name improved detection of that face, suggesting that face modules are not wholly impenetrable to cognitive influences outside their domain (however, see Norris, 1995, for a critique of this interpretation). On the other hand, evidence of implicit face recognition (without awareness) in patients with prosopagnosia indicates that face-recognition modules exist whose operation is cognitively impenetrable to central systems concerned with consciousness (McNeil & Warrington, 1991; Renault, Signoret, Debruille, Breton, & Bolger, 1989; Sergent & Signoret, 1992b; Tranel & Damasio, 1985; see also reviews by Bruyer, 1991; Nachson, 1995; Young, 1994). Such evidence is consistent with the fact that even when the person has detailed semantic knowledge about individuals whose faces she or he cannot recognize explicitly, such knowledge does not penetrate the workings of the module to make it yield the information that would lead to adequate explicit recognition. As well, studies on implicit recognition in prosopagnosia suggest that names are strongly linked to those modules so that they can be activated automatically by the module's output without the subject being aware of the link between them. As a result, faces that are recognized only implicitly can prime names that then are read or learned more quickly (Bruyer et al., 1983; de Haan et

al., 1987) than unprimed names and that lead to differential skin conductance responses when the names do not correspond to the faces (Bauer, 1984). The close, automatic association between names and faces suggest either that they are part of the propositional knowledge contained in the face-recognition module or that the module's shallow output automatically activates representations for the corresponding names.

The studies on implicit face recognition in prosopagnosia can also be interpreted as providing evidence of shallow output. The output of the face module that supports implicit recognition is not interpreted at a deep level, which implies access to a body of stored knowledge about the face and the ability to relate that knowledge to other information. Rather, the output is shallow in that it contains information only about the structural representation of the face and the name with which it is associated, and even that information is not accessible to consciousness.

This argument in favor of shallow output, although defensible, is weak. To make it stronger, we need to demonstrate that (1) the face-recognition system will accept any stimulus that satisfies its required input properties without regard to how that stimulus is constructed and (2) that its output only represents information pertinent to the structural (nonsemantic) properties of faces; information about other, nonface aspects of the stimulus should be lost. That is, once activated, the face-recognition system should give rise to a percept of a face without the awareness of the possibly nonfacial nature of the stimuli that make up the face.

Consider, for example, a face made up of fruit, books, or parts of a woman's body arranged in such a manner that they resemble a face. The Italian artist Arcimbaldo [ILLUSTRATION FOR FIGURE 13 OMITTED] was noted for painting such composite faces, and other artists, such as Escher, Dali, and Magritte, also tried their hands at it. Apart from their artistry, what makes these paintings compelling and whimsical is that we are aware of both the face and the nonface stimuli of which they are composed. If a face-recognition system is truly modular and its output is shallow, it should not by itself give rise to the type of double awareness that we all experience.

According to the argument we have been developing in this paper, this type of double awareness is mediated by the activation of at least two separate systems - one for face recognition and one for object recognition. In the absence of an intact object-recognition system, the face-recognition system should not be able to support this double awareness. Moreover, face perception should not be prone to possible interference that arises from competition with an object-recognition system for the focus of attention in perception. We tested these hypotheses in the next two experiments.

Experiment 18: Recognition of Faces and the Nonface Objects of which the Faces are Composed: The Arcimbaldo Effect

An intriguing question was how CK would perceive faces made up of objects, such as a face with cherries for eyes, a pear for a nose and a banana for a mouth. This question gets at the heart of the issue of modularity. Without an intact object-recognition system, it is

possible that CK simply would not recognize the objects and, therefore, be unable to assemble them into a face. But if the face-recognition system were truly modular, it should make no difference whether the objects are recognized or not: The system should be activated if the objects are placed in the proper relation to each other to satisfy the configurational properties of a face, as we have identified them in this study. When that occurs, the system should be triggered and face perception should be mandatory.

The anecdotal evidence in this regard is strong. We immediately perceive any such arrangement as a face, although we are aware as well of the objects of which it is composed. In a more experimental vein, Jeffreys and Tukmachi (1992), using ERPs, showed that an early waveform, P190, which distinguished faces from other objects, also was evident when an individual viewed faces comprised of objects. Once the facial configuration was broken, the facial percept, as well as the identifying waveform, was lost.

Based on the anecdotal and experimental evidence, we were confident that CK would perceive the object-composite faces as faces. The more interesting question was whether he would be able to see past the face to the objects that comprised it. Although CK is impaired at recognizing objects, he is not impaired at detecting them or at giving a detailed description of their parts. We, therefore, expected him to see the faces, as well as to be able to say that they were composed of nonface objects, although we thought he might have some difficulty in identifying exactly what those objects were. However, in accordance with the criterion of shallow output, and without an intact object-recognition system to fall back on, CK should be aware only of the face itself and not of the objects of which it is composed. The results we obtained were exactly what an extreme modularity hypothesis predicted: The shallow output of the face-recognition system delivered information only about the structural properties of faces but not about the nonface elements from which the output was derived.

To test our hypotheses, CK and 12 control participants viewed and described eight composite faces taken from the paintings of Arcimbaldo ([ILLUSTRATION FOR FIGURE 13 OMITTED] for a sample). About 8 months later, CK was shown five additional, less detailed paintings from other artists, as well as the initial nine Arcimbaldo figures in an inverted format.

Results and Comment. All the control participants saw each of the pictures as both a face and as the objects that comprised it (Table 18). For CK this was true of only two of the faces, the tree-face of the painting Inverno and the animal face of Terra, and partially true of another one, Natura, a vegetable face. Even for these, it took a great deal of prompting and quite some time before CK came to the realization that the faces were composed of objects. For controls, the realization was immediate and some faces, like Terra, were perceived at first only as objects and only later did the face emerge. For all the other faces, CK could describe the face clearly but had no awareness that it was composed of objects, despite being prompted to report everything of which he was aware. For the five extra faces that were presented only to CK, he did not detect the nonface parts of any of them. When the faces were inverted, he immediately noted the nonface parts but could not

apprehend any faces. These results are consistent with the hypothesis that the output of the face-recognition system is shallow, as befits a module in Fodor's sense. They also add support to the hypothesis that once a face is inverted, the face-recognition system no longer apprehends it ([ILLUSTRATION FOR FIGURE 13 OMITTED], Natura).

Table 18. Number of Faces and Nonface Aspects of Painting That Were Recognized As Being Formed of Objects When Upright and Inverted (Max. = 8).

	Upright		Inverted	
	Faces	Nonface	Faces	Nonface
Controls	8	8	–	–
CK	2	8	8	8

CK's reports of his perceptual experience are instructive. When viewing Figure 13a, he described the individual as "a happy-looking man, facing to the right with the eyes looking slightly in the other direction. The cheeks are red and he has a large nose. He also seems to have some bags under his eyes" (pointing at the cherries). When prompted to describe anything else he saw, he noted, "He seems to have a pear and some other fruit in his hair, some lettuce for sleeves, and something that looks like a sea urchin on his chin." When prompted further, he could not offer any more information. Indeed, in two other cases, he was able to note the objects if they did not occupy the location of the eyes, nose, or mouth. If the objects fell on the external regions, he seemed more likely to identify them. This finding is consistent with our observation that it is the configuration of internal features that drives the face-recognition system. Because the external features fall outside the domain of the system, objects there are more likely to be detected by those functionally limited nonface mechanisms that survived damage.

For another picture [ILLUSTRATION FOR FIGURE 13C OMITTED], CK noted immediately that it was a man with a mustache facing right. When prompted to provide a richer description, CK began naming, often erroneously, some of the animals in his hair and then moved toward the center. when he named the elephant, he exclaimed, "Wait a minute. This man's face is made up of animal heads. That's fantastic." when the experiment was over, we informed him that all the other pictures he had viewed were constructed similarly. He was incredulous. We returned to show him that this was the case, beginning with Rudolfo, the fruit man. Examine it as he would, he could still not "see" the fruit in the internal features, although he had no difficulty identifying the pear in the man's hair. When asked why he thought it was that he could see the "animal man" but not the "fruit man," he offered the following explanation, which we think is correct. "The animals all have faces, and it is the faces that I see. Although they aren't people faces, they look like faces and that may be enough." Although animal heads are not as effective at activating the face-recognition system as human faces, partial activation may occur because animal heads share some features with human faces (see Experiments 11 and 14). Alternatively, animal heads may activate some residual object-recognition system that allowed CK to identify them better than other objects.

Two observations lead us to believe that a contributing factor to CK's dim, or absent, awareness of the objects that comprise a face is the competition between an intact face-recognition system and an impaired object-recognition system for access to a central system that supports awareness. when the same Arcimbaldo faces were shown to him in an inverted orientation 8 months later, he never mistook them for a face, and he was able to identify the general class (though hardly ever the specific exemplars) that comprised the face. Indeed, after viewing these inverted "faces," he seemed to have recognized the style and spontaneously asked whether he could see those colorful faces he had viewed on the previous session.

The second observation concerned the composite of Magritte's painting "The Rape," which consists of a woman's nude body framed by hair. The picture is stark and its lines simple, unlike Arcimbaldo's detailed, Baroque paintings. Despite this, CK immediately perceived it only as a face without recognizing the parts. When prompted for a more detailed description, he said. "The person is exothalmic. (What do you mean exothalmic?) I mean the person has goiter. See, look at those big, bulging eyes (referring to the breasts). And that thick neck (referring to the thighs) caused by an overdeveloped thyroid gland. (Anything else?) The person has a beard (referring to the pubis) but there is a dot (navel) where the nose should be." At this point, one of us covered the beard and nose while CK was viewing the picture. His head literally snapped back in surprise at the objects that now revealed themselves.

As long as the intact face-recognition system was operating, the output of the impaired object-recognition system could not compete effectively for access to conscious awareness. Thus, CK was aware only of the shallow output of the face module and not of the nonface elements that comprised the face. Once activation of the face module was removed, by inverting the face, occluding some of the crucial features, or having objects occupy locations outside the internal features, CK could detect the objects and identify the class to which they belong if not the specific item.

The notion of competition between object- and face-recognition systems for conscious awareness leads to the prediction that CK should perform better than controls in circumstances in which such competition leads to interference. We tested this idea in the next experiment.

Conclusion. When viewing upright, intact faces composed of objects, CK could recognize the face but could not even detect the objects that make up the face when they comprised the internal features; this evidence of shallow output bolsters the case for considering the face-recognition system to be a module in Fodor's sense.

Experiment 19: The Interfering Effect of Object Recognition on Face Recognition: Not Seeing the Faces for the Trees

Finding hidden targets in a picture or drawing is a captivating and engrossing children's game that even adults find challenging [ILLUSTRATION FOR FIGURE 14 OMITTED].(4) Such hidden targets are difficult to detect because they share contours or

components with other, more salient stimuli that compete with the target for perceptual awareness. If, however, the competing stimuli are less salient or perceptible, the hidden targets should be easier to detect.

This was our reasoning when we administered a hidden stimulus test with faces as targets and objects as distracters for CK and control participants. The purpose of the test was to find all the hidden faces in a visual scene of a forest clearing containing nonface objects [ILLUSTRATION FOR FIGURE 14 OMITTED]. The faces were all composed of objects such as the trees, rocks, and streams, which form the subject matter of the picture. The configurations that these objects formed activate the face-recognition system. Because the output of the face-recognition system is shallow, as Experiment 18 suggested, and because CK lacked an intact object-recognition system, the objects were less likely to capture his perceptual awareness than that of control subjects. As a result, the "faces" in the forest were likely to be detected more easily by CK, who was not as prone to interference from the output of a competing object-recognition system.

Results and Comment. CK was able to detect faces more quickly and in greater number than the best control participant (Table 19). One face simply jumped out at him during the first second, three other were seen in the first minute. Most control participants could discover no faces at all for over a minute and, over the entire 5 min duration, the typical participant discovered fewer faces than CK did in 1 min.

The results are consistent with the hypothesis that the effectiveness of this version of the hidden target effect depends, in part, on the competition between the output of a face- and an object-recognition system for perceptual attention. Successful competition from the object-recognition system prevents normal people from detecting the hidden faces. As CK's performance demonstrates, once that competition is reduced or eliminated, in his case because the object-recognition system is damaged (and because the output of the face-recognition system is shallow), faces are detected more easily.

Conclusion. CK's superior ability to detect, in a visual scene, hidden target faces composed of objects indicates that the outputs of the face-recognition system, which are shallow, compete with the outputs of the object-recognition system for perceptual awareness; when the latter system is damaged, interference with the face-system output is reduced and its access to conscious awareness is facilitated.

Discussion of Experiments 18 and 19

These experiments provide a compelling demonstration of shallow output of a module. Relying primarily on the output of his face-recognition system, CK is unaware of the nonface elements that serve as the input to that system. In Experiment 18, this prevents him from seeing the objects that comprise the face, and in Experiment 19, it allows him to see past the objects to the hidden faces, which normal people have difficulty finding among the interfering objects.

One of the factors that may have contributed to the especially compelling effects that we observed in the two experiments is that the input to the face- and object-recognition systems were derived from the same spatial location. Under these conditions, competition between the two systems may be particularly intense since a location with overlapping stimuli demands that a single interpretation be assigned to the stimuli at any given time. Such a situation is not unlike that found in ambiguous pictures in which two figures appear as either background or foreground (see Experiments 13 and 14). One cannot see both figures at once but is forced to switch from one percept to another. When one of the recognition systems is damaged, as is the case for object recognition in CK, its output may be too weak to dislodge the shallow output of the intact face-recognition system and capture attention for that spatial location. If the two systems do not receive input from overlapping spatial locations, there may still be some competition between their outputs for attention, but it is not of a strong, mutually exclusive variety. Indeed, CK can detect objects outside the region over which the face-recognition system operates, and he can identify objects if the face-recognition system is temporarily made inactive by occlusion or inversion, thereby eliminating the domain-specific input it needs to be activated.

Using commissurotomized patients and the composite faces from our study, Gazzaniga (1997) recently provided evidence regarding the neural substrates underlying the competition observed in our study. When Arcimbaldo faces were presented to the two hemispheres, the left hemisphere reported seeing only the objects that comprised the face, whereas the right hemisphere reported seeing the face but not the objects. Gazzaniga's finding suggests that CK's face-recognition performance reflects the operation of an intact right-hemisphere holistic mechanism. Object recognition, on the other hand, depends on the operation of an analytic left-hemisphere mechanism.

A possible interpretation of CK's performance based on the global/local hypothesis is that CK may have intact global processing but impaired local processing, which accounts for his good global face recognition but poor ability to identify the local objects that comprise the face (Navon, 1977; Robertson & Delis, 1986; Sergent & Hellige, 1986). Whether or not CK is good at processing [TABULAR DATA FOR TABLE 19 OMITTED] global but not local information is a matter for future research to determine. Whatever the outcome of that research, we do not think that his performance on our task can be explained only in terms of the global/local hypothesis. CK had no difficulty in identifying the local features of the face - he just could not see them as objects. Moreover, he had no difficulty in detecting the objects if they occupied the peripheral contours of the face, where such local features may have come in direct conflict with global processing. Rather, CK only had difficulty detecting objects if they occupied the location of internal features of faces.

Experiments 18 and 19 present evidence that the output of the face recognition system is shallow. CK's performance also suggests that face recognition is mandatory if the stimulus input, however bizarre, satisfies the domain-specific conditions we identified in previous experiments as necessary to activate the face-recognition system. If we also consider evidence of informational encapsulation/cognitive impenetrability from studies on recognition without awareness in prosopagnosia, face recognition satisfies the main

criteria of modularity (Fodor, 1983; Moscovitch & Umilta, 1990): domain specificity, informational encapsulation/cognitive impenetrability, and shallow output. For the time being we are satisfied merely to note this. We postpone a fuller treatment of issues related to modularity for the general discussion.

GENERAL DISCUSSION

Our objective in this study was to gain a better understanding of face recognition by studying CK, a person with visual object agnosia and acquired alexia but intact face recognition. Following a suggestion from Moscovitch and Umilta (1990), we reasoned that such a study would provide us with the opportunity to investigate the processing limits of a face-recognition system with little contamination from processes that contribute to recognition of other types of complex visual patterns. In this way we hoped to answer some questions that have remained unresolved in the literature despite decades of study. Based on the result of our experiments, we think the outcome has repaid our efforts, and we hope that the reader agrees. Because each experiment, as well as each set of experiments, was followed by extensive discussion, for the most part we limit ourselves in this final section to summarizing the main empirical findings and theoretical implications. We conclude with a brief discussion of modularity and of the elements of a neuropsychological model of face recognition.

Aspects of Recognition of Upright Faces

If faces are upright, CK recognizes and remembers them as well or better than neurologically intact people (Experiments 1 through 5). He detects family resemblance and appreciates transformations that occur in development and aging, as well as those that occur across shifts in viewpoint and changes in lighting. In short, for recognition of upright, whole faces, CK's performance suggests strongly that an isolated face recognition system is sufficient.

What Stimulus Features Are Necessary for Activating a Face-Recognition System?

We list the stimulus attributes our results suggest are crucial for face recognition. These constitute the domain-specific information necessary for activating the face-recognition system.

1. The face must be upright. Inverted faces are not recognized by CK (Experiments 9 through 11).

2. The crucial information for facial identity is carried by a spatial configuration formed of the internal features of a face - the eyes, the nose, and the mouth (Experiment 12). No single component is necessary, but information regarding the spatial relation between any two of the three elements is sufficient for face identification (Experiment 17). External features seem to be much less important.

3. Faces are identified on the basis of the spatial relations of the internal features to each other (secondary relational features; Experiments 13 through 16) and with respect to their deviation from the representation of a facial prototype (norm-based coding) (Experiment 8).

4. Information about specific, internal facial features, not merely their spatial relations, is also represented in the face-recognition system (Experiments 3 and 17).

5. The particular elements of which a face is composed are immaterial as long as the required configurational properties of the face are preserved. Thus, any configuration that is facelike will do, whether it is a caricature (Experiment 6), a cartoon (Experiment 7), an object in the shape of a face (Experiment 19), or a face whose separate features are composed of objects (Experiment 18).

We think that it is important to add to this list another point concerned with the output from the system rather than the input to it:

6. The output of the system is a structural description of a specific, individual face that does not preserve information about nonface aspects of the stimulus that activated the system.

As this list indicates, the domain over which the face-recognition system operates is highly constrained and the output it emits is shallow. For these and other reasons we discussed earlier, the face-recognition system would seem to satisfy the criteria for modularity.

Before we return to the problem of modularity, we wish to deal with a related issue - what makes face recognition possible when the highly constrained conditions that satisfy the input properties of the face system are not met? One of the major findings in our series of experiments is the observation that under such conditions neurologically intact people continue to recognize faces, albeit somewhat less well, by using mechanisms suited for object recognition to supplement the face-recognition system. Because those mechanisms are damaged in CK, his ability to recognize faces under these conditions is severely compromised. This is seen clearly when faces are inverted.

The Inversion Effect

Our findings argue strongly that inverted faces cannot engage the face-recognition system directly but rather require mediation by the object-recognition system. Lacking an intact object-recognition system, CK cannot recognize even cartoon faces when they are inverted, something control participants found trivially easy. These results do not support any variations of the discriminability hypothesis, which posits that no additional mechanisms need to be invoked for processing inverted, as compared to upright, faces (Valentine, 1991). At the same time, we wish to stress that our result cannot be interpreted to mean that the face-recognition system can be dispensed with during recognition of inverted faces. Although direct proof is lacking, we believe that the face-

recognition system is still needed to process the information about inverted faces that is gathered by the object-recognition system. How information from one domain is mapped onto another is one of the central issues that future research must address.

When and What Does the Object-Recognition System Contribute to Face Recognition?

As we noted, the object-recognition system is needed when any of the stimulus conditions listed in points 1 through 4 are violated. Inverting whole faces or simply their internal features, altering the spatial relations among internal features, or occluding more than one of the internal features deprives the face-recognition system of the configurational, domain-specific information necessary for its operation. Under such circumstances, face recognition relies on the contribution of the object-recognition system.

The precise nature of that contribution, however, has yet to be determined. Because the face-recognition system requires a configurational input that is orientation-specific, the ideal system to complement it would be one that takes parts and contours as its input and that can integrate and normalize them (Farah, 1990, 1991; Harman & Moscovitch, 1996; Humphreys, Riddoch, Donnelly, Freeman, Boucart, & Muller, 1994). As a result, a number of investigators have suggested that face perception becomes part-based once it is dependent on the object-recognition system (Farah, Tanaka, & Drain, 1995), although exactly what that entails is not clear. One possibility is that faces are now recognized on the basis of individual features, but there is no direct evidence that this is the case. Our own observation with CK indicates that there must be more to it than that. First, he can identify individual features of the face, sometimes as well as controls (Experiment 17), suggesting that the face-recognition system can also represent individual features. Second, when CK could identify inverted or fractured faces, he sometimes did so by identifying the separate features if they appeared to be salient, such as Michael Douglas' dimpled chin or Prince Charles' large ears. His appreciation of salient features is consistent with his recognition of caricatures (Experiments 6 through 8). Third, the deficit CK has in object recognition is not in identifying parts but in integrating them into a whole. When he did have difficulty in identifying parts of faces, as he sometimes did in dealing with inverted cartoon faces, and as he always did in making sense of overlapping objects and Mooney figures, it is because the proper interpretation of the parts depended on appreciating their relation to the whole.

Consistent with the nature of CK's disorder, what these observations suggest is that the contribution of the object-recognition system is not primarily in identifying facial features. Rather, the object-recognition system is needed for integrating facial features into a reasonable facsimile of a face when the gestalt is violated. This facsimile may be needed to gain access to the face-recognition system. How the object-recognition system accomplishes this function is not known. Because it has access to information about external features or facial contours (Experiment 12), the object-recognition system may use them as a frame or point of reference for aligning the internal features when they are inverted or fractured. In this way, a sufficiently reasonable facsimile of the face may be constructed. Although speculative, this proposal is consistent with similar proposals

advanced by Humphreys et al. (1994) to explain the process of integration involved in object recognition.

The Interaction of Face-Recognition and Object-Recognition Systems

Even when face recognition is dependent on the operation of an intact object-recognition system, there is good reason to believe that optimal recognition depends on some interaction between the two systems. Prosopagnosic patients with relatively preserved object recognition perform very poorly at face recognition (see Farah, 1990; Farah, Wilson, et al., 1995; Bauer, 1984). Indeed, their performance typically is much worse than that of our control subjects, who perform at about 60 to 80% correct when faces were inverted or fractured, both being conditions we identified as demanding the object-recognition system. Together, these results indicate that one cannot rely only on the object-recognition system for even moderately successful face recognition. Instead, these results suggest that information about faces that is derived by the object-recognition system needs to be shared with the face-recognition system if good recognition is to be achieved. However suggestive the data, direct evidence on this point is lacking but sorely needed. As yet, we know little about the interaction between processes involved in object and face recognition. Our studies indicate that research on this topic is crucial for a full understanding of face recognition.

Models and Theories of Face Recognition

Of the various models of face recognition, our results are most consistent with the configurational, norm-based model, which states that faces are recognized on the basis of variations of second-order relations with respect to a norm or prototype (Experiment 8). Although the model emphasizes the importance of spatial relations among facial features, we wish to add that the internal features themselves are also represented. CK had no difficulty in selecting the appropriate missing feature from another one that fit equally well into the face (Experiment 17). Because specific features may also be encoded by the object-recognition system, it is difficult to disentangle the contribution of the two types of systems in neurologically intact people who possess both (see discussion in Rhodes et al., 1993). The problem is all the more difficult to resolve because the object-recognition system may also be involved in integrating parts into a whole. Given these considerations, it is understandable that inconsistencies will be found in the literature on configural and feature processing in faces.

In the introduction, we suggested that the gestalt or template hypothesis complements the norm-based hypothesis rather than rivaling it. Our findings allow us to be a little more specific. According to the gestalt or template hypothesis, faces are represented as gestalts or templates in which the features cannot be coded or processed independently but rather derive their identity from the gestalt in which they are embedded (see Yuille, 1991, for a computational model of deformable face templates). Results from our studies on fractured (Experiment 15) and misaligned (Experiment 16) faces indicate that for faces to be recognized, the spatial relations among internal features must be preserved and that this information is orientation-specific. It is this information that forms the facial gestalt.

Access to the face-recognition system can be gained only via this configurational, spatial information. That is why separate, single features or even a collection of all the features cannot be processed by the face-recognition system and why it makes identification on this basis difficult even for controls, let alone for CK. This does not mean that information about single features is not represented by the face-recognition system or that such features cannot be inspected in isolation. Having accessed the representation of the face via configural information, CK can identify the appropriate feature of the face even when it is presented in isolation. In other words, he can move from the unit to its components but not vice versa. In this regard, his performance with faces is analogous to his performance with objects. Having preserved imagery, he can draw an object and its separate components based on his internal representation of them, but he cannot gain access to that representation from vision because object perception necessarily is part-based (Behrmann et al., 1994).

The Modularity of Face Recognition

We presented evidence of our own and from other investigators that face recognition satisfies the main criteria of modularity in the strong Fodorian sense: domain specificity, informational encapsulation/cognitive impenetrability, and shallow output (Fodor, 1983; Moscovitch & Umilta, 1990). There is even some evidence that face recognition is mandatory (Experiment 19; Farah, Wilson, et al., 1995), although we do not consider this to be one of the main criteria of modularity. An alternative to the face-module hypothesis is the holistic hypothesis, which states that there is no face module but rather a specialized holistic processor whose domain is not restricted to faces but is open to all stimuli that can be processed and represented holistically. Another alternative to the face-module hypothesis is the individuation or within-category discrimination hypothesis (Damasio et al., 1982; Logothetis & Sheinberg, 1996), which states that face recognition depends on a processor whose general purpose is to distinguish among exemplars of homogeneous stimuli. Although none of the experiments we reported was designed to test these alternatives, evidence from other studies argues in favor of modularity.

A number of investigators have reported cases who presented with impaired within-category discrimination of faces but not of cars (Sergent & Signoret, 1992), glasses, coins (De Renzi, 1986b; De Renzi et al., 1994; Farah, 1995), or even of living creatures, such as sheep (McNeil & Warrington, 1993). Conversely, CK could identify faces easily but not glasses (Farah, 1995), dogs, cars (Moscovitch & Harman, unpublished observations), or "greebles" (Gauthier & Tarr, 1997), a class of artificially created stimuli that have the configural properties of faces although they look little like them (Gauthier, Behrmann, & Tart, unpublished observation).

In an attempt to see if CK could identify holistic stimuli normally, we tested his perception and recognition of geons (Biederman, 1987). Our rationale was that if geons are the building blocks of object perception, as some investigators have argued, they should act as perceptual wholes or gestalts. If CK's face recognition is mediated by a specialized holistic processor rather than a face module, CK's processing of other holistic stimuli, like geons, should be normal. In fact, he was grossly impaired in both perception

and recognition (Suzuki, Peterson, Moscovitch, & Behrmann, 1997) whenever the geon was rotated in depth or in the picture plane even though he could recognize faces rotated in depth normally (see Experiment 5).

At the moment, the evidence favors the face-module hypothesis, although it is not conclusive. What is needed for testing the holistic hypothesis properly is a satisfactory definition of what it means for stimuli other than faces to be holistic. If we accept the definition that patterned stimuli are holistic if they are identified on the basis of second-order relational properties, as greebles are purported to be, or that they are complex but indivisible units such as geons, the evidence supports the idea that face recognition is modular because CK is impaired at perceiving these stimuli but not faces. There is always the possibility, however, that this definition is wrong (perhaps Garner's, 1974, definition of integrality is more appropriate) and that some class of stimuli will be discovered that are processed as well as faces by what is presumed to be an isolated face-recognition system.

Another alternative is that second order relations exist as latent properties of the stimulus which become manifest only after extensive experience with it (Diamond & Carey, 1986; Bruyer & Crispeels, 1992; Ross & Turkewitz, 1982). It is the interaction of the potential properties of a stimulus with experience that makes a stimulus configural or holistic. According to this view, faces are not special but just the most common of potentially holistic stimuli. Studies have shown that other stimuli, such as dogs and greebles, are also treated holistically by experts but not by novices (Diamond & Carey, 1986; Gauthier & Tarr, 1997; Tanaka & Gauthier, 1997). In the hands of experts, such stimuli show the inversion, fracturing, and isolated-part effects that distinguish faces from objects, and on functional neuroimaging studies, they may even activate an area in the fusiform gyrus that overlaps with the area that is activated by faces (Gauthier, Tarr, Anderson, & Gore, 1997).

Although the studies on expertise may help explain how people come to treat faces differently from objects, they do not refute the face-module hypothesis. As we noted in the introduction, there is evidence of double dissociation between face recognition and recognition of other complex visual stimuli such as dogs, cows, sheep, and cars, even when the individual was an expert in recognizing them (see discussion on p. 557). In CK's case, though he could recognize faces normally, he no longer could recognize airplanes and tin soldiers, though he was an expert at both before his accident. Until such dissociations are shown to be spurious or artifactual, our working assumption is that face recognition is modular.

Neural Substrates of Face Recognition

The neurological and neurophysiological evidence is consistent with the modularity hypothesis. Beginning with Gross et al.'s (1972) single-unit recording studies in nonhuman primates, numerous other investigators have reported that there are neurones in the fusiform gyrus of the inferotemporal cortex that respond selectively to faces (Gross, 1992; Gross et al., 1993; Perrett et al., 1992; Wright & Roberts, 1996 and references

therein). Damage to homologous regions in humans leads to prosopagnosia. Although typically the damage is bilateral and often also causes object agnosia and dyslexia (Meadows, 1974; Benton, 1980; McCarthy & Warrington, 1990), a number of cases of relatively pure prosopagnosia have been reported following lesions restricted to the right hemisphere (De Renzi, 1986a, 1986b; De Renzi, et al., 1994 and references therein). Gazzaniga's (1997) observation of the Arcimbaldo effect in a commissurotomized patient attests to the right-hemisphere locus of the face-recognition system as do many laterality studies in people who are neurologically intact (for review see Bryden, 1982; Moscovitch, 1979; Rhodes, 1993) or have brain lesions (Benton, 1980; Farah, 1990; Milner, 1980).

The most impressive neurological evidence for modularity, however, comes from recent electrophysiological and neuroimaging studies in humans. Recording from scalp electrodes, Jeffreys and Tukmachi (1992) found an early positive-evoked response at a latency of about 190 msec (P 190) to faces and facelike arrangements of objects. Responses evoked by nonface objects had a similar scalp distribution but were smaller and usually later, whereas responses to inverted faces were simply delayed. More recent studies by Bentin et al. (1996) found an earlier negative potential, N 170, which they believe is distinguishable from the P 190 both by its scalp distribution and the properties of the stimulus that evokes it. The N 170 is larger over the posterior temporal region and more pronounced on the right hemisphere. It is evoked by human faces but not by hands, animal faces, or other inanimate or animate objects. However, unlike the P 190, which closely reflects behavioral sensitivity to inversion and configural information, the N 170 is also evoked by inverted faces and by facial features, especially the eyes, either in isolation or in a mixed collection of features.

Face-specific evoked responses have also been obtained from subdural, chronically implanted electrodes. A peak negative responses, N 200, was found over discrete regions of the inferior occipito-temporal cortex when subjects viewed faces but not when they viewed scrambled faces, letter strings, animal or cars (Allison et al., 1994a, 1994b; Puce et al., 1995; Puce, Allison, Asgari, Gore, & McCarthy, 1996). The resemblance in latency and response characteristics between the subdural N 200 potentials and Jeffreys and Tukmachi's (1992) scalp-recorded P 190 potentials suggest that they are generated by common structures in the inferotemporal cortex.

The presumed location of these generators coincide with regions of activation to faces that have been reported in a number of PET and fMRI studies (Grady et al., 1995; Haxby et al., 1994; Kanwisher et al., 1996a, 1996b, in press; Puce et al., 1995; Sergent et al., 1992). In all cases, the fusiform gyrus, particularly on the right, has been implicated. When the task requires that the faces be identified semantically rather than simply perceived perceptually, activation is also observed in more anterior regions of the inferotemporal cortex bordering on the parahippocampal gyrus on the right (Sergent et al., 1992) and on the left if naming is also involved (Sergent, MacDonald, & Zuck, 1994).

Although Sergent et al. (1992) reported a left-right difference in activation between objects and faces, with the left being favored for objects, most PET and fMRI studies have not attempted to distinguish clearly between areas involved in recognition of objects

as compared to faces. Because regions activated by objects (Kohler, Kapur, Moscovitch, Winocur, & Houle, 1995; Malach, Reppas, Benson, Kwong, Jiang, Kennedy, Ledden, Brady, Rosen, & Tootell, 1995; Moscovitch, Kapur, Kohler, & Houle, 1995) often partially overlap those activated by faces, these results are open to the interpretation that faces and objects are processed by a common region concerned with the perception of complex visual patterns.

This interpretation has been refuted recently by a number of fMRI studies that have confirmed the degree of specificity in brain organization that earlier studies hinted at. Kanwisher et al. (1996a, 1996b, in press) and McCarthy et al. (in press) found that a circumscribed region in the right lateral fusiform gyrus is preferentially activated to faces as compared to scrambled faces and animate or inanimate objects even when within category (subordinate level) discriminations are required. It has yet to be determined, however, whether this region is sensitive to upright faces that retain the integrity of spatial relations among facial features or whether this is a region that is sensitive to detection of facial features, particularly the eyes, as in Bentin et al.'s (1996) study.

Components of a Face-Recognition System

Taken together, the results of the electrophysiological and functional neuroimaging studies suggest that there is a distinct neural system dedicated to face recognition. The system consists of three interrelated components: (1) a region in the right occipital-temporal sulcus that is sensitive to facial features, particularly the eyes, (2) a region in the right lateral fusiform gyrus that is sensitive to configurational (second-order relational) features, and (3) an anterior region on the border of the fusiform and parahippocampal gyri on the right that is sensitive to semantic/physiognomic aspects of facial identity, and on the left, that is sensitive to names. This system must also receive input from adjacent, object-recognition regions that contribute to face recognition when face stimuli lack the domain-specific information needed to activate the face system directly (see also Tovee & Cohen-Tovee, 1993).

The infero-temporal cortex and the occipito-temporal junction (lingual gyrus) is also the site where investigators have observed selective impairment in object recognition and reading following brain lesions (McCarthy & Warrington, 1990) and selective activation to objects and to words on neurophysiological and neuroimaging studies. Although there is some overlap among the regions, the precise location of each differs from one another and from that observed for faces (Allison et al., 1994a, 1994b; Kanwisher et al., 1996a, 1996b, in press; Nobre, Allison, & McCarthy, 1994; McCarthy et al., in press; Puce et al., 1995). Such findings reinforce the view that the infero-temporal cortex is a region where modules for perception of complex visual patterns can develop or be created (Jacobs & Jordan, 1992). It likely holds this privileged position by virtue of its location at the interface between the primary visual cortex and early extrastriate cortex involved in feature and shape analysis on the one hand, and the multimodal association cortex in the lateral and medial temporal lobe involved in semantic processing and memory on the other (Gross, 1992; Gross et al., 1993).

The proximity of these various modules to one another explains why damage to this region typically leads to deficits in all aspects of higher-order visual pattern recognition and why it is only in rare circumstances, such as CK's, that damage is highly selective. By taking advantage of such rare cases, neuropsychologists have been able to delineate the functional and structural characteristics of these various systems and thereby provide the theoretical framework and empirical base from which cognitive neuroscientists operate.

Conclusion

The strategy most often employed by neuropsychologists in their investigations has been to study the defective function that results from circumscribed lesions. We hope that our study with CK has shown that there is much to be gained by studying the function that is spared. By investigating his preserved face recognition in relative isolation from the influence of processes involved in object recognition and reading, which are defective, we have been able to distinguish those aspects of face recognition that are mediated by a face module from those that are dependent on input from other systems. Our goal now is to characterize more precisely the processes underlying face recognition and to understand more fully how the different systems interact.

METHOD

Experiment 1

Participants were shown 140 photographs of famous people. All of them were cut out of magazines, mounted in plastic to preserve them, and inserted in two binders (Sets A and B), each containing 70 faces. As much as possible, photos were chosen in which identifying nonface features were absent. For Set A, faces were first shown one at a time in the normal, upright orientation and a few months later they were shown in an inverted orientation or as disguised (see Experiment 9 below). The order was reversed for Set B and the lag until the upright orientation was tested was only minutes. Recognition of faces in each binder was tested 2 years apart. Participants were given 10 sec to identify the picture. A response was scored correct if the name or some identifying information was supplied.

Experiment 2

There were 50 photos in all with a range of 3 to 6 photos per person. The photos were presented in order of typicality, the most typical being shown last. For each person who had to be identified, the points awarded depended on how early in the sequence of typicality the person was identified and corresponded to the position the photo occupied in the sequence. Thus, if there was a sequence of six photos including the target, and the person was identified after the third photo from the end, the participant received four points [ILLUSTRATION FOR FIGURE 1 OMITTED].

Experiment 3

The stimuli were taken from a color photo spread in Time magazine in which faces of seven males from different nationalities were combined in sequence with those of seven females to produce 49 offspring who had 50% of their features from each parent. On each trial, a photo of an offspring was presented and the participant had to choose from an array of the seven male and seven female faces which ones were combined to form the offspring [ILLUSTRATION FOR FIGURE 2 OMITTED]. A point was awarded for each parent who was identified correctly.

Experiment 4

In each display, a single front-view photograph of a face is presented above a display of six faces. On the first 6 trials, the six possible choices are all front-view photographs and the participant must choose the single corresponding item. On a further 24 trials, the six choices are three-quarter view photos. This condition ensures that the identification is not done on a simple matching of features in the photographs. On the remaining 24 trials, the six choices include the target under three different lighting conditions together with three distracter items. Male and female faces appear equally often.

The target is presented with the six choices for an unlimited exposure duration. On the first 30 trials, the participant matches the single target with the counterpart from the array, and on the remaining 24 trials, the participant identifies all three choices in the array that correspond to the target. A point is awarded for such correct response.

Experiment 5

The control participants were 18 young (19 to 23 years old) and 18 older (65 to 72 years old) adults who were part of another study on memory. Sixteen photographs taken from Moran, Kimble, and Mefferd (1960) were shown to participants in sets of four faces at a time They were instructed to study the faces so as to recognize them on a subsequent memory test. Each set was exposed for 16 seconds. Immediately after the last set was studied, an array consisting of 32 faces was presented and the participant's task was to choose the 16 faces that were studied. Twenty minutes later, the same array of 32 faces was presented, and the participant again had to choose the faces that were studied from among the distracters. This procedure was repeated three times with a different set of targets and distracters each time.

Experiment 6

We presented CK and controls with a total of 29 simple and detailed caricatures of famous people and had them try to identify them as in previous experiments. If they failed to identify a caricature, we then presented the target name along with three distracters, one phonemically related, one semantically related, and one unrelated [ILLUSTRATION FOR FIGURE 3 OMITTED]. The latter manipulation was uninformative because participants rarely confused the caricature with the lures. Therefore, no more will be said about it. Last, we presented a photo of the person for

those caricatures that a participant could not identify to ascertain that it is caricature identification, and not person identification, that we were measuring.

Experiment 7

We presented CK and controls with 31 cartoons of faces of characters that we clipped from magazines and children's books. Each cartoon was presented individually on a page and participants were given 10 sec to identify the cartoon. It was not necessary to be accurate in naming the character, but it was necessary to provide strong identifying information (e.g., for Minnie Mouse, saying it was Mickey's wife or girl friend would suffice). The cartoons were presented after participants attempted to identify each of them when they were presented in inverted orientation.

Experiment 8

We used virtually the same procedure described by Carey et al. (submitted, Experiment 3). CK was read the names of 15 famous men. After each one, he was asked to imagine the man's face as clearly as possible. Once all 15 names were presented, we then showed CK 60 line drawings comprising four different line drawings of each person: veridical, caricature, anti-caricature, and lateral caricature. The drawings were presented in a predetermined but counterbalanced order so that, within a set of 15 trials, each face appeared once and the type of deviation varied from one face to the next in the order veridical, caricature, anti-caricature, and lateral. Thus, in each of the four sets of fifteen, no face was repeated, and across sets, each type of distortion for each face appeared only once. CK's performance was compared with that of controls reported in Carey et al. (submitted). After completing the caricature condition, CK was shown photos of the individuals who were represented as caricatures. He identified all the photos correctly.

The drawings were created by using Brennan's (1985) program for generating caricatures based on deviations of features of each individual face from the norm, which was created by combining features of all faces used in the experiment. By comparing the veridical drawing with the norm, the program generated caricatures by increasing the difference between the two by a proportion of 0.50 along the same vector. Anti-caricatures were created by decreasing the distance by 0.50 along the same vector, and lateral caricatures were created by increasing the distance by 0.50 along a vector at right angles to that vector (for details see Carey et al., submitted).

Experiment 9

We presented CK and 12 controls with two sets, A and B, each consisting of 70 photos of famous people. The photos in each set were first presented in two blocks of 35, each block consisting of upright or disguised faces. For each face, there were two simultaneous disguises (e.g., mustache and glasses) that could be removed one at a time if identification was not possible with both disguises present [ILLUSTRATION FOR FIGURE 5 OMITTED]. The blocks (upright and disguised) in both sets were counterbalanced across subjects. After the photos were viewed in the two conditions, the

same photos, in each set, were shown as upright and undisguised to see if the people were indeed familiar to participants. For CK, Set B was presented in the same order as for controls, whereas Set A was presented in the reverse order with upright faces shown first and the disguised and inverted faces shown a few months later. Participants were given about 5 sec to identify each face. As before, a response was considered correct if participants supplied the name or proper identifying information about the person. In the disguised condition, the scores reported were for identification when both disguises were present.

Experiment 10

Participants viewed 31 inverted cartoon faces. These faces were the same as the ones used in Experiment 7, but in that experiment they were viewed first in the inverted orientation before they were seen upright. Otherwise, the procedure was the same as that in Experiment 7.

Experiment 11

The faces were derived from the same large set of high school graduation photos created by Moran et al. (1960), a different subset of which was used in Experiment 5. Participants were shown a set of 16 target photos, all displayed on one card and below them was a test card consisting of 32 photos that included the targets plus 16 lures. Participants were given unlimited time to select the targets. Moran et al. had designed similar sets of photos of the heads of different species of dogs and cats and these were presented for perceptual matching in the same manner as the human faces. For the human faces, there were three matching conditions: target and test upright, target and test inverted, and target inverted and test upright. Different faces appeared in each condition. Because we had only two different sets of photos of animal heads, we tested them only in the upright-upright and inverted-upright condition.

Experiment 12

Participants were shown two sets, one consisting of 11, and the other of 12, color photos of famous people. The internal parts of the face consisting of the eye, nose, and mouth were cut out as a unit and inverted. To achieve this a line was drawn to the uppermost part of each eye and extended to the outside (temporal) portion of the eye. A line equal and parallel to it was drawn through the bottom of the mouth (lower lip). The two lines were joined and the rectangle formed by that was cut out. Thus the internal parts of the face retained the spatial relation to each other but not to the external contour [ILLUSTRATION FOR FIGURE 7 OMITTED]. Subjects viewed first one set of photos with the internal feature inverted and the external upright and vice versa for the second set. After viewing the distorted photos, the full intact photo was presented upright for identification. The identification criteria and the time allotted was the same as in previous experiments.

Experiment 13

We used line drawings of caricatures of famous people because we already knew that CK could identify caricatures normally. We presented nine sets of caricatures, each consisting of three overlapping figures, which participants were given 10 sec to identify as in the previous experiments. If identification was not perfect, participants were then required to choose the three targets from a set of five individually separate figures that were presented alongside the overlapping set [ILLUSTRATION FOR FIGURE 8 OMITTED]. Thus, in total there were 27 overlapping caricatures and a recognition set consisting of an additional 18 lures.

Experiment 14

From a standard test booklet containing 50 pictures (Mooney Closure Test), we selected all 7 faces and another 9 objects that pilot testing indicated were of comparable difficulty. Participants viewed each of them for 10 sec and tried to identify them. For the faces, they were also required to indicate the age and sex of the depicted individual and the direction in which the person was facing. They were scored correct only if they could provide all the information.

Experiment 15

Participants viewed 40 fractured faces of famous people, one at a time, and were given 10 sec to identify each one as in the previous experiments. After viewing the fractured faces, participants saw the same faces intact for identification.

Experiment 16

Participants viewed 20 new faces of famous people. In half, the left side of the face was misaligned vertically with the right half along a line from the top to the tip of the nose so that the eye on one side was aligned with the mouth on the other [ILLUSTRATION FOR FIGURE 11 OMITTED]. In the other 10 pictures, the face was misaligned horizontally along a line through the middle of the nose so that the nose at the top half was aligned with the edge of the face in the bottom half. Participants were required to identify each misaligned face as in previous experiments. After viewing the misaligned faces, participants were shown the intact faces for identification.

Experiment 17

Fifteen black and white photos of famous faces were scanned into a Macintosh computer. For a third of them, the eyes were removed; for a third, the nose; and for a third, the mouth [ILLUSTRATION FOR FIGURE 12 OMITTED]. Participants viewed each of these incomplete faces in random order and attempted to identify them as in the other experiments. After viewing each face, participants were shown the missing part along with a lure that fit as well into the empty space as the correct part did. They had to select the correct part. After completing both these portions of the task, the intact face was shown for identification.

Experiment 18

Participants were shown composite faces taken from the paintings of Arcimbaldo. Each face consisted either of fruit (Rudolfo), flowers (Primavera, Flora), books (Biblico), trees (Inverno), vegetables (Natura), animals (Terra), or landscape with houses, trees, and streams (Anthro). All participants were asked simply to describe what they saw. If they saw a face, they were asked to provide information about the type of face it was, its emotion, sex, and other identifying features. If they also saw the objects, they were asked to identify each object as best they could. If they claimed to see only one or the other, they were prompted to describe as much about the picture as they could ("Tell me everything you see").

Experiment 19

CK and 12 control participants were presented with the picture "Faces in the Forest" by Beverely Doolittle [ILLUSTRATION FOR FIGURE 14 OMITTED]. They were asked to examine the picture and to indicate all the faces that they could detect in 5 min by pointing to their location.

Notes

1. After this manuscript had been typeset, it was brought to our attention that the lateral caricatures used in Carey et al.'s (submitted; see also Carey, 1992, p. 99) study and in ours were incorrectly scaled to the norm before executing the "lateral" manipulation which produced the caricature (Rhodes, personal communication, August 1997). When the procedure was corrected, the resulting new lateral caricatures were recognized better than the anti-caricatures but still not as well as the true caricatures and veridical drawings (Rhodes, personal communication, August 1997). In light of these developments, the

conclusion we draw in favor of the norm-based hypothesis (which is based on the old lateral caricatures), lacks support and may not be valid. Therefore, the issue is still not settled as to which hypothesis, norm-based coding or density alone (or noise), best accounts for the data.

It is important to note that these considerations do not invalidate any of the other conclusions regarding face recognition that emerge from our study.

2. Preliminary data indicate that a deviation between 45 to 90 from the upright is sufficient to prevent a face from engaging face-recognition mechanisms. This applies also to rudimentary faces with highly salient features such as are depicted in cartoons (Experiment 10).

3. The peripheral-central dichotomy provides a short-hand description of the nature of the agnosia and roughly mirrors the original Lissauer (1890) classification. The terms are operationally defined to refer to deficits that affect either the derivation of a coherent structural description (peripheral agnosia) or the assignment of meaning to a well-structured percept (central agnosia). The terms do not map precisely onto underlying neuroanatomical substrates although they approximately follow the axis from posterior occipital into more ventral temporai regions. More recent and fine-grained nomenclatures which consider a range of deficits along a continuum are also available (for discussion see Farah, 1990; Humphreys et al., 1994), but for ease of description, a binary classification is often used. We think however, that CK's deficit is best classified as an integrative agnosia (Riddoch & Humphreys, 1987) in which intermediate levels of visual processing are impaired.

4. Hirschfeld provides a celebrated example of this type of game in his weekly New York Times caricatures of actors and actresses. By including the name Nina followed by a number in each caricature, he challenges the observer to find those number of instances in which he has hidden the name Nina in the caricature.

REFERENCES

Allison, T., Ginter, H., McCarthy, G., Nobre, A. A., Puce, A., Luby, M., & Spencer, D. D. (1994a). Face recognition in human extrastriate cortex. Journal of Neurophysiology, 71, 821-825.

Allison, T., McCarthy, G., Nobre, A., Puce, A., & Belger, A. (1994b). Human extrastriate visual cortex and the perception of faces, words, numbers, and colors. Cerebral Cortex, 5, 544-554.

Assal, G., Favre, C., & Anderes, J. P. (1984). Non-reconnaisance d'animaux familiers chez un payson: Zoo-agnosie ou prosopagnosie pour les animaux. Revue Neurologique, 140, 580-584.

Bartlett, J. C., Hurry, S., & Thorley, W (1984). Typicality and familiarity of faces. Memory and Cognition, 12, 219-228.

Bartlett, J. C., & Searcy, J. (1993). Inversion and configuration of faces. Cognitive Psychology, 25, 281-316.

Bauer, R. M. (1984). Autonomic recognition of names and faces in prosopagnosia: A neuropsychological application of the guilty knowledge test. Neuropsychologia, 23, 457-469.

Behrmann, M., Moscovitch, M., & Winocur, G. (1994). Intact visual imagery and impaired visual perception in a patient with visual agnosia. Journal of Experimental Psychology: Human Perception and Performance, 30(5), 1068-1087.

Behrmann, M., Winocur, G., & Moscovitch, M. (1992). Dissociations between mental imagery and object recognition in a brain-damaged patient. Nature, 359, 636-637.

Benson, P. J., & Perrett, D. I. (1991). Perception and recognition of photographic quality caricatures: Implications for the recognition of natural images. European Journal of Cognitive Psychology, 3, 105-135.

Benson, P. J., & Perrett, D. I. (1994). Visual processing of facial distinctiveness. Perception, 23, 75-93.

Bentin, S., Allison, T., Puce, A., & McCarthy, G. (1996). Electrophysiological studies of face perception in humans. Journal of Cognitive Neuroscience, 8, 551-565.

Benton, A. L. (1980). The neuropsychology of facial recognition. American Psychologist, 35, 176-186.

Benton, A. L., Hamsher, K. des, Varney, N. R., & Spreen, O. (1978). Facial recognition: Stimulus and multiple choice pictures. New York: Oxford University Press.

Benton, A. L., & van Allen, M. W (1973). Test of facial recognition. Neurosensory Center Publication No. 287. Iowa City, IA: University Hospitals, Department of Neurology.

Beymer, D., & Poggio, T. (1996). Image representation for visual learning. Science, 272, 1905-1909.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. Psychological Review, 94, 115-147.

Bodamer, J. (1947). Die Prosop-agnosie. Archiv fur Psychiatrie und Nervenkrankheiten, 179, 6-53.

Bobes, M. A., Valdes-Sosa, M., & Olivares, E. (1994). An ERP study of expectancy violation in face perception. Brain and Cognition, 26, 1-22.

Bradshaw, J. L., & Nettleton, N. C. (1981). The nature of hemispheric specialization in man. The Behavioural and Brain Sciences, 4, 51-92.

Bradshaw, J. L., & Wallace, G. (1971). Models for the processing and identification of faces. Perception and Psychophysics, 9, 443-448.

Brennan, S. E. (1985). The caricature generator. Leonardo, 18, 170-178.

Bruce, C. J., Desimone, R., & Gross, C. G. (1981). Visual properties of neurons in a polysensory area in superior temporai sulcus of macaque. Journal of Neurophysiology, 46, 369-384.

Bruce, V. (1988). Recognizing faces. London: Erlbaum.

Bruce, V., & Humphreys, G. (1994). Recognizing objects and faces. Visual Cognition, 1, 141-180.

Bruce, V., & Young, A. (1986). Understanding face recognition. British Journal of Psychology, 77, 305-327.

Bruyer, R. (Ed.) (1986). The neuropsychology of face perception and facial expression. Hillsdale, NJ: Erlbaum.

Bruyer, R. (1991). Covert face recognition in prosopagnosia: A review. Brain and Cognition, 15, 223-235.

Bruyer, R., & Crispeels, G. (1992). Expertise in person recognition. Bulletin of the Psychonomic Society, 30, 501-504.

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

Bryden, M.P. (1982). Laterality: Functional asymmetry in the intact brain. New York: Academic Press.

Carey, S., & Diamond, R. (1994). Are faces perceived as configurations more by adults than by children? Visual Cognition, 1(2/3), 253-274.

Carey, S. Rhodes, G., Diamond, R., & Hamilton, J. (submitted). Norm-based coding of faces: Evidence from studies of caricature recognition.

Cohen, M. E., & Carr, W. J. (1975). Facial recognition and the von Restorff effect. Bulletin of the Psychonomic Society, 6, 383-384.

Corballis, M. C. (1991). The lopsided ape: The evolution of the generative mind. New York: Oxford University Press.

Damasio, A. R. (1990). Category-related recognition defects as a clue to the neural substrates of knowledge. Trends in Neuroscience, 13, 95-98.

Damasio, A. R., Damasio, H., & Van Hoesen, G. W (1982). Prosopagnosia: Anatomic basis and behavioral mechanisms. Neurology, 32, 331-341.

Damasio, A. R., Tranel, D., & Damasio, H. (1990). Face agnosia and the neural substrates of memory. Annual Review of Neuroscience, 13, 89-109.

de Haan, E. H. E, & Hay, D. C. (1986). The matching of famous and unknown faces given either the internal or the external features: A study on patients with unilateral brain lesions. In H. D. Ellis, M. A. Jeeves, F. Newcombe, & A. W Young (Eds.), Aspects of face processing (pp. 302-309). Dordrecht: Martinus Nijhoff.

de Haan, E. H. E, Young, A. W, & Newcombe, F. (1987). Face recognition without awareness. Cognitive Neuropsychology, 4, 389-419.

De Renzi, E. (1986a). Prosopagnosia in two patients with CT scan evidence of damage confined to the right hemisphere. Neuropsychologia, 24, 385-389.

De Renzi, E. (1986b). Current issues in prosopagnosia. In H. D. Ellis, M. A. Jeeves, F. Newcombe, & A. W. Young (Eds.), Aspects of face processing (pp. 243-252). Dordrecht: Martinus Nijhoff.

De Renzi, E., Faglioni, P, Grossi, D., & Nichelli, P. (1991). Apperceptive and associative forms of prosopagnosia. Cortex, 27, 213-221.

De Renzi, E., Perani, D., Cartesimo, G. A., Silveri, M. C., & Fazio, F. (1994). Prosopagnosia can be associated with damage confined to the right hemisphere - An MRI and PET study and a review of the literature. Neuropsychologia, 32, 893-902.

Desimone, R. (1991). Face-selective cells in the temporal cortex of monkeys. Journal of Cognitive Neuroscience, 3, 1-8.

Desimone, R., Albright, T. D., Gross, C. G., & Bruce, C. (1984). Stimulus selective properties of inferior temporal neurons in the macaque. Journal of Neuroscience, 4, 2051-2062.

Diamond, R., & Carey, S. (1986). Why faces are and are not special: An effect of expertise. Journal of Experimental Psychology: General, 115, 107-117.

Ellis, H. D., Jeeves, M. A., Newcombe, F., & Young, A. W. (Eds.) (1986). Aspects of face processing. Dordrecht: Martinus Nijhoff.

Ellis, H. D., Shepherd, J. W., & Davies, G. M. (1979). Identification of familiar and unfamiliar faces from internal and external features: Source implication for theories of face recognition. Perception, 8, 431-439.

Farah, M. J. (1990). Visual agnosia: Disorders of object recognition and what they tell us about normal vision. Cambridge, MA: MIT Press.

Farah, M. J. (1991). Patterns of cooccurrence among the associative agnosias: Implications for visual object representation. Cognitive Neuropsychology, 8, 1-19.

Farah, M. J. (1995). A dissociable system for visual recognition: A cognitive neuropsychology approach (pp. 101119). In S. M. Kosslyn & D. N. Osherson (Eds.), Invitation to cognitive science: Visual cognition. Cambridge, MA: MIT Press.

Farah, M. J., Levinson, K. L., & Klein, K. L. (1995). Face perception and within-category discrimination in prosopagnosia. Neuropsychologia, 33, 661-674.

Farah, M. J., Tanaka, J. W., & Drain, H. M. (1995). What causes the face inversion effect? Journal of Experimental Psychology: Human Perception and Performance, 21, 1-7.

Farah, M. J., Wilson, K. D., Drain, H. M., &Tanaka, J. R. (1995). The inverted face inversion effect in prosopagnosia: Evidence for mandatory, face-specific perceptual mechanisms. Vision Research. 35, 2089-2093.

Farah, M. J., Wilson, K. D., Drain, H. M., Tanaka, J. R. (submitted). What is "special" about face perception?

Feinberg, T. E., Schindler, R. J., Ochoa, E., Kwan, P. C., & Farah, M. J. (1994). Associative visual agnosia and alexia without prosopagnosia. Cortex, 30, 395-412.

Fodor, J. (1983). The modularity of mind. Cambridge, MA: MIT Press.

Fodor, J. (1985). Precis of The modularity of mind. Behavioral & Brain Sciences, 8, 1-42.

Garner, W. R. (1974). The processing of information and structure. Hillsdale, NJ: Erlbaum.

Gauthier, I., & Tarr, M. (1997). Becoming a "greeble" expert: Exploring the face recognition mechanism. Vision Research, 37, 1673-1682.

Gauthier, I., Tarr, M.J., Anderson, A., & Gore, J. (1997) Expertise training with novel objects can recruit the fusiform face area. Society for Neuroscience Abstract. New Orleans, LA.

Gazzaniga, M. S. (1997). Pieces of Mind. Scientific American Frontier, Connecticut Public Television/Scientific American Magazine (narrator, Alan Alda).

Going, M., & Read, J. E. (1974). Effects of uniqueness, sex of subject, and sex of photograph on facial recognition. Perceptual and Motor Skills, 39, 109-110.

Grady, C. L., Mcintosh, A. R., Horwitz, B., Maisog, J. M., Ungerleider, L. G., Mentis, M. J., Pietrini, P., Schapiro, M. B., & Haxby, J. V. (1995). Age-related reductions in human recognition memory due to impaired encoding. Science, 269, 218-221.

Gross, C.G. (1973). Visual functions of inferotemporal cortex. In V. Jung (Ed.), Handbook of sensory physiology, VII/3. Berlin: Springer.

Gross, C. G. (1992). Representation of visual stimuli in inferior temporal cortex. Philosophical Transactions of the Royal Society of London, series B, 335, 3-10.

Gross, C. G., Rocha-Miranda, C. E., & Bender, D. B. (1972). Visual properties of neurons in inferotemporal cortex of the macaque. Journal of Neurophysiology, 35, 96-111.

Gross, C. G., Rodman, H. R., Gochin, P M., & Colombo, M. W. (1993). Inferior temporal cortex as a pattern recognition device. In E. Baum (Ed.), Proceedings of the Third NEC Conference Research Symposium, Computational Learning and Cognition (pp. 44-73). Philadelphia: SIAM.

Gross, C. G., & Sergent, J. (1992). Face recognition. Current Opinion in Neurobiology, 2, 156-161.

Hamm, J. P., & McMullen, A. (1996). Effects of orientation on the identification of rotation objects depend on the level of identity. Journal of Experimental Psychology: Human Perception and Performance. In press.

Hancock, P. J. B., Burton, A.M., & Bruce, V. (1996). Face processing: Human perception and principal component analysis. Memory & Cognition, 24, 26-40.

Harman, K., & Moscovitch, M. (1996). The inversion effect for faces and objects: Part-based/holistic encoding and the effects of format congruence. Paper presented at the Annual Meeting of the Canadian Society for Brain Behavior and Cognitive Science. Montreal, Quebec, Canada.

Harmon, L. D. (1973). The recognition of faces. Scientific American, 227, 71-82.

Haxby, J. V., Grady, C. L., Horwitz, B., Ungerleider, L. G., Mishkin, M., Carson, R. E., Herscovitch, P., Schapiro, M. B., & Rapaport, S. I. (1991). Dissociations of spatial and object visual processing pathways in human extrastriate cortex. Proceedings of the National Academy of Sciences, USA, 88, 1621-1625.

Haxby, J. V., Horwitz, B., Ungerleider, L. G., Maisog, J. M., Pietrini, P., & Grady, C. L. (1994). The functional organization of human extrastriate cortex: A PET rCBF study of selective attention to faces and locations. Journal of Neuroscience, 14, 6336-6353.

Hecaen, H., & Angelergues, R. (1962). Agnosia for faces (prosopagnosia). Archives of Neurology, 7, 24-32.

Hecaen, H., Goldblum, M. C., Masure, M. C., & Ramier, A.M. (1974). Une nouvelle observation d'agnosie d'objet: Deficit de l'association ou de la categorisation, specifique de la modalite visuelle. Neuropsychologia, 12, 447-464.

Heywood, C. A., & Cowey, A. (1992). The role of the "face-cell" area in the discrimination and recognition of faces by monkeys. Philosophical Transactions of the Royal Society of London (B), 335, 31-38.

Humphreys, G. W., Riddoch, M. J., Donnelly, N., Freeman, T., Boucart, M., & Muller, H. M. (1994). Intermediate visual processing and visual agnosia. In M. J. Farah & G. Ratcliff (Eds.), The neuropsychology of high-level vision (pp. 63101). Hillsdale, NJ: Erlbaum.

Jackson, J. H. (1874, 1915). On the nature of the duality of the brain. Medical Press and Circular, 1, 19, and Brain, 38, 80-103 (1915).

Jacobs, R. A., & Jordan, M. I. (1992). Computational consequences of a bias toward short connections. Journal of Cognitive Neuroscience, 4, 323-336.

Jeffreys, D. A. (1993). The influence of stimulus orientation on the vertex positive scalp potential evoked by faces. Experimental Brain Research, 91, 340-350.

Jeffreys, D. A., & Tukmachi, E. S. A. (1992). The vertex-positive scalp potential evoked by faces and objects. Experimental Brain Research, 96, 163-172.

Johnston, R. A., Milne, A. B., Williams, C., & Hosie, J. (1997). Do distinctive faces come from outer space? An investigation of the status of multidimensional face-space. Visual Cognition, 4, 59-67.

Jolicoeur, P. (1985). The time to name disoriented natural objects. Memory and Cognition, 13, 289-303.

Kanwisher, N., Chun, M M., McDermott, J., & Ledden, P. J. (1996a). Functional imaging of human visual recognition. Cognitive Brain Research, 5, 55-67.

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

Kanwisher, N., Woods, R. P., Iacoboni, M., & Mazziota, J. C. (1996b). A locus in human extrastriate cortex for visual shape analysis. Journal of Cognitive Neuroscience, 9, 133-142.

Kohler, S., Kapur, S., Moscovitch, M., Winocur, G., & Houle, S. (1995). Dissociation of pathways for object and spatial vision: A PET study in humans. Neuroreport, 6, 1865-1868.

Landis, T., Cummings, J. L., Christen, L., Bogen, J. E., & Imhof, H. G. (1986). Are unilateral right posterior cerebral lesions sufficient to cause prosopagnosia? Clinical and radiological findings in six additional patients. Cortex, 22, 243-252.

Levy-Agresti, J., & Sperry, R. W. (1968). Differential perceptual capacities in major and minor hemispheres. Proceedings of the National Academy of Sciences, USA, 61, 1151.

Light, L. L., Kayra-Stuart, F., & Hollander, S. (1979). Recognition memory for typical and unusual faces. Journal of Experimental Psychology: Human Learning and Memory, 5, 212-219.

Lissauer, H. (1890). Ein fall von seelenblindheit nebst einem beltrag zur theorie derselben. Archiv fur Psychiatrie, 21, 222-270. [Edited and reprinted in translation by Jackson, M. (1988). Lissauer on agnosia. Cognitive Neuropsychology, 5, 155-192.]

Logothetis, N., & Sheinberg, D. L. (1996). Visual object recognition. Annual Review of Neuroscience, 19, 577-621.

Malach, R., Reppas, J. B., Benson, R. R., Kwong, K. K., Jiang, H., Kennedy, W. A., Ledden, P. J., Brady, T. J., Rosen, B. R., & Tootell, R. B. H. (1995). Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. Proceedings of the National Academy of Sciences, USA, 92, 8135-8139.

McCarthy, G., Puce, A., Gore, J. C., & Allison, T. (in press). Face-specific processing in the human fusiform gyrus. Journal of Cognitive Neuroscience, 9, 605-610.

McCarthy, R. A., & Warrington, E. K. (1990). Cognitive neuropsychology: A clinical introduction. San Diego, CA: Academic Press.

McMullen, P. A., Fisk, J. D., & Phillips, S. J. (submitted). Spared visual face processing in a severely appreciative visual object agnosic.

McNeil, J., & Warrington, E. K. (1991). Prosopagnosia: A reclassification. Quarterly Journal of Experimental Psychology, 43A, 267-287.

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

Meadows, J. C. (1974). The anatomical basis of prosopagnosia. Journal of Neurology, Neurosurgery, and Psychiatry, 37, 489-501.

Michel, F., Poncet, M., & Signoret, J.-L. (1989). Les lesions responsables de la prosopagnosie sont-elles toujours bilaterales? Revue Neurologique, 146, 764-770.

Milner, B. (1980). Complementary functional specialization of the human cerebral hemispheres. In R. Levy-Montealeini (Ed.), Neurons, transmitters and behavior. Vatican City: Pontificiae Academiae Scientiarum Scripta Varia.

Mooney, C. M. (1956). Closure with negative after images under filtering light. Canadian Journal of Psychology, 10, 191-199.

Moran, L.J., Kimble, J. P., Jr., & Mefferd, R. B., Jr. (1960). Repetitive psychometric measures: Memory for faces. Psychological Reports, 7, 407-413. (Repetitive psychodiagnostic measures: Social memory. Hogg Foundation for Mental Health. University of Texas).

Moscovitch, M. (1979). Information processing and cerebral hemispheres. In M. S. Gazzaniga (Ed.), Handbook of behavioral neurobiology, Vol. 2 (pp. 379-446). New York: Plenum Press.

Moscovitch, M., Kapur, S., Kohler, S., & Houle, S. (1995). Distinct neural correlates of visual long-term memory for spatial location and object identity. Proceedings of the National Academy of Sciences USA, 92, 3721-3725.

Moscovitch, M., & Radzins, M. (1987). Backward masking of lateralized faces by noise, pattern, and spatial frequency. Brain and Cognition, 6, 72-90.

Moscovitch, M, Scullion, D., & Christie, D. (1976). Early vs. late stages of processing and their relation to functional hemispheric asymmetries in face recognition. Journal of Experimental Psychology: Human Perception and Performance, 2, 401-416.

Moscovitch, M., & Umilta, C. (1990). Modularity and neuropsychology: Modules and central processes in attention and memory. In M. Schwartz (Ed.), Modular deficits in Alzheimer's disease (pp. 1-59). Cambridge, MA: MIT Press/Bradford.

Nachson, I. (1995). On the modularity of face recognition: The riddle of domain specificity. Journal of Clinical and Experimental Neuropsychology, 17, 256-275.

Nachson, I., Moscovitch, M., & Umilta, C. (1995). The contribution of external and internal features to the matching of unfamiliar faces. Psychological Research, 58, 31-37.

Nakayama, K., Shimojo, S., & Silverman, G. H. (1989). Stereo-scopic depth: Its relation to image segmentation, grouping, and the recognition of occluded objects. Perception, 18, 55-68.

Navon, D.C. (1977). Forest before the trees: The precedence of global features in visual processing. Cognitive Psychology, 9, 353-383.

Newcombe, F., Mehta, Z., & de Haan, E. H. F. (1994). Category specificity in visual recognition. In M. J. Farah & G. Ratcliff (Eds.), The neuropsychology of high-level vision (pp. 103132). Hillsdale, NJ: Erlbaum.

Nobre, A. C., Allison, T., & McCarthy, G. (1994). Word recognition in the human inferior temporal lobe. Nature, 372, 260-263.

Norris, D. (1995). Signal detection theory and modularity: On being sensitive to the power of bias models of semantic priming. Journal of Experimental Psychology: Human Perception and Performance, 21, 935-939.

Palmer, S., & Rock, I. (1994a). Rethinking perceptual organization: The role of uniform connectedness. Psychonomic Bulletin & Review, 1, 29-55.

Palmer, S., & Rock, I. (1994b). On the nature and order of organizational processing: A reply to Peterson. Psychonomic Bulletin & Review, 1, 515-519.

Perrett, D. I., Hietanen, J. K., Oram, M. W., & Benson, P. J. (1992). Organization and functions of cells responsive to faces in the temporal cortex. Philosophical Transactions of the Royal Society of London, 335, 23-30.

Perrett, D. I., Rolls, E. T., & Caan, W. (1982). Visual neurones responsive to faces in the monkey temporal cortex. Experimental Brain Research, 47, 329-342.

Peterson, M. A. (1994a). Shape recognition can and does occur before figure-ground organization. Current Directions in Psychological Science, 3, 105-111.

Peterson, M. A. (1994b). The proper placement of uniform connectedness. Psychonomic Bulletin & Review, 1, 509-514.

Peterson, M. A., & Gibson, B. S. (1994a). Must shape recognition follow figure-ground organization: An assumption in peril. Psychological Science, 9, 253-259.

Peterson, M. A., & Gibson, B. S. (1994b). Object recognition contributions to figure-ground organization: Operations on outlines and subjective contours. Perception and Psychophysics, 56, 551-554.

Puce, A., Allison, T., Asgari, M., Gore, J. C., & McCarthy, G. (1996). Differential sensitivity of human visual cortex to faces, letterstrings and textures: A functional magnetic resonance imaging study. Journal of Neuroscience, 16, 5205-5215.

Puce, A., Allison, T., Gore, J. C., & McCarthy, G. (1995). Face-sensitive regions in human extrastriate cortex studied by functional MRI. Journal of Neurophysiology, 74, 1192-1199.

Renault, B., Signoret, J.-L., Debruille, B., Breton, F., & Bolger, F. (1989). Brain potentials reveal covert facial recognition in prosopagnosia. Neuropsychologia, 27, 905-912.

Rhodes, G. (1988). Looking at faces: First-order and second-order features as determinants of facial appearance. Perception, 17, 43-63.

Rhodes, G. (1993). Configural coding, expertise and the right hemisphere advantage for face recognition. Brain and Cognition, 22, 19-41.

Rhodes, G., Brake, S., & Atkinson, A. P (1993). What's lost in inverted faces? Cognition, 47, 25-57.

Rhodes, G., Brennan, S., & Carey, S. (1987). Identification and ratings of caricatures: Implications for mental representation of faces. Cognitive Psychology, 19, 473-497.

Rhodes, G., & McLean, I. G. (1990). Distinctiveness and expertise effects with homogeneous stimuli: Towards a model of configural coding. Perception, 19, 773-794.

Rhodes, G., & Tremewan, T. (1993). The Simon then Garfunkel effect: Semantic priming, sensitivity and modularity of face recognition. Cognitive Psychology, 25, 147-187.

Rhodes, G., & Tremewan, T. (1994). Understanding face recognition: Caricature effects, inversion and the homogeneity problem. Visual Cognition, 1(2/3), 275-311.

Riddoch, M. J., & Humphreys, G. W. (1987). A case of integrative visual agnosia. Brain, 110, 1431-1462.

Robertson, L. C., & Delis, D. C. (1986). "Part-whole" processing in unilateral brain damaged patients: Dysfunction of hierarchical organization. Neuropsychologia, 24, 363-370.

Rolls, E. T. (1992). Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas. Philosophical Transactions of the Royal Society of London (B), 335, 11-21.

Ross, P., & Turkewitz, G. (1981). Individual differences in cerebral asymmetries for facial recognition. Cortex, 17, 199-214.

Ross, P., & Turkewitz, G. (1982). Changes in hemispheric advantage in processing facial information with increasing stimulus familiarization. Cortex, 18, 489-499.

Searcy, J. H., &Bartlett, J. C. (1996). Inversion and processing of component and spatial relational information in faces. Journal of Experimental Psychology: Human Perception and Performance, 23, 904-915.

Sergent, J. (1984). An investigation into component and configurational processes underlying face perception. British Journal of Psychology, 75, 221-242.

Sergent, J., & Hellige, J. B. (1986). Role of input factors in visual-field asymmetries. Brain and Cognition, 5, 174-179.

Sergent, J., MacDonald, B., & Zuck, E. (1994). Structural and functional organization of knowledge about faces and proper names: A positron emission tomography study. In C. Umilta & M. Moscovitch (Eds.), Attention and performance XV: Conscious and nonconscious information processing (pp. 203-228). Cambridge, MA: MIT Press/Bradford.

Sergent, J., Ohta, S., MacDonald, B. (1992). Neuroanatomy of face and object processing. Brain, 115, 15-36.

Sergent, J., & Signoret, J. L. (1992a). Varieties of functional deficits in prosopagnosia. Cerebral Cortex, 2, 375-388.

Sergent, J., & Signoret, J. L. (1992b). Implicit access to knowledge from unrecognized faces in prosopagnosia. Cerebral Cortex, 2, 389-400.

Suzuki, S., & Cavanagh, P (1995). Facial organization blocks access to low-level features: An object inferiority effect. Journal of Experimental Psychology: Human Perception and Performance, 21, 901-913.

Suzuki, S., Peterson, M. A., & Moscovitch, M., & Behrmann, M. (1997). Viewpoint specificity in the identification of single objects is evident in control participants and very exaggerated in visual object agnosia. Paper presented at the Annual Meeting of the Society for Cognitive Neuroscience. Boston, March 1997.

Takane, Y., & Sergent, J. (1983). Multidimensional scaling models for reaction times and same-different judgements. Psychometrika, 48, 393-423.

Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. Quarterly Journal of Experimental Psychology, 46A, 225-245.

Tanaka, J. W., & Sengco, J. A. (in press). Parts and their configuration in face recognition. Memory and Cognition.

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

Tovee, M. J., & Cohen-Tovee, E.M. (1993). The neural substrates of face processing models: A review. Cognitive Neuropsychology, 10, 505-528.

Tranel, D., & Damasio, A. R. (1985). Knowledge without awareness: An autonomic index of facial recognition by prosopagnosics. Science, 228, 1453-1454.

Turk, M., & Pentland, A. (1991). Eigen faces for recognition. Journal of Cognitive Neuroscience, 3, 71-68.

Tversky, B. T., & Baratz, D. (1985). Memory for faces: Are caricatures better than photographs? Memory & Cognition, 13, 45-49.

Ungerleider, L. G. (1995). Functional brain imaging studies of cortical mechanisms for memory. Science, 270, 769-775.

Valentine, T. (1988). Upside-down faces: A review of the effects of inversion upon face recognition. British Journal of Psychology, 79, 471-491.

Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. The Quarterly Journal of Experimental Psychology, 43A, 161-204.

Valentine, T., & Bruce, V. (1986a). Recognizing familiar faces. The role of distinctiveness and familiarity. Canadian Journal of Psychology, 40, 300-305.

Valentine, T., & Bruce, V. (1986b). The effect of race, inversion, and encoding activity upon face recognition. Acta Psychologica, 61, 259-273.

Van Essen, D.C., Anderson, C. H., & Felleman, D. J. (1992). Information processing in the primate visual system: An integrated systems perspective. Science, 225, 419-423.

Warrington, E. K., & James, M. (1967). An experimental study of facial recognition in patients with unilateral cerebral lesions. Cortex, 3, 317-326.

Warrington, E. K., & Rabin, P, (1970). Perceptual matching in patients with cerebral lesions. Neuropsychologia, 8, 475-487.

Whiteley, A. M., & Warrington, E. K. (1977). Prosopagnosia: A clinical psychological and anatomical study of three patients. Journal of Neurology, Neurosurgery, and Psychiatry, 40, 395-403.

Wright, A. A., & Roberts, W. A. (1996). Monkey and human face perception: Inversion effects for human faces but not for monkey faces or scenes. Journal of Cognitive Neuroscience, 8, 278-290.

Yamane, S., Kaji, S., & Kawano, K. (1988). What facial features activate face neurons in the inferotemporal cortex of the monkey? Experimental Brain Research, 73, 209-214.

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

Yin, R. K. (1970). Face recognition by brain-injured patients: A dissociable ability? Neuropsychologia, 8, 395-402.

Young, A. W. (1994). Conscious and nonconscious recognition of familiar faces. In C. Umilta & M. Moscovitch (Eds.), Attention and performance XV: Conscious and nonconscious information processing (pp. 153-178). Cambridge, MA: MIT Press/Bradford.

Young, A. W., Hay, D.C., McWeeney, K. H., Flude, B. M., & Ellis, A. W. (1985). Matching familiar and unfamiliar faces on internal and external features. Perception, 14, 737-746.

Young, A. W., Hellawell, D., & Hay, D.C. (1987). Configurational information in face perception. Perception, 16, 747-759.

Young, M.P., & Yamane, S. (1992). Sparse population coding of faces in the inferotemporal cortex. Science, 256, 1327-1331.

Yuille, A. L. (1991). Deformable templates for face recognition. Journal of Cognitive Neuroscience, 3, 59-70.