CHAPTER THREE

Integrative agnosia

Glyn W. Humphreys
Cognitive Science Research Centre, School of Psychology, University of Birmingham, Edgbaston, Birmingham, UK

HISTORICAL BACKGROUND

Visual object recognition is undoubtedly a complex process. This is testified by the many attempts to build robust computer vision systems that are capable of recognising objects across a range of views, under different lighting conditions and in different contextual environments (see Brady, 1997; Lowe, 1987 for examples). It is also a commonplace assumption in the computer vision community that visual recognition comprises a number of distinct steps—these include: edge extraction, grouping of local image features, segmentation of objects from the background and from other objects that may be present, formation of a structural description of the object, and accessing stored structural and semantic information from the description assembled from the image (e.g. see Biederman, 1987; Marr, 1982). If the brain follows similar steps to achieve recognition, and if these steps are at least to some degree localised in different neural regions, we might expect visual recognition to break down in a variety of ways, according to the nature of the component processes involved.

In contrast to computational arguments for there being several necessary substages in object recognition, the neurological and neuropsychological literature on recognition disorders has traditionally adopted a dichotomous approach. This approach originates in the pioneering work of Lissauer (1890). Lissauer distinguished between two forms of recognition disorder or visual agnosia: apperceptive and associative. Apperceptive agnosia was diagnosed as an impairment that disrupts the formation of a normal percept for the visual stimulus, although sensation of the basic properties of the image should be spared (e.g. brightness perception). Associative agnosia was diagnosed as being an impairment of the processes involved in retrieving stored memories from objects, despite perceptual processing being intact (see Chapters 4 and 9 for further discussion). Clinically, the distinction has often relied on the presence or absence of the ability of patients to copy objects that cannot be recognised. The label "associative agnosia" is applied to patients who can copy objects, and the term "apperceptive agnosia" is given to patients who fail to copy as well as recognise. This means of classifying patients is still in vogue today (e.g. see Behrmann, Moscovitch, & Winocur, 1994).

Despite the continuing popularity of the apperceptive–associative distinction, however, case studies over the past ten years have indicated that a finer-grained analysis of patients is possible. For example, several patients have now been documented who show impaired visual access to knowledge about the associative or functional properties of objects and yet can perform difficult object decision tasks at a high level (e.g. Hillis & Caramazza, 1995; Riddoch & Humphreys, 1987a; Sheridan & Humphreys, 1993; Stewart, Parkin, & Hunkin, 1992). Such object decision tasks require discrimination between real objects and non-objects formed by interchanging the parts of real objects to create unfamiliar, though perceptually "good" stimuli. Since non-objects may not be rejected from their general perceptual attributes, good performance is contingent on access to stored visual memories for familiar objects. In these patients, then, there can be access to stored visual memories without access to semantic information that defines the "meaning" of the stimulus. Such patients appear to represent "true" associative agnosics in the sense defined by Lissauer, since any deficit occurs after perceptual access to some forms of stored knowledge have taken place. Other patients perform relatively poorly at object decision but they are nevertheless able to carry out many apparently high-level perceptual tasks without difficulty—this can include matching objects presented in different views, where invariant perceptual properties must be extracted from objects (e.g. Forde, Francis, Riddoch, Rumiati, & Humphreys, 1997; Humphreys & Rumiati, 1998; Sartori & Job, 1988). Yet other patients are impaired at unusual view matching, though they are able to perform a variety of perceptual tasks at a reasonable level (e.g. finding a figure on a complex background, counting the number of three-dimensional figures present in two-dimensional line drawings with occluding parts etc.; Warrington & James, 1988; see Chapter 4). In some further cases, perceptual judgements about even rudimentary aspects of form can be severely impaired (judged by poor copying, impaired matching of line orientations, object sizes etc.), leading to the recognition problem. However, the same patients may be able to use the same properties of form for making actions (e.g. they show an appropriately scaled grasp aperture when reaching to objects of different size, despite being poor at perceptual judgements of size; see Milner, Perrett, Johnston, Benson, Jordan et al., 1991; Milner & Goodale, 1995, for example). Thus their impairment cannot be attributed to poor sensory discrimination, but rather perhaps some form of dissociation between visual information used for recognition and perceptual judgements, and visual information used for action. These dissociations indicate that within the broad distinction between apperceptive and associative agnosia, a number of different forms of recognition disorder can be found. Several attempts have been made to capture these different disorders within multi-stage models of vision (e.g. Humphreys & Riddoch, 1987a, 1993; Humphreys et al., 1994; Warrington, 1982, 1985).

One particular disorder, due to poor perceptual integration of form information, was identified in a single case study reported by Riddoch and Humphreys (1987b). In this chapter I review the initial evidence for "integrative agnosia" along with other data collected subsequently from the same patient. I relate the disorder found in the original patient to findings with other patients in the literature, and I discuss the implications of the results for understanding visual object recognition.

DEFINING INTEGRATIVE AGNOSIA

The patient studied by Riddoch and Humphreys (1987b), H.J.A., suffered an infarct of the posterior cerebral artery. This resulted in bilateral damage to the occipito-temporal regions of the cortex, involving the lingual and fusiform gyri (see Riddoch et al., 1999, for an MRI scan). There was a superior altitudinal defect for both visual fields, but brightness detection within his lower fields was preserved. Following the lesion, H.J.A. was profoundly impaired at a variety of vision-dependent tasks: object recognition, face recognition, word recognition and reading, colour perception and finding his way around his environment. These problems were modality-specific. H.J.A.'s tactile recognition of objects was good, and his ability to name objects from definition, and to give definitions of objects from their names, was entirely normal.

Like a number of other agnosic patients documented in the literature (e.g. Goldstein & Gelb, 1918; Grossman, Galetta, & D'Esposito, 1997; Sirigu, Duhamel, & Poncet, 1991; Wapner, Judd, & Gardner, 1978), H.J.A.'s attempts to identify objects were characterised by piecemeal descriptions of the forms. For example, when presented with a paintbrush H.J.A. remarked: "it appears to be two things close together but obviously it is one thing or else you would have told me." When presented with a line drawing of a pig he described each part of the object in turn and then deduced that it was a pig from the shape of its tail: "there is a round head joining what looks like a powerful body; there are four shortish legs; it doesn't say anything to me; ah but there is a small and curly tail so I think it must be a pig." When asked to describe how he went about identifying faces, H.J.A. said: "recently I've been going on the eyebrows but they don't help very much." These errors indicate that H.J.A. had little sense of familiarity for objects he failed to identify, and that he had some difficulty in perceiving objects as perceptual wholes. Indeed, one common tendency was for him to over-segment stimuli, so that parts of the same object became parsed as separate stimuli (as with the paintbrush example). Identification, when it occurred, was typically based either on the presence of some diagnostic local feature (the tail of the pig) or on a long process of deductive reasoning.

H.J.A.'s naming errors were always visually related to target objects and never related purely in terms of their semantic association (e.g. he named a line drawing of a nose as "a soup spoon" [due to the line representing the contour of the nose having an upturn at the bottom], a violin was named as "a mechanical tool with a turning bit" [the pegs], etc.). It is unlikely that such a pattern, of "pure" visual errors, reflects a deficit after access to forms of stored knowledge has been achieved (see Plaut & Shallice, 1993, for simulations of naming errors after lesioning different levels of a model of object naming). H.J.A. was also better able to identify real objects (at around a 60% level) than photographs (around 40%) and he was worst at identifying line drawings (around 30%, depending upon the items). Adding surface detail, and 3D depth information (via stereo, with real objects), benefitted performance. When surface detail was present, H.J.A. was less inclined to segment objects into separate parts. Surface information from objects thus seems to interact with the processes involved in integrating form elements together to form coherent perceptual wholes. When objects were mis-identified, H.J.A. was never able to indicate their use, by gesturing. He was also poor at matching tests requiring access to semantic information from objects (e.g. judging whether a hammer is used with a screw or a nail, when the stimuli were presented as pictures) and at object decision tasks. The problem was not one of naming but of recognition.

Yet, despite the indications of there being a visual locus for the deficit in H.J.A. (e.g. the visual errors and the effects of surface detail on recognition), he performed well on many standardised tests of perceptual processing. For example, he was able to reproduce highly accurate drawings of objects that he failed to identify. Figure 3.1 shows H.J.A.'s copy of an etching of St Paul's cathedral in London, which is highly accurate despite the picture being complex and containing many edge segments. Nevertheless, this particular picture took six hours to reproduce! So, although H.J.A.'s copies were accurate, they took an abnormally long time to complete. Also, H.J.A. often drew lines in an unusual order when copying, instead of following the parts of a single object in a coherent way (e.g. if one line fell across and occluded two parts of an object, he would move from one part of the object to follow the line instead of first reproducing the two parts of the same object together). Due the time taken and the unusual order of pencil strokes, Riddoch and Humphreys suggested that H.J.A.'s drawings were not necessarily reflecting a normal perceptual process, but rather a process of serially following each line without necessarily organising the lines into objects. It follows that, though the end product of copying may be good, we should be cautious in accepting accurate copies as evidence of intact perception.

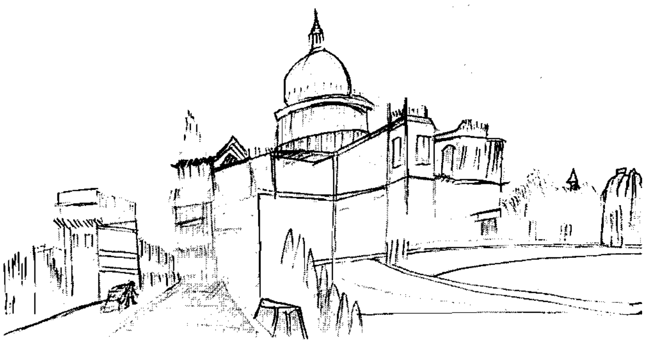

FIG. 3.1 Copy by H.J.A. of an etching of St Paul's cathedral, London, during the blitz in World War II.

FIG. 3.2 Examples of stimuli from the Efron shape test (after Efron, 1968). The task requires squares to be discriminated from rectangles, matched for brightness.

In H.J.A.'s case, though, arguments about his perceptual processing did not rest solely on copying; he was in addition good at a test of shape perception, used to diagnose shape coding problems in other patients. This test, used originally by Efron (1968), requires that the patient judge whether shapes such as those shown in Fig. 3.2 are squares or rectangles. The squares and rectangles are equated for area and brightness, so the stimuli cannot be discriminated from these properties; differences between the shapes are then varied systematically to provide a sensitive measure of shape discrimination. H.J.A. performed normally on this task (Humphreys, Riddoch et al., 1992), whereas patients with impaired encoding of basic properties of form are impaired (Benson & Greenberg, 1969; Campion, 1987; Davidoff & Warrington, 1993; Milner et al., 1991). Furthermore, H.J.A. could make orientation and size-matching judgements at a normal level (Humphreys & Riddoch, 1984). And, in visual search tasks requiring detection of a target differing in the two-dimensional orientation of some of its lines relative to distractors, H.J.A. manifested (normal) flat search functions (where reaction times [RTs] are affected only minimally by the number of distractors present; Humphreys, Riddoch et al., 1992). Flat search functions are typically interpreted as evidence of parallel processing supporting the discrimination of targets from distractors. This last test is of some interest since it provides an on-line measure of processing efficiency. At least when the differences between targets and distractors were quite salient, H.J.A. demonstrated evidence of processing form information in a spatially parallel manner.

Thus the initial tests of visual processing were consistent with there being relatively good, parallel encoding of displays, and consequently they were consistent with the inference that the recognition deficit must be post-perceptual in nature: that is, H.J.A. had associative agnosia. The only contra-indication of this was with H.J.A.'s copying, though even here the finished drawings were satisfactory. Also it must be acknowledged that copying is a fickle behaviour to measure, and it is possible that H.J.A.'s slow, line-based strategy may occur too sometimes in normal subjects. To infer that there is some underlying perceptual deficit, more probing tests of visual processing are required.

Riddoch and Humphreys (1987b) reported data from a number of such additional tests.

1. Object decision performance was contrasted with line drawings and silhouettes. Silhouettes preserve the global outline shape of objects, but lose internal line details. Normal subjects find silhouettes more difficult than line drawings, presumably because they are able to use the extra details present in line drawings to identify the objects (or discriminate the objects from the non-objects, in object decision). In contrast to this, H.J.A. tended to perform better with silhouettes than with line drawings. This suggests that the internal details in line drawings disrupted rather than enhanced H.J.A.'s perception; for instance, internal lines may serve as segmentation cues which H.J.A. is abnormally sensitive to. Such cues, when present, lead to H.J.A. parsing the shapes incorrectly. More recently, Lawson and Humphreys (1999) have found similar effects in a picture-word verification task. Normal subjects are slow to verify silhouettes relative to line drawings, especially as objects are rotated away from a prototypical orientation. H.J.A. showed no sign of this disruption with silhouettes. Note that stored information about objects should be the same whether accessed by line drawings or silhouettes. The fact that performance can be somewhat better with silhouettes indicates an effect occurring at a pre-recognition stage.
2. H.J.A. was tested using sets of overlapping figures, with performance measured relative to baselines in which the same figures were presented alongside each other. H.J.A. was slowed disproportionately with overlapping figures relative to the non-overlapping baselines.
3. Object identification was compared across a range of presentation durations, using a set of line drawings that H.J.A. was often able to identify in free vision. H.J.A.'s identification performance decreased dramatically as the exposure duration shortened. With an unlimited exposure he named around 80% of the drawings correctly, with a 500ms exposure this decreased to around 30% and with a 100ms exposure only around 15% of the objects were named.

On all three tests in which the visual properties of the displays were made more difficult to assimilate, H.J.A. performed worse than controls. These results pointed to the presence of an underlying perceptual deficit, even though H.J.A. could copy objects and make basic shape discriminations. What characterises the tasks where H.J.A. did well and those where abnormalities were detected? The tasks where H.J.A. succeeded (copying, single shape discrimination) could all be done in a serial manner (e.g. with parts of objects being encoded one at a time), and they had unlimited presentation times (so there were no costs in accuracy due to the encoding of parts being serial). The tasks where he was impaired (1–3 above) (a) constrained the opportunity for serial encoding (e.g. by reducing the exposure time), and (b) used stimuli with multiple internal segmentation cues (with line drawings and overlapping figures, containing numerous T-junctions).

To account for the pattern of performance, Riddoch and Humphreys proposed that H.J.A. could process basic, local visual elements in a relatively normal way (e.g. as indicated by search for orientation-defined target lines), in parallel across the visual field. However, the processes involved in integrating those elements into perceptual wholes, by grouping, were impaired. Due to this poor grouping of visual elements, H.J.A. was abnormally sensitive to segmentation cues, and tended to parse stimuli inappropriately into separate parts. Human recognition is limited to just one object at a time (Baylis & Driver, 1993; Duncan, 1984). Segmentation processes in vision act to deliver a parsed visual field in which separate objects can be identified in turn. Thus segmentation may be viewed as the counterpart of grouping—the tendency to group elements together into a single object description competing against processes that act to segment displays into separate objects. Impairments to grouping will consequently lead to over-segmentation of the visual array, and poor recognition. Recognition processes may then operate in a piecemeal way, and be strongly affected by display time.

ANALYSES USING VISUAL SEARCH

Supportive evidence for H.J.A. having a problem in grouping local form elements was reported by Humphreys, Riddoch et al. (1992). They used visual search tasks which required the detection of a "form-conjunction" target from amongst distractors made up of similar local elements—an example would be to search for an inverted T amongst upright T distractors, all of which contain horizontal and vertical form elements, with targets distinguished from distractors by the way in which these elements conjoin. In studies of normal observers, Duncan and Humphreys (1989) and Humphreys, Quinlan, and Riddoch (1989) had demonstrated that search for form conjunctions was strongly affected by grouping relations between the stimuli. Search is relatively efficient (showing only weak effects of the numbers of distractors present) when distractors are homogeneous; search is inefficient and linearly related to the numbers of distractors present when distractors are heterogeneous (e.g. Ts rotated 90° left and right as well as upright). Heterogeneous distractors tend not to group with one another, and any grouping that does operate will be as strong between the distractors and the target as it is between distractors. Hence targets and distractors will not be segmented easily on the basis of parallel grouping operations. Search may then depend on serial selection of one stimulus at a time, leading to linear search functions. In contrast, homogeneous distractors (being identical) will tend to group together and be segmented from the target, making search more efficient. Indeed, RTs can be particularly fast to "target absent" trials with homogeneous displays, due to subjects responding to a homogeneous group of distractors.

H.J.A. was given similar tasks and manifested an unusual pattern of performance. He was as good as the control subjects on the tasks that controls find difficult—search for a target amongst heterogeneous distractors. He was poor at the normally easy task of searching for a target amongst homogeneous distractors. His RTs were affected by the number of distractors present, he made numerous errors, and there was no evidence of "fast absent" responses to homogeneous displays of distractors.

The fact that H.J.A. was no worse than the controls in the difficult search task indicates that he has no problems in serial search across visual displays. The selective deficit with homogeneous displays, however, is consistent with H.J.A. having impairments to a process used by control subjects to make search of these displays efficient—parallel grouping between the elements. Humphreys, Riddoch et al. (1992) replicated these findings with abstract forms as well as with letter stimuli, demonstrating that the effects were not confined to the use of letter-like forms (Ts at various orientations). One other point to note is that, when conducting serial search with heterogeneous displays, H.J.A. made few errors—his error rate was raised only with homogeneous displays. This suggests that H.J.A. attempted to process homogeneous displays in parallel (since with serial search his error rate would be low), but he was simply poor at doing this. It appeared that he could not prevent his visual system from attempting to group elements even when he would have benefitted by treating each element individually.

These results, showing a selective deficit in search with homogeneous (groupable) displays, have been simulated by Humphreys, Freeman, and Müller (1992). They used the SERR model of visual search, in which visual stimuli are selected by activating stored "templates" for targets and distractors used in search tasks (e.g. there might be templates for Ts at various orientations, to simulate the above studies). Elements in the visual field group together by virtue of their having identical local feature combinations, and different items compete with each other if they fall at the same location (e.g. two Ts at different locations would group and support one another, but a T would compete with an inverted

T to be represented at a given location). Templates are activated according to which items are represented most strongly in the visual field, and the strength of activation varies as a function of grouping. Normally, homogeneous items group and activate their template efficiently, enabling search to be efficient. Heterogeneous items compete, making search protracted. If an activated template belongs to a distractor rather than a target, then linked items are rejected and search continues until the target is detected. This leads to serial search functions being generated, matching the standard pattern of search found with normal subjects (Humphreys & Müller, 1993). Humphreys et al. (1992) modified the model by adding noise to the activation functions. This resulted in incorrect features in local elements sometimes being activated transiently. Once this occurred, however, grouping between homogeneous items could be disrupted and the model behaved as if heterogeneous items were present. Search with homogeneous items became slow and error-prone. Interestingly, there was relatively little effect with heterogeneous distractors since items tended to compete rather than group in any case. The results mimic the data from H.J.A. In the model, disruption to a specific process, distractor grouping, selectively affects search with homogeneous displays; this provides an existence proof that a similar impairment could underlie H.J.A.'s impairment.

ENCODING WHOLES AND PARTS

Much but not all of the data reported in Humphreys, Riddoch et al. (1992) and Riddoch and Humphreys (1987b) highlighted a problem for H.J.A. in grouping local features to form visual "gestalts". However an exception to this was his tendency to perform better with silhouettes than with line drawings (Riddoch & Humphreys, 1987b; see also Lawson & Humphreys, 1999). Earlier we attributed this to a tendency to over-segment visual objects into their parts, when internal line cues for segmentation were present (with line drawings but not with silhouettes). The result also suggests, though, that H.J.A. is sensitive to some more wholistic information in vision; for example, his object decisions with silhouettes were above chance, consistent with his using overall shape outline on at least some occasions. How might such wholistic representations be formed?

Humphreys, Riddoch and Quinlan (1985) investigated this issue using compound letters (i.e. large "global" letters formed from smaller "local" letters; see Navon (1977) for a first example). The task was to discriminate, on different blocks of trials, whether the global or the local stimuli were Ss or Hs. When responses were made to the global forms, the local forms could be consistent or inconsistent with the response (e.g. both Ss or 1S and 1H), or they could be neutral (e.g. an O). The same manipulation occurred when responses were made to the local forms (when the identity of the global forms could be consistent, inconsistent or neutral). Under appropriate conditions (e.g. when there is some uncertainly concerning the location of the target and the local elements are sufficiently dense), normal subjects respond faster to the global than to the local forms and the identity of the global forms affects local responses (e.g. there is interference when their identities are inconsistent) (see Navon, 1977). This provides one example of when the global "forest" seems to be identified before the local "trees".

Under equivalent conditions H.J.A., like control subjects, responded more quickly to the global than to the local letters (for H.J.A. RTs were as much as 300ms faster than global letters). Thus he was able to discriminate the global letters as perceptual wholes. In fact, H.J.A.'s responses to global letters were relatively normal and it was his responses to local letters that were slowed (when the letters appeared in the context of the global shape). However, unlike controls, H.J.A. showed no indication of any global interference on local responses. It might be argued that the lack of global interference arose precisely because H.J.A.'s local responses were slow; for example, any initial activation of a response by the global letter may have decayed by the time the local letter was identified. However, data with normal subjects show that interference effects occur across quite wide variations in local and global response times (Lamb & Robertson, 1988, 1989), making this account unlikely. Further, subsequent to this initial study Lamb, Robertson and Knight (1990) found somewhat similar results in a group study of patients with lesions to the superior temporal gyrus (STG) (of either hemisphere), though they showed a much smaller overall RT advantage for global letters (i.e. there was less disruption in responding to local forms; patients with right hemisphere lesions in fact showed a local advantage). This unusual pattern, of a global advantage without interference, suggests instead that the global and more local aspects of the forms may be processed independently, with interference arising only when the two forms of information are integrated perceptually. In H.J.A., and perhaps also patients with lesions of the STG, global aspects of shape can be derived, but these representations are not embellished efficiently with more detailed local form information. This may facilitate selective attention to the local and global aspects of form, minimising interference effects. In H.J.A.'s case the disruption in deriving local form information may mean that local and more global representations are never fully integrated, contributing to his recognition deficit. The STG communicates with the inferior occipito-temporal region via area MT (Kaas, 1988), and so STG lesions may disconnect this region from inferior temporal regions concerned with perceptual integration (which are damaged bilaterally in H.J.A.'s case).

What form might H.J.A.'s global representations take? At least two possibilities suggest themselves. One is that his global representations are based on low spatial frequency components in displays. The other is that they are based on coding the positions (but not the identities) of the local elements. If global forms are derived from position-based coding, then at least one form of grouping would seem to operate—grouping by proximity—even if other forms of grouping are impaired (e.g. grouping by similar identities). Whatever the case, information about the local identities of parts will not be specified (as they might be if grouping by similarity took place). To derive sufficient local information for object identification to operate, H.J.A. may then attempt to process parts serially, leading to piecemeal naming responses.

Other evidence supporting the proposal that H.J.A. can encode global shape information, but that this information is impoverished relative to the information derived by normal subjects, was reported by Boucart and Humphreys (1992). They had H.J.A. make perceptual matches to fragmented line drawings. The fragments in the line drawings could be aligned and collinear or they could be misaligned so that they were no longer collinear, but the overall shape had the same low spatial frequency components as before. Normal subjects are advantaged when they match forms with collinear segments. H.J.A. showed no evidence for this. He was sensitive to the orientation of the global form, however; he could better discriminate items whose global orientation differed than items with the same global orientation. Again it appears that there was impaired grouping by collinearity to support the global information that could be derived, either from low spatial frequency components or from a position-base analysis.

AGNOSIA AND SIMULTANAGNOSIA

The findings from H.J.A. indicate a deficit in integrating local elements into articulated representations of perceptual wholes, with perception breaking down into a parts-based analysis of objects on many occasions. One might ask, what is the relation between such a disorder and the syndrome of simultanagnosia, in which patients seem again to have limited ability to process visual information in parallel but this co-occurs with a relatively good ability to identify single objects (e.g. Balint, 1909; Coslett & Saffran, 1991; Humphreys & Price, 1994; Kinsbourne & Warrington, 1962)?

In behavioural terms, H.J.A. manifests few signs of simultanagnosia. He is able to report on the presence of several objects simultaneously (e.g. when asked to decide how many objects are present his RTs are relatively unaffected for up to four objects; Humphreys, 1998; Humphreys et al., 1985). He negotiates his environment successfully, picking up objects correctly and avoiding collisions (Humphreys & Riddoch, 1987b). In this respect, H.J.A. behaves quite differently from patients with simultanagnosia following bilateral lesions of the parietal lobes (Balint, 1909).

Humphreys (1998), as others before him, proposed that dorsal and ventral areas of the brain perform separate computational functions in vision (e.g. see also Milner & Goodale, 1995). According to Humphreys, ventral regions deal with the analysis of parts within objects; dorsal with the representation of at least a limited number of separate objects. It is this representation of a limited number of objects, within the dorsal visual stream, that provides us with some awareness of the spatial structure of the visual environment. Dorsal visual structures remain intact in H.J.A., and presumably enable him to move successfully in the environment (even if he does not recognise the objects present!). Within his ventral system, however, there is a limitation in the parallel grouping of visual forms, impairing object recognition. The opposite pattern of impairment may be found after bilateral lesions of dorsal visual areas. In this case, patients show poor awareness of the spatial structure of their environment. However, in tasks requiring the recognition of single objects such patients can show good performance and even evidence of processing visual parts in parallel (e.g. word identification can be unaffected by the number of letters present; Humphreys, 1998).

Farah (1990) distinguished between two forms of simultanagnosia, according to whether patients had lesions affecting ventral or dorsal visual areas. According to Farah, "ventral" simultanagnosia may be due to a limited visual short-term memory and "dorsal" to impaired disengagement of attention from objects. H.J.A., though, showed few deficits in visual short-term memory, at least as assessed in enumeration tasks. Whilst agreeing with the distinction between different functional deficits after ventral and dorsal lesions, it remains my contention that they are better characterised in terms of impairments in the construction of different forms of spatial representation (parts within objects and separate objects). Patients classed as ventral simultanagnosics following unilateral left ventral lesions (e.g. Kinsbourne & Warrington, 1962) may simply have a reduced version of the deficit suffered by H.J.A. after bilateral lesions. A unilateral left deficit may particularly impair the parallel grouping of parts within objects that are represented within the left hemisphere (words).

LONG-TERM VISUAL MEMORY

Although H.J.A. was severely impaired in visually recognising objects, he performed remarkably well at tasks designed to tap aspects of his visual memory for form. For example, his drawings from memory were as accurate as those produced by control subjects, and he produced detailed descriptions of objects from memory, including information about their visual properties (size, shape etc.)(Riddoch & Humphreys, 1987b). The only clear deficit on initial testing of long-term visual memory was with colour knowledge, which was quite poor. Such results suggested that, at least as far as form information is concerned, visual perceptual processes necessary to object recognition can be separated from long-term visual memory and the imagery processes that support performance in drawing and long-term recall (see Chapter 5 for further discussion of this argument). In this respect, colour knowledge may be somewhat different, and rely on the re-activation of perceptual representations of colour.

Subsequent studies, however, have indicated that H.J.A.'s long-term visual knowledge is not perfect, when probed using tasks similar to those where he shows a deficit in perception. Young, Humphreys, Riddoch, Hellalwell, and de Haan (1994) investigated H.J.A.'s long-term visual knowledge of faces. He showed good recall of individual features of faces but poor memory for more "configural" properties, which may require features to be integrated. This subtle deficit is consistent with perception and long-term memory recall tapping at least some common processes.

AGNOSIA 16 YEARS ON

H.J.A. suffered his brain lesions in 1981. In 1997, Riddoch et al. (1999) retested his performance on many of the original tasks used to diagnose his agnosia. H.J.A.'s ability to identify real objects showed some improvement. However, there was little change in his identification of line drawings, and he continued to be impaired on tests stressing the integration of form information, such as identifying overlapping figures and line drawings compared with silhouettes. This suggests that the basic underlying visual impairment had remained relatively constant, though he had become better able to use other forms of stimulation (e.g. stereo and texture cues, in real objects). A more pronounced improvement with real objects than with line drawings has been noted before in follow-ups of agnosic patients (e.g. Wilson & Davidoff, 1993). At a more detailed level, subsequent tests revealed that H.J.A. was still selectively impaired with overlapping figures, and he continued to perform relatively better with silhouettes than line drawings. The basic symptoms of integrative agnosia remained.

Interestingly, H.J.A.'s ability to recall the visual properties of objects did show some deterioration. For example, his line drawings of objects were less easy for control subjects to identify and his definitions specified fewer visual features. This was not due to some overall drop in performance, however; in fact his definitions contained more verbal detail than previously. These results suggest that long-term visual memory interacts with visual perception, at least in the sense that visual memories deteriorate unless updated by intact perceptual descriptions. Over the longer term, an integrative perceptual problem can also contribute to the loss of long-term visual knowledge about objects.

RELATIONS TO OTHER PATIENTS

As I noted above, the types of visual identification error made by H.J.A. resemble those described in several other case reports in the literature, in which patients attempt to name objects via the serial identification of their parts. The patient reported by Butter and Trobe (1994), when asked to describe how many objects were present when given single line drawings, even stated that there were several objects present, identifying parts as separate items. The patient reported by Shelton, Bowers, Duara and Heilman (1994) showed poor copying of objects when the parts had to be spatially related. Such responses are consistent with some form of problem in perceptual integration. Butter and Trobe also assessed their patient with overlapping figures and silhouettes. Similarly to Riddoch and Humphreys (1987b), there was a marked impairment with overlapping figures and relatively better performance with silhouettes than line drawings. Thus this pattern of performance is not unique to H.J.A.

DeRenzi and Lucchelli (1993) reported data on a patient who, like H.J.A., demonstrated relatively good shape perception on the Efron task (Fig. 3.2), along with poor performance on overlapping figures tests and on several tests requiring access to stored memory from line drawings (e.g. object decision). DeRenzi and Lucchelli's patient also found it very difficult to discriminate realistic from impossible figures, created by making local parts structurally inconsistent with one another. Such a task requires that the parts be integrated together. On similar tasks, H.J.A. too performed poorly. DeRenzi and Lucchelli's patient, however, was also impaired in recalling the perceptual details of objects from long-term memory. For example, she was poor at drawing objects from memory (though copying was relatively good) and she was often unable to describe the perceptual difference between two objects, when given their names. In such a case, the perceptual impairment seems to co-occur with a disorder of stored visual knowledge. Now, whilst it can be argued that, when probed, H.J.A. too had a deficit in recalling the kinds of visual attributes he had difficulty in perceiving (e.g. facial configurations; Young et al., 1994), the severity of any memorial deficit was less pronounced than his perceptual impairment. It seems likely that patients can have associated lesions, which generate substantial problems in long-term visual memory, in addition to any perceptual deficit (cf. DeRenzi & Lucchelli, 1993). A similar argument can be applied to the patient described by Grailet, Seron, Bruyer, Coyette, and Frederix (1990). Like H.J.A., the copies produced by this patient indicated some problems in parsing visual stimuli. For instance, surface reflectance properties were reproduced as if they were parts of objects. Drawing from memory, though, was impaired (this patient tended to reproduce general associative knowledge in his drawings; for example he introduced a container into the body of an animal when asked to draw a camel), and he was deficient in naming to visual definitions. The patient was in addition impaired in tactile object recognition. Grailet et al. proposed that their case had a central (cross-modal) problem in integrating parts into wholes, and suggested that H.J.A.'s deficit was at an earlier stage of binding visual features together. However the presence of a memory deficit in this patient could contribute to the cross-modal nature of his problem.

Riddoch and Humphreys (1987b) argued that, for H.J.A., there was relatively good encoding of basic properties of shape, along with poor integration of parts to wholes. They suggested that the processes of grouping parts into wholes is also necessary for accurate figure–ground coding to occur. For example, with overlapping figures the ability to link parts to one object may help in segmenting it from the background. The parts that enter into this integration process include correctly computed local contours. Subsequent to this, Davidoff and Warrington (1993) have posited that deficits in shape coding and deficits in figure–ground formation can doubly dissociate. They report a patient who was able to describe simple overlapping geometric forms, who could judge the number of three-dimensional line drawings present in a display, and who could make perceptual judgements about whether contours were aligned or not. However, the patient was poor at shape discrimination tasks, such as the Efron shape test. They argued that figure–ground coding was relatively intact in this patient, though shape coding was impaired. However, it is possible that figure–ground processes still rely on outputs from shape coding mechanisms, and for this pattern of deficit to occur. This would be the case if intact figure–ground coding processes can recover from poor shape input, at least when figure–ground is not taxed or measured under real-time conditions. Interestingly, in some of the tests used by Davidoff and Warrington to demonstrate intact figure–ground coding H.J.A. too performs at a high level. For example, he can discriminate letter fragments shown against a background of visual noise (the shape detection test from the VOSP battery; Warrington & James, 1991), and he can count the number of three-dimensional line drawings present (again using stimuli from VOSP). Deficits in grouping and segmentation are nevertheless apparent when time-based measures are used or time restrictions imposed.

CONCLUSIONS

H.J.A.'s case demonstrates that a form of visual agnosia can exist even though a patient shows relatively good basic coding of shape and even though stored knowledge of objects is largely preserved. The deficit appears to affect a stage of visual processing intermediate between basic shape coding and visual access to memory representations, concerned with parallel perceptual grouping and the integration of perceptual parts into wholes. It is revealed most strikingly under conditions that stress visual segmentation and grouping. It indicates that Lissauer's original distinction between apperceptive and associative agnosia needs to be fractionated further, to reflect the sub-processes involved at the different processing stages (see Humphreys & Riddoch, 1987a; Humphreys et al., 1994).

ACKNOWLEDGEMENTS

This work was supported by grants from the Medical Research Council (UK), the EU and the Human Science Frontier Programme.

REFERENCES

Balint, R. (1909). Seelenahmung des "Schauens": Optische ataxie, raumliche Storung der Aufmerkamsamkeit. *Monatschrift für Psychiatrie und Neurologie*, *25*, 51–81.

Baylis, G., & Driver, J. (1993). Visual attention and objects: Evidence for hierarchical coding of location. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 451–70.

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*, *20*, 1068–87.