Welford, A. T. (1952). The "psychological refractory period" and the timing of high speed performance—A review and a theory. *British Journal of Psychology* 43, 2–19.

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica* 41, 67–85.

Wolfe, J. M., K. R. Cave, and S. L. Franzel (1989). Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance* 15, 419–433.

Yantis, S., and J. C. Johnston (1990). On the locus of visual selection: Evidence from focused attention tasks. *Journal of Experimental Psychology: Human Perception and Performance* 16, 135–149.

Chapter 3

Dissociable Systems for Visual Recognition: A Cognitive Neuropsychology Approach

Martha J. Farah

Parsimony is a guiding principle in cognitive science as in other sciences. Of course, it is not an infallible principle; nature is sometimes more complex than we expect it to be. In the study of visual pattern recognition, most people's initial assumption is that we have a single, general purpose system for recognizing all the different types of stimuli in our visual world. After all, there is no obvious reason why a system that can recognize a face should not also be able to recognize, say, an armchair or an airplane, and a single, general purpose system has the advantage of parsimony. However, in this case the assumption of parsimony appears to be wrong. As the studies reviewed in this chapter will demonstrate, the recognition of faces and common objects appear to be functions of distinct subsystems with separate neural substrates and different ways of representing shape.

The idea that face recognition is special is not new. In addition to the neuropsychological support I discuss in this chapter, evidence from normal subjects suggests that face recognition is different from other types of object recognition. For example, infants are born with a preference for gazing at faces rather than at other objects. At just thirty minutes of age, they will track a moving face farther than other moving patterns of comparable contrast, complexity, and so on (see Morton and Johnson 1991, for a review of this and other studies of infant face perception). The "face inversion effect" to be discussed in more detail later, provides another indication that face recognition is special. Whereas most objects are only a bit harder to recognize upside down than right side up, inversion makes faces dramatically harder for normal adult subjects to recognize. (See Valentine 1988, for a review of this research.)

The writing of this chapter was supported by ONR grant N00014-93-I0621, NIMH grant R01 MH48274, NINDS grant R01 NS34030, Alzheimer's Association/Hearst Corporation Research Grant PRG-93-153, a University of Pennsylvania Research Foundation grant and an NSF STC grant to the Institute for Research in Cognitive Science at the University of Pennsylvania. I would like to thank Jim Tanaka for his collaboration in developing the concept of holistic representation for faces, and my other co-authors on the papers cited herein for vital input into this research.

These findings from normal subjects indicate two clear differences between face recognition and the recognition of other objects: face recognition has earlier developmental precursors and is more orientation sensitive than other types of object recognition.

These differences need not, however, imply that different systems are involved. How might one individuate different systems? In this chapter I will use three commonsensical criteria. To be considered different, two systems must: (1) be functionally independent, such that either can operate without the other; (2) be physically distinct (which is not necessarily redundant with the first criterion—two programs running on the same computer can be functionally independent); and (3) process information in different ways, so that it is not merely a physical duplicate of another. By these criteria, the foregoing data on face tracking and inversion effect do not tell us whether face and object recognition are accomplished by different systems. Faces could be the first type of shape a general purpose system represents. Similarly, faces might require a special orientation-sensitive type of shape representation derived within a physically unitary and functionally indivisible system.

In this chapter, the hypothesis that face recognition depends upon a specialized system will be tested with data from both brain-damaged and normal subjects. Although disease and injury do not normally confine their damage to functionally defined subdivisions of the brain, occasionally a person does sustain a brain injury with relatively selective effects on one cognitive system. Such individuals have much to teach us about the functional architecture of the mind. They have been called "experiments of nature" because their brain damage can be viewed as an (unfortunate) experimental manipulation that eliminates one component of the cognitive architecture and allows us to observe the results. Normal subjects figure in this research in two ways: as control subjects who provide baseline performance data against which to measure patients' impairments, and as subjects of interest in their own right. In the latter case, some hypotheses that arise from the context of neuropsychological research can be more conveniently tested with normal subjects.

3.1 The Visual Agnosias: Impairments of Visual Recognition

The most relevant neuropsychological impairment for present purposes is *visual agnosia*. The term *agnosia* refers to an impairment of object recognition that is not attributable to a loss of general intellectual ability or to an impairment in such elementary visual perceptual processes as brightness, acuity, depth, and color (see Farah 1990, for a detailed overview). Thus by definition, agnosics retain full knowledge of the nonvisual aspects objects and can recognize them by touch, hearing characteristic sounds, or listening to a verbal definitions. They can also perceive at least some of visual properties.

In *associative agnosia*, perception can be remarkably well preserved, to the extent that the person may be able to draw a good copy of a drawing or object he or she cannot recognize. Indeed, the term was coined in the nineteenth century because it seemed that perception was intact in these cases and that the problem must therefore lie in associating perception with knowledge of the objects. Our understanding of vision has now progressed to the point where we can identify different levels of visual representation—from those early and intermediate representations that make explicit the edges and surfaces of an image to higher-level representations that make explicit the more stable shape properties of the distal object (see Chapter 4). Associative visual agnosia is currently viewed by most neuropsychologists as an impairment at the highest levels of visual representation, rather than as an inability to associate normal visual representations with other types of knowledge. According to this view, the ability of associative agnosics to draw the object results from their use of lower-level visual representations, whereas recognition requires higher-level representations. The observation that associative visual agnosics tend to make visual errors—that is, that they mistake objects for visually similar objects—is also at least suggestive of an impairment in visual perception (see Farah 1990).

Associative visual agnosia does not always seem to affect the recognition of all types of stimuli equally. The selectivity observed in some cases of agnosia lends support to the hypothesis that there are specialized systems for recognizing particular types of stimuli. The best-known example of this is *prosopagnosia*, the inability to recognize faces after brain damage.

Prosopagnosics cannot recognize familiar people by their faces alone and must rely on other cues for recognition, such as a person's voice or distinctive clothing or hairstyles. The disorder can be so severe that the patient will not even recognized close friends and family members. One prosopagnosic described sitting in his club and wondering why another member was staring at him so intently. When he asked a steward to investigate, he learned that he had been looking at himself in a mirror (Pallis 1955)! Although many prosopagnosics also experience some degree of difficulty recognizing objects other than faces, in other cases the deficit appears to be strikingly selective for faces (e.g., DeRenzi 1986).

3.2 Prosopagnosia: Damage to a Specialized Recognition System?

The most straightforward interpretation of prosopagnosia is that the highest levels of visual representation are subdivided into specialized systems, and prosopagnosics have lost the specific system that is necessary for recognizing faces but not essential—or at least less necessary—for recognizing other types of objects. However, it is possible that faces and other types of objects are recognized using a single, general purpose recognition system but that faces are simply the most difficult type of object to recognize. Prosopagnosia could then be explained as a mild form of agnosia in which the impairment is detectable only on the most taxing form of recognition task. This account has the appeal of parsimony in that it requires only a single type of visual recognition system. Perhaps for this reason, it has gained considerable popularity (see e.g., Damasio, Damasio, and Van Hoesen 1982).

To determine whether prosopagnosia is truly selective for faces, and hence whether the human brain has specialized mechanisms for recognizing faces, we must therefore assess the prosopagnosic performance on faces and nonface objects *relative* to the difficulty of these stimuli. One technical difficulty encountered here is that normal subjects invariably perform both face and nonface recognition tasks nearly perfectly. The resultant ceiling effect thus masks any differences in difficulty that might exist between tasks and makes it pointless to test normal subjects in the kinds of recognition tasks traditionally administered to prosopagnosic patients.

With this problem in mind, Karen Klein, Karen Levinson, and I looked for a visual recognition task that would allow us to manipulate task difficulty for normal subjects, with the goal of setting normal performance at a moderate level (Farah, Klein, and Levinson, in press). The performance of a prosopagnosic subject on face and nonface stimuli could then be assessed relative to normal performance on the same tasks, answering the question of whether the subject was disproportionately impaired at face recognition.

Our subject, LH, was a forty-year-old man who has been prosopagnosic since an automobile accident in college. He is profoundly prosopagnosic, unable to recognize reliably his wife, children, or even himself in a group photograph. Yet he is highly intelligent, has no difficulty recognizing printed words, and only minimal difficulty recognizing objects. Although he has a degree of impairment with recognizing objects in drawings, this appears less severe than his impairment with faces.

We employed a recognition memory paradigm in which subjects first studied a set of photographs of faces and nonface objects, then performed an old/new judgment on a larger set of photographs, half of which were old. In a first experiment, we compared the recognition of faces to the recognition of a variety of nonface objects, which were paired with very similar foils, (as shown in Figure 3.1). In this experiment, we succeeded in equating the difficulty levels of the two sets of stimuli for a set of normal undergraduate subjects at approximately 85 percent correct. LH was given additional study time with the stimuli to ensure that he would

Figure 3.1
Examples of faces and objects used in recognition memory study with normal subjects and a prosopagnosic subject. The top item in each triple was studied, and the bottom two items were test items.

(d) (e) (f)

Figure 3.1 (cont.)

perform above a chance level. LH showed a significantly larger performance disparity for the two stimulus sets than the normal subjects, achieving only 62 percent correct for faces and 92 percent correct for objects.

In a second experiment, we attempted to test a particular version of the hypothesis that face recognition is just harder than object recognition, a view that has recently been promoted by Damasio and his colleagues (1982). Accordingly to this account, it is the fact that faces are highly similar exemplars all belonging to the same category (namely *face*) that makes them particularly taxing. We tested this hypothesis by comparing recognition of exemplars of the category face with an equivalent number of highly similar exemplars from a single nonface category, namely eyeglass frames. Examples of the stimuli are shown in Figure 3.2.

The faces and eyeglass frames were divided evenly into sets of old items, which appeared in both the study and the test phases of the experiment, and sets of new items, which appeared only in the test phase. Similar-looking eyeglass frames were separated into old and new sets to make the task more challenging; for example, there were both old and new horn-rims, and old and new aviator-style frames. As before, LH was disproportionately impaired at face recognition relative to nonface recognition compared to normal subjects. In this experiment, normal subjects found face recognition considerably easier than eyeglass frame recognition. Normal undergraduates achieved, on average, 87 percent correct responses on faces and 67 percent correct on eyeglass frames. A second group of normal subjects matched in age and education level with LH showed the same disparity; they achieved, on average, 85 percent correct on faces and 69 percent correct on eyeglass frames. LH showed significantly less face superiority in this task than normal subjects, achieving 64 percent correct for faces and 63 percent correct for eyeglass frames. Like the first experiment, this one also suggests that LH's impairment in face recognition cannot be attributed to a general problem with object recognition. The results also suggest that the problem does not lie with the recognition of specific exemplars from any visually homogeneous category but is specific to faces.

A final experiment was undertaken to address the specificity of LH's face-recognition impairment. In essence, the first two experiments compared LH's performance with faces and his performance with stimuli that are similar to faces—in their recognition difficulty and their membership in a visually homogeneous category—but are not processed by the hypothesized face-specific recognition mechanism. Stating the experimental design in this way suggests the ideal nonface comparison stimulus: upside-down faces. As mentioned earlier, inverting a face makes it much harder for normal subjects to recognize. On the basis of the face-inversion effect, it is generally assumed that if a specialized face-recognition mechanism exists,

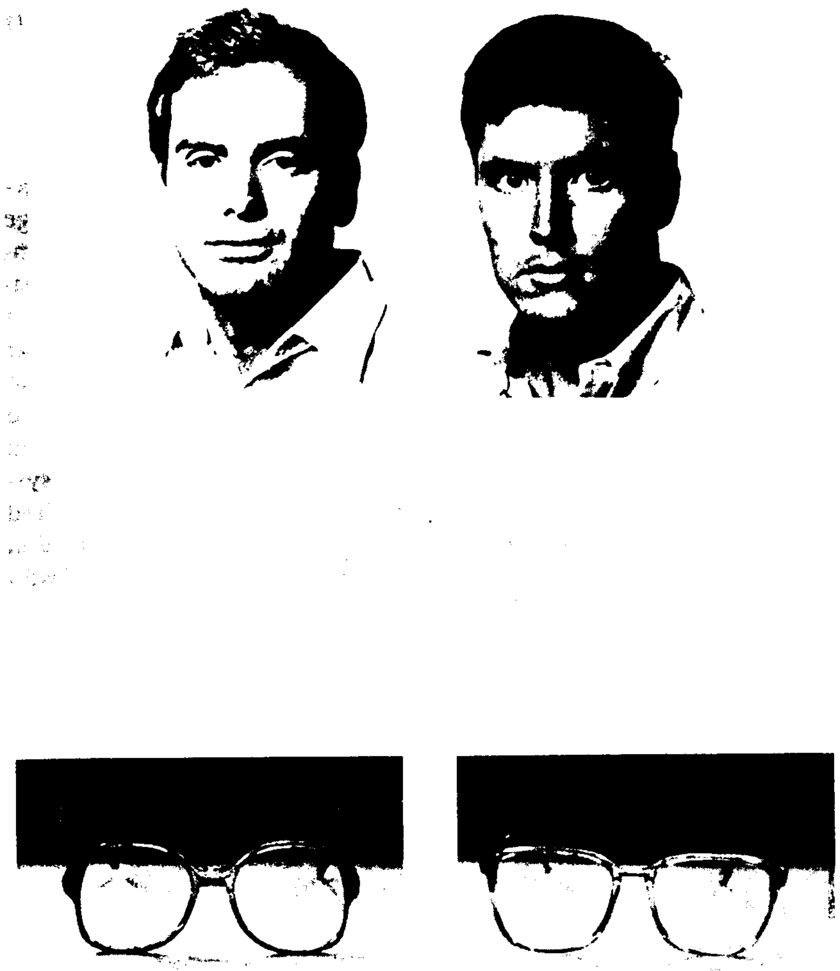

Figure 3.2
Examples of faces and eyeglass frames used in recognition memory study with normal subjects, prosopagnosic subject, and object-agnosic subject.

it is specialized for the processing of upright faces. Inverted faces, therefore, constitute ideal comparison stimuli: They are equivalent to upright faces in virtually all physical stimulus parameters, including complexity and inter-item similarity, but they do not engage (or engage to a lesser extent) the hypothesized face-specific processing mechanisms.

My colleagues and I reasoned that if LH's underlying impairment was not face-specific, he would show a normal face-inversion effect (Farah, Wilson, Drain, and Tanaka, in press). In other words, he would perform normally with upright faces relative to his performance on inverted faces. In contrast, if he had suffered damage to neural tissue implementing a specialized face-recognition system, he would show an absent or attenuated face-inversion effect. That is, he would be disproportionately impaired with upright faces relative to his performance on the comparison stimuli, inverted faces.

LH and normal subjects were tested in a sequential matching task, in which an unfamiliar face was presented, followed by a brief interstimulus interval, followed by a second face, to which the subject responded "same" or "different." The first and second faces of a trial were always in the same orientation, and upright and inverted trials were randomly intermixed. As expected, normal subjects performed better with the upright than with the inverted faces, replicating the usual face inversion effect: 94 percent versus 82 percent correct, respectively.

LH's results were more surprising. He was significantly more accurate with inverted faces, achieving 58 percent correct for upright and 72 percent correct for inverted faces! This outcome was not among the alternatives we had considered. We had assumed that if he had an impaired face processor, it would simply not be used in this task and he would, therefore, show an absent or attenuated face-inversion effect. Instead, it appears, he has an impaired face-specific processor, which is engaged by the upright but not by the inverted faces, and used even though it is impaired and thus disadvantageous. This result was confirmed in additional studies, which invariably showed either statistically significant or nonsignificant trends in the same direction.

Two major conclusions follow from LH's "inverted inversion effect." First, LH's prosopagnosia results from damage to a specialized face-recognition mechanism. Inverted faces are the perfect control stimulus for equating faces and nonface objects for such factors as complexity and inter-item similarity. LH's disproportionate impairment on upright relative to inverted faces is therefore strong evidence that an impairment of face-specific processing mechanisms underlies his prosopagnosia.

A second, and unexpected, finding was that LH's specialized face-perception system was contributing to his performance, even though it was impaired and clearly maladaptive. This demonstrates the involuntary nature of the specialized face system and provides very direct neuropsychological support for Fodor's (1983) characterization of special-purpose perceptual systems ("modules") mandatorily engaged by their inputs.

The general conclusion of these three studies with LH is that prosopagnosia represents the selective loss of visual mechanisms needed for face recognition, and not needed (or less necessary) for other types of object recognition. There is, therefore, specialization within the visual recognition system in which faces are recognized differently than other objects.

3.3 Selective Impairment of New Face Learning

Prosopagnosics such as LH are equally impaired at learning new faces and recognizing previously familiar faces, as we would expect from damage to the substrates of face representation. My colleagues and I recently encountered someone with an even more selective impairment. CT is impaired at learning new faces, but his ability to recognize previously familiar faces and to learn other nonface visual objects is relatively intact. (Tippett, Miller, and Farah, in preparation). This pattern of performance is consistent with a disconnection between intact face representations and an intact medial-temporal memory system. As such, it provides additional evidence that the neural substrates of face representation are distinct from the representation of other objects, as they can be selectively disconnected from the substrates of new learning.

CT's face perception was normal on a variety of measures, including the face-inversion task used with LH. His overall level of performance was also good relative to normal subjects, and he showed a normal inversion effect. His learning of verbal material and even visual material other than faces is also normal. However, when given the face- and eyeglass-learning task, he performed about as well as LH, achieving 58 percent correct for faces and 63 percent correct for eyeglasses. Additional evidence of his inability to learn faces comes from his identification of famous faces. For people who were famous prior to his head injury, CT performed within the range of eight age-matched control subjects on a forced choice famous/not famous task; whereas for more recently famous individuals he performed at chance level. One celebrity allowed us to make an especially interesting comparison between premorbid and current face recognition. In the case of Michael Jackson, the singer's extension plastic surgery following CT's injury provides us with a "within-celebrity" comparison of face recognition. Despite the greater popularity and media exposure of Michael Jackson in recent years, CT recognized an older picture of the celebrity and failed to recognize an up-to-date photograph.

3.4 Object Agnosia with Preserved Face Recognition: Further Clues to the Functional Architecture of Visual Recognition

Some associative agnosics appear to have more difficulty with object recognition than with face recognition, presenting us with the mirror image of the prosopagnosic's impaired and spared abilities. This pattern of impairment is interesting for two reasons. First, it offers further disconfirmation of the hypothesis that prosopagnosia is just a mild disorder of a general purpose object recognition system, with faces simply being harder to recognize than other objects. If this were true, how could a person do better with faces than with other objects? Second, it distinguishes two possible relationships that might hold between the specialized face system and the nonface object system. As illustrated in Figure 3.3, the two systems could be arranged in parallel, so that a stimulus would be recognized by one or the other. Alternatively, the two systems could be arranged in series, so that all stimuli would first be processed by the one system, with faces then receiving further processing by the other system.

Given the intuition that face recognition requires processing that is somehow more elaborate or demanding than object recognition, which presumably motivated the alternative accounts described in the last section, one might expect the latter, serial arrangement to hold. According to

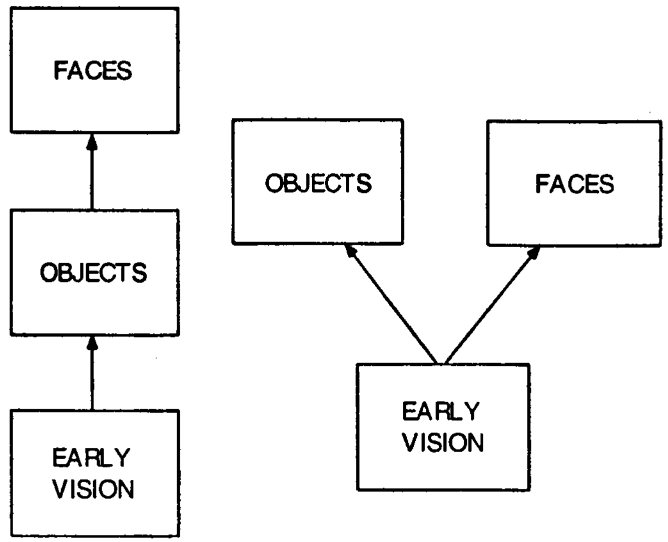

Figure 3.3
Schematic diagram of two different ways in which face recognition could be distinct from object recognition.

this view, there is a specialized face system but it is not functionally independent of the object system; it requires input from the object system and performs further processing on that input. In other words, face recognition involves normal object recognition plus some additional processing. This arrangement contrasts with the first one, according to which earlier visual processes deliver their products to two parallel, independent systems, one required to recognize faces and the other objects.

If there are indeed associative visual agnosics with relatively intact face recognition, then the systems subserving object and face recognition must be arranged in parallel. Marlene Behrmann and I recently set out to confirm experimentally the clinical observation that recognition of faces can be disproportionately spared. We used the same faces and eyeglasses experiment earlier administered to LH.

The subject in this experiment was CK, a thirty-five-year-old man who sustained a head injury in an automobile accident. An MRI showed bilateral thinning of the occipital lobes but no other focal abnormality. CK is agnosic for objects and printed words. His pattern of performance differs from normality in the direction opposite to LH's: He is 98 percent correct for faces and only 48 percent correct for eyeglasses. He shows a larger superiority of faces over eyeglasses than expected on the basis of either of the sets of normative data collected for the experiments with LH. This result, taken together with the earlier findings from LH, implies that the systems specialized for face and object recognition are functionally independent. Put more precisely, there are two systems—one more important for face recognition than for nonface object recognition and another system (or set of systems) more important for nonface object recognition than for face recognition. And they are arranged in parallel.

3.5 Functional Differences between Face and Nonface Processing

Having concluded that there are at least two specialized subsystems underlying visual recognition, let us now turn to the question of what these specialized systems might be specialized for, in terms of the kinds of visual information processing they carry out. Before addressing this question, it would be helpful to review a bit of what we know about nonface object recognition.

A recent review of published cases of associative visual agnosia suggests that face, object, and printed-word recognition are all pairwise dissociable, but that not all possible three-way combinations of impaired and spared face, object, and word recognition are possible (Farah 1991). Object recognition, in particular, was found to be impaired only if either face

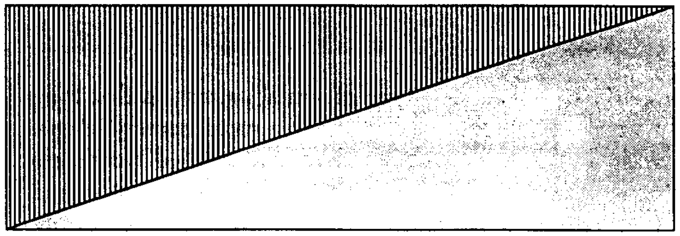

Figure 3.4
Graphic representation of the relations between two hypothesized types of pattern recognition. (See text.)

recognition or word recognition was also impaired.[1] This leads to the following hypothesis concerning the number of specialized recognition systems and their domains of applicability: There are two systems, one of which is essential for face recognition, useful for common object recognition, and not at all needed for printed word recognition, and the other of which is essential for printed-word recognition, useful for common object recognition, and not at all needed for face recognition. Figure 3.4 illustrates the inferred contributions of these two systems to face, object, and word recognition. The two hypothesized systems can account for the disproportionate face-recognition impairment of subjects like LH (damage to the first system), and the disproportionate object- (and word) recognition impairment of subjects like CK (a more severe degree of damage, to the second system). In contrast, there is no way to damage the two systems to produce an impairment in common object recognition alone, which explains the apparent absence of such cases in the literature.

This two-system interpretation of the patterns of co-occurrence of the different types of associative visual agnosia offers a clue to the nature of the nonface object recognition system. Whatever type of visual information processing this system performs, this processing is more taxed by printed words than by common objects.

1. As discussed in the original review, one case report mentioned "mild object agnosia" with no accompanying reading or face recognition difficulties in a table, but referred to the same subject as nonagnosic in the text. The information in the table represents the only violation of the pattern found by me in the literature. Rumiati, Humphreys, Riddoch and Bateman (1994) recently reported that they found another violation of this pattern. However, their subject does not appear to be a visual agnosic: His errors in naming objects are overwhelmingly semantic (e.g., "cup" for saucer) rather than visual (e.g., "record" for saucer), and he has similar problems interpreting object *names* as well as object pictures. It is not clear why such a pattern of performance would be interpreted in terms of an impairment in visual object recognition.

One salient property of words, as a type of visual pattern, is that they are composed of numerous individually recognizable parts, namely letters. In fact, research with normal subjects has shown that printed words are recognized by first recognizing their letters. For example, Johnston and McClelland (1974) found that tachistoscopic word recognition was significantly more disrupted by a mask made up of letters than by one made up of letter fragments. This finding is consistent with the idea that a necessary stage in word recognition is the explicit recognition of component letters.[2] There is also evidence that the underlying impairment in subjects who have lost the ability to recognize printed words consists of an inability to recognize multiple shapes. Such individuals typically resort to reading letter by letter, as if they can only recognize one part of the word at a time. For a fuller discussion of the role of visual perception in acquired impairments of printed-word recognition, see Farah and Wallace (1991).

Like words, most objects can also be subdivided into component parts. In fact, as reviewed by Biederman in Chapter 4, many current theories of object recognition hypothesize some form of structural description that is a representation of object shape in terms of parts, which are themselves explicitly represented as shapes in their own right. The more extensive the part decomposition, the more parts there will be in an object's representation, but the simpler those parts will be. The less the part decomposition, the fewer parts there will be in an object's representation, but the more complex those parts will be.

The conjecture being put forth here is that word recognition involves extensive part decomposition and, therefore, requires the ability to represent a large number of parts; face recognition, on the other hand, is holistic in that it involves virtually no part decomposition, and hence requires the ability to represent complex parts. Common objects are represented using a mixture of the two types of representation.

3.6 Face Recognition and Holistic Shape Representation: Empirical Tests

The patterns of co-occurrence among disorders of face, object, and word recognition suggest the existence of two complementary systems of shape representation. Consideration of the types of representations underlying word recognition led, above, to a conjecture about those underlying face recognition. Specifically, if the system that is essential for recognizing words is specialized for the representation of numerous but relatively simple parts, then the system that is essential for recognizing faces might be specialized for the representation of complex but relatively few parts. In collaboration with James Tanaka and others, I carried out several tests of this hypothesis.

In one set of studies, we reasoned as follows (Tanaka and Farah 1993): To the extent that some portion of a pattern is explicitly represented as a part for purposes of recognition, then when that portion is presented later in isolation, subjects should be able to identify it as a portion of a familiar pattern. In contrast, if a portion of a pattern does not correspond to the way the subject's visual system parses the whole pattern, then that portion presented in isolation is less likely to be recognized. Tanaka and I taught subjects to identify a set of faces, along with a set of nonface objects, and then assessed their ability to recognize both the whole patterns and their parts. Examples of study and test stimuli are shown in Figure 3.5. Relative to the recognition of houses, face recognition showed a greater disadvantage for parts relative to wholes: Subjects achieved, on average, 81 percent and 79 percent accuracy for parts of houses and whole houses, respectively, and 65 percent and 77 percent for parts of faces and whole faces, respectively. This is what we would expect if the representations underlying face recognition do not explicitly represent parts or do so to a lesser degree than nonface objects. Similar results were obtained with inverted faces and scrambled faces as the nonface comparison stimuli.

My collaborators and I recently adapted Johnston and McClelland's masking paradigm (mentioned earlier in connection with word recognition) to a new test of the hypothesis that face perception involves less part decomposition than the perception of such other stimuli as words, houses, or inverted faces (Farah, Wilson, Drain, and Tanaka, 1995). Recall that Johnston and McClelland found that word perception was more disrupted by a mask composed of letters, compared with one composed of letter fragments; they inferred that a necessary stage in word recognition is letter recognition. In our first experiment, word and face perception were assessed in a sequential same/different matching task in which the first stimulus (word or face) was presented only briefly and was followed by a mask. We used two kinds of masks. In the "part mask" condition, either letters or facial features were presented in spatial arrangements that did not make real words or faces. In the "whole mask" condition, a word or a face was used to mask the first stimulus. We predicted that if faces are perceived holistically, without explicit representations of their parts, the part masks should not be very disruptive of face perception compared to the disruption caused by a whole face mask. With word perception, on

2. The word superiority effect, by which letters embedded in words are perceived better than words presented in nonwords or alone, might appear to imply that words are perceived holistically, without decomposition into letters. However, its implications are weaker than this. It implies only that, in addition to individual letter representations, word or letter-cluster representations are also activated, and that the activation states of the latter representations influence those of the former.

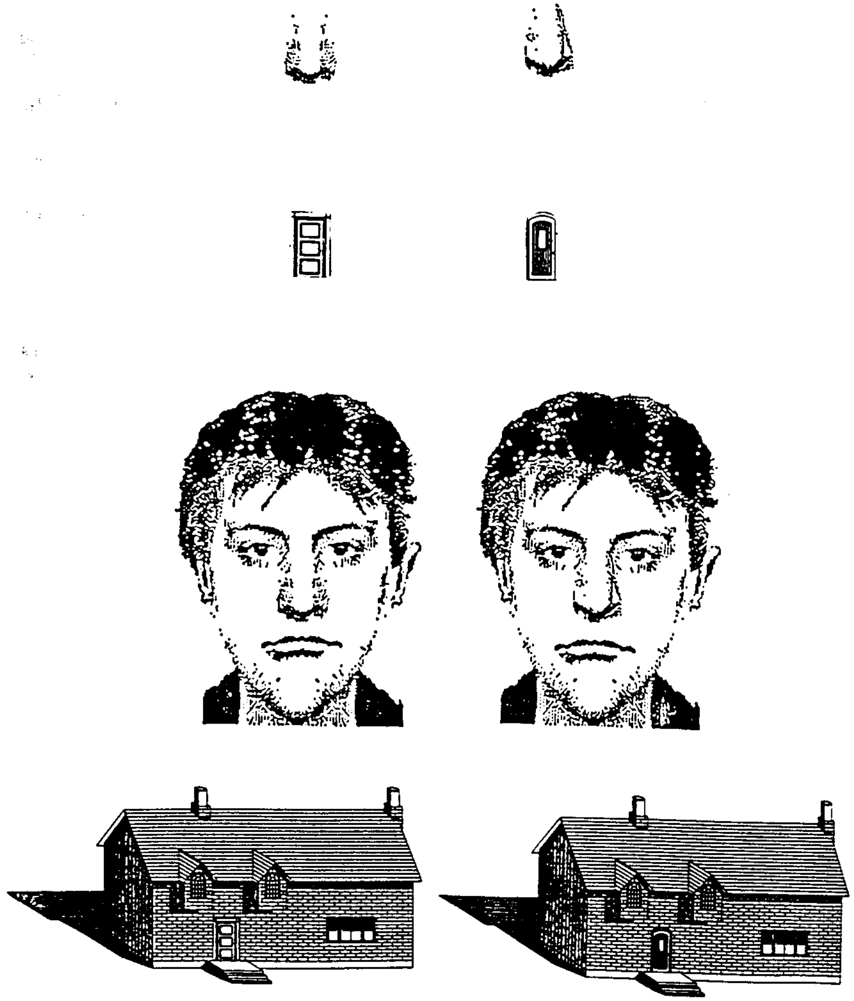

Figure 3.5
Examples of face and house stimuli used in memory study with normal subjects.

the other hand, part masks should be effective. This is what we found: Whereas subjects correctly judged 78 percent of the faces with a part mask, their accuracy dropped to 73 percent with a whole mask. In contrast, their performance for words with parts and whole masking was 78 percent and 77 percent, respectively. In subsequent experiments, we found that the difference between part and whole masks is found only for upright faces; inverted faces show no difference. We also found that the perception of houses showed an intermediate degree of sensitivity to part masks. These results accord well with the hypothesis of a parts-based system and a holistic system, used together for the recognition of objects such as houses and used separately for the perceptions of words and faces, respectively.

In the final experiment to be described, we bring the research back to prosopagnosia and the neural bases of face recognition. The neuropsychological results described earlier imply that there is some neurologically distinct subsystem that is more important for face recognition than for other kinds of object recognition. The results of the experiments described imply that normal subjects perceive faces more holistically than they perceive other kinds of objects. Taken together, these findings suggest that the face-recognition system damaged in prosopagnosia is a one of relatively holistic representation. The final experiment tests this hypothesis directly.

Tanaka, Drain, and I compared the relative advantage of whole faces over face parts for normal subjects and for the prosopagnosic LH. Our initial plan was to administer the same task Tanaka and I used with the normal subjects to LH; but despite intensive effort, LH could not learn to recognize a set of faces. We therefore switched to a short-term memory paradigm in which a face was presented for study, followed by a blank interval, followed by a second presentation of a face. The subject's task was to say whether the first and second faces were the same or different. There were two different conditions for presentation of the first face: it was either "exploded" into four separate frames containing the head, eyes, nose, and mouth (in their proper relative spatial position within each frame), or presented intact. The second face was always presented in the normal format, so that the two conditions can be called parts-to-whole and whole-to-whole. Normal subjects performed better in the whole-to-whole than in the part-to-whole condition; they averaged 93 percent and 74 percent correct answers, respectively, thus providing further evidence that their perception of a whole face is not equivalent to the perception of its parts. LH showed abnormally little difference between the two conditions, scoring 74 percent and 73 percent correct answers, respectively. This finding is consistent with the hypothesis that he can no longer benefit from seeing faces as wholes.

3.7 General Conclusions

Starting with the clinical observations of people with brain damage, hypotheses about the functional architecture of visual object recognition were formulated and tested in controlled experiments with brain-damaged subjects. Questions about the nature of shape representation within this architecture were initially addressed using normal subjects, but as soon as some preliminary answers were obtained in this way, the experimental paradigms could be adapted for use with brain-damaged subjects and the linkage between types of representation and neural systems could be tested directly. Because our interest is in the functioning of the normal system, experiments with both brain-damaged and normal subjects are relevant to testing these hypotheses. Our decision to use a given population depends on theoretical and practical considerations.

Let us summarize what we have learned from the foregoing experiments. The selective impairment of face recognition in a prosopagnosic subject, LH, suggests that we are endowed with a specialized system for recognizing faces. This system is not necessary for (or is less important for) recognizing common objects, even when such objects form a large and visually homogeneous category. Furthermore, the system is anatomically distinct, in that it can be selectively damaged by head injury. Studies of subject CT suggest that this system can also be selectively disconnected from other brain areas. A pattern of impairment opposite to LH's was observed in an agnosic subject, CK, suggesting that the face recognition system does not merely elaborate the processing of the object system, but rather processes stimuli in parallel with it, and is at least partially functionally independent of the other system. How many specialized systems are there? Patterns of co-occurrence among disorders of face, object, and printed word recognition over many cases suggest that there are two underlying systems of representation. According to this interpretation, LH has moderate damage to one system—the one essential for face recognition, used for object recognition, and not needed for word recognition. CK has severe damage to the other system—the one essential for word recognition, used for object recognition, and not necessary for face recognition.

The two systems can be distinguished by the way they represent shape. Previous research has suggested that word recognition requires the ability to represent numerous shapes and that impaired visual word recognition results from a reduction in the number of shapes that can be represented within a short time. Research with normal subjects suggests that faces are recognized as single complex wholes that are not decomposed into separately represented parts. A final study with LH showed that the quality of his face perception was not dependent on the opportunity to perceive the face as a whole, which is consistent with the idea that he has an impairment in the holistic perception of faces. Referring back to the issue raised at the outset of this chapter, we can now offer a tentative answer: Face recognition and common object recognition depend on different systems that are anatomically separate, functionally independent, and differ according to the degree of part decomposition used in representing shape.

Suggestions for Further Reading

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

Farah, M. J., and G. Ratcliff (eds.). (1994). *The neuropsychology of high-level vision: Collected tutorial essays*. Hillsdale: Erlbaum Associates.

Heilman, K. M., and E. Valenstein (eds.). (1993). *Clinical neuropsychology* (3rd ed.). New York: Oxford University Press.

Shallice, T. (1988). *From neuropsychology to mental structure*. New York: Oxford University Press.

Problems

3.1 The neuropsychological studies described in this chapter all involved single cases, rather than groups of subjects. What are the advantages and disadvantages of this practice?

3.2 Why might the visual system have evolved different systems for recognizing different types of stimuli?

References

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

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

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

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

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

Farah, M. J., and M. A. Wallace (1991). Pure alexia as a visual impairment: A reconsideration. *Cognitive Neuropsychology* 8, 313–334.

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

Johnston, J. C., and J. L. McClelland (1980). Experimental tests of a hierarchical model of word identification. *Journal of Verbal Learning and Verbal Behavior* 19, 503–524.

Morton J., and M. H. Johnson (1991). CONSPEC and CONLERN: A two-process theory of infant face recognition. *Psychological Review* 98, 164–181.

Pallis, C. A. (1955). Impaired identification of faces and places with agnosia for colors. *Journal of Neurology, Neurosurgery and Psychiatry* 18, 218–224.

Tanaka, J. W., and M. J. Farah (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology*.

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