PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

Recognizing Faces [and Discussion]

Vicki Bruce and M. E. Le Voi

Phil. Trans. R. Soc. Lond. B 1983 **302**, 423-436
doi: 10.1098/rstb.1983.0065

Phil. Trans. R. Soc. Lond. B **302**, 423–436 (1983)
Printed in Great Britain

Recognizing faces

By Vicki Bruce
Department of Psychology, University of Nottingham, Nottingham NG7 2RD, U.K.

Face recognition is an area where research has increased considerably in recent years, yet theoretical progress has been slow. Here it is argued that by considering the perception and recognition of familiar faces, as well as episodic memory for unfamiliar faces, a functional framework for face recognition can be developed. Experiments using faces, that include tasks analogous to 'visual search' and 'lexical decision' are described, and the processes in operation are compared with those occurring in word recognition. The results allow us to distinguish a number of possible subcomponents for a functional model of face recognition.

Introduction

Human faces are visual configurations that are structurally complex. Monozygotic twins apart, no two faces are identical, yet the physical differences distinguishing one face from another may be subtle and can defy verbal description. Nevertheless most people can recognize the faces of many hundreds or even thousands of individuals encountered in their everyday life or through the media. Our facility to recognize faces is almost certainly given by their multiple social significance. Faces serve to identify an individual and to portray or betray his or her emotional attitude, and expressive movements may serve to disambiguate information during comprehension. Despite the essentially social importance of faces, the processes involved in encoding, storing and recognizing faces should be explicable within a cognitive framework.

Research into face recognition has increased dramatically during the last decade (the most comprehensive review is given by Davies *et al.* 1981). Some stimulus for this activity came from a general increase in interest in non-verbal memory after demonstrations of the apparent superiority of pictorial over verbal recognition memory (Shepard 1967). Early picture recognition memory studies were criticized (see, for example, Goldstein & Chance 1974) because the pictures used were highly heterogeneous, and other researchers therefore chose to explore recognition memory for homogeneous materials. It was generally (though not universally (see Ellis 1975)) observed that recognition memory for pictures of human faces was superior to that demonstrated for other homogeneous sets of items, such as canine faces, buildings (Scapinello & Yarmey 1970), inkblot pictures or snowflakes (Goldstein & Chance 1971).

The highly accurate performance shown in laboratory tests of facial memory seemed to contrast sharply with the apparent fallibility of the eyewitness to a crime, as evidenced by public concern over cases of mistaken identity (see, for example, Devlin 1976). The paper presented by Elizabeth Loftus at this meeting has illustrated the activities of cognitive psychologists in the area of eyewitness testimony. This applied problem provided a second impetus for laboratory research into the processes of facial recognition, as well as into ways of making reconstructions of faces by witnesses more accurate (Davies 1981).

In general, therefore, research into face recognition expanded because of subjects' accuracy (in the laboratory) or inaccuracy (in the field) of selecting from a set of distractors that face or those faces that had previously been briefly viewed. The research activity generated has arguably been unsystematic (Goldstein & Chance 1981) and has certainly largely been conducted without any guiding theoretical framework (Ellis 1975, 1981; Bruce 1979). Nor has much theory emerged from this research, with some recent exceptions (e.g. Hay & Young 1982).

Theoretical progress in face recognition might be achieved if we adopt a systematic functional approach to facial memory. We need to differentiate different kinds of coding that may be applied to faces and that may mediate their subsequent recognition. In the area of verbal memory it is commonplace to describe the system in terms of different functional subcomponents, whether these are conceived of as separate 'stores' (see, for example, Baddeley and Hitch & Halliday, this symposium) or different kinds of processing (see, for example, Craik, this symposium). Indeed there have been some attempts to apply theoretical ideas from the verbal memory literature to face memory. Thus the 'levels of processing' framework (Craik & Lockhart 1972) was applied to memory for faces by Bower & Karlin (1974) and stimulated some interest (Patterson & Baddeley 1977; Winograd 1978; Parkin & Hayward 1983). However, such ideas have come from the literature on verbal episodic memory. Arguably, to understand how we remember faces in laboratory 'episodes' we need also to understand how they are perceived and recognized in more everyday activities. Thus perception and memory for faces should be considered together (and indeed the major thrust of the original levels of processing framework was that memory should be seen as a by-product of perceptual processing). A functional approach to face recognition requires that we consider the everyday process of recognizing or searching for the faces of friends and acquaintances, and to see how such processes may be applied in rather more artificial tasks of remembering faces in the laboratory. In this paper some episodic memory experiments will be described, but the emphasis will be on studies more analogous to those found in the literature on 'semantic memory' and 'word recognition'.

Here it will be shown that we can usefully distinguish several different types of code which may be involved in storing, recognizing and identifying faces. The description of these different coding processes has been possible through experiments in which decision latency as well as accuracy has been recorded, and in which the processing of familiar as well as unfamiliar faces has been studied. This paper will describe this experimental work and end by making comparisons between the coding processes that appear to be implicated in facial identification and memory, and those proposed elsewhere in the literature on visual and verbal memory.

Episodic recognition memory for faces: pictorial, structural and associative coding

In many studies of facial recognition memory, the subject's task is to recognize those photographs that were earlier studied from a larger set containing distractor (unstudied) items. The problem with presenting identical photographs at study and test is that facial memory is confounded with pictorial memory or 'stimulus recognition' (Bruce 1982; Hay & Young 1982). The subjects may be able to perform the task on the basis of recognizing superficial characteristics of a particular portrait rather than recognizing the *face* depicted. Outside the laboratory, our ability to recognize the faces of friends in novel poses must depend on some stored encoded representations of the faces, perhaps capturing information that remains invariant under natural transformations (cf. Gibson 1966), or to which 'rules' about how faces

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

are distorted by natural transformation can be applied. Such a representation, which captures the essential structural aspects of a face that allow it to be distinguished from other faces, will here be termed a 'structural description' or a 'structural code' for a face. To study the formation and use of structural codes for faces, we must change the pictures of the faces between study and test.

Some experiments suggested that subjects were unaffected by moderate changes in pose between presentation and test (Patterson & Baddeley 1977; Davies *et al.* 1978), suggesting that the pictorial memory component in such face-recognition tasks is minimal. However, more recent experiments (see, for example, Bruce 1982) have failed to confirm this. Bruce (1982, expt 1) showed that faces that were changed either in expression (e.g. smiling to unsmiling) or in pose (full face to three-quarters), or both, between study and test, were less accurately and more slowly recognized than those that were identical study and test, and the decline in performance was greater if both expression and pose were changed than if only one change was made.

This suggests that 'pictorial' memory plays a part in recognizing photographs of faces, but leaves open the question of whether or not we should regard some 'pictorial code' as a separate subcomponent in face recognition. A possible explanation of such results is that subjects' recognition performance is based on some visual code of the initially viewed faces, which contains both 'pictorial' and 'structural' detail, and that later recognition is dependent on the degree of overlap between the pictorial and structural features of the test face and those contained in the stored trace. However, the results of a second experiment suggest that we may be able to think of 'pictorial' and 'structural' codes for faces as distinct.

In this study (Bruce 1982, expt 2) recognition memory for familiar faces, as well as unfamiliar faces, was tested, and the pictures of the faces were either identical or changed in both pose and expression at test. The familiar faces were those of the academic staff in my department, so that it was possible to obtain systematically varied pictures of each individual, which is difficult to do when the faces are those of celebrities. The subjects were student and other members of the same department.

Subjects were shown a set of 24 faces (half of them familiar) for 8 s each and later tested on a set of 48 faces (again half of them familiar). Twenty-four of these were those studied, with half tested in the same view and half tested in a different view, in which head angle and expression were changed. Subjects were warned before studying the faces, and again before the test, that the pictures of the faces might be changed. Both recognition accuracy and latency were recorded.

Table 1 gives the results of this experiment. First it can be seen that while recognition accuracy for the familiar faces does not appear to suffer if these are transformed at test, there is a statistically significant effect of changing the faces on recognition latency. The fact that accuracy is not affected is probably due to a ceiling effect. The important point is that there is a demonstrable effect on recognition latency, suggesting that some aspects of the original pictures shown must have been encoded. This suggests that we can distinguish a 'pictorial' from a 'structural' code, since something corresponding to the latter must already be stored in the case of familiar faces. It is the structural codes for familiar faces, whatever their nature, that allow the faces of the academic staff to be recognized when originally studied. We may also note that table 1 confirms, with different subjects, the effect of transforming unfamiliar faces that was demonstrated earlier.

The second main observation on these results is that recognition memory for the familiar faces is much more accurate and rapid than for the unfamiliars (though it is important to note that there is no significant interaction between familiarity and picture change in the analysis of the latency data). A simple dual coding theory (Paivio 1971) could account for the superior recognition memory for familiar faces. Both a visual and a verbal representation (i.e. a name) are available for a familiar but not for an unfamiliar face. Thus recognition memory for the familiar faces can be mediated by 'dual' rather than 'single' codes.

TABLE 1. MEAN PERCENTAGE HITS AND DECISION LATENCIES TO PREVIOUSLY PRESENTED FAMILIAR AND UNFAMILIAR FACES (BRUCE 1982)

	type of change...	none	angle and expression
familiar faces	correct (%)	95.8	94.5
	latency/ms	987	1066
unfamiliar faces	correct (%)	88.8	54.8
	latency/ms	1325	1557

The false positive rates in this experiment were 7.7% to familiar and 6.3% to unfamiliar faces.

However, as hinted in the introduction, it may be more profitable to think in terms of multiple rather than dual coding systems. As we have already seen, it may be possible to consider more than one visual code on which recognition memory can be based. Likewise the notion that familiar faces are coded 'verbally' may be an oversimplification.

Rather than thinking of familiar faces simply in terms of their ease of naming, we can think of them as being represented in semantic memory. One need not have a *name* for a familiar face to have available a set of specific semantic associations (e.g. 'She's the lady who keeps the sweet shop that stays open late, who always has a half in the local at about 10 p.m.'). These specific semantic associations, which may include, or allow access to, a name, and which are independent of the particular view of the face seen, can provide information that can mediate episodic recognition decisions. However, unfamiliar faces may also arouse associations which may be rather general (male/female; happy/sad) or more specific ('He looks just like my Uncle Henry.'). At least some of these associations may be dependent on a particular view of a face. The advantage that familiar faces enjoy in recognition memory may be given by their property of always producing specific semantic associations, which may include a name, in addition to any general, or affective, associations made to the pictures shown. The latency data gathered in this experiment can be fairly well accounted for by assuming that the more codes that can be accessed at test, the faster will be the decision. (See Bruce (1982) for the details of this argument.)

SEARCHING FOR FAMILIAR FACES

In the previous section I distinguished between different types of visual code and different kinds of semantic or other associative code that might be involved in recognizing faces. In this section I move away from episodic recognition studies to consider how different coding systems may be involved in tasks where subjects are asked to decide whether or not each face that they encounter is one of a small set of familiar 'targets', a task rather more like the everyday processing of faces.

Searching for a familiar person in a crowd or series of faces involves testing each inspected

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

face for those visual characteristics that are remembered about the target person. Empirical research with alphanumeric displays (see, for example, Neisser 1963), as well as common sense, would dictate that we could find a target more easily if distractor faces were visually dissimilar on a number of features than if they were highly similar. It should be easier to find your grandfather at a disco than in an old-time dance hall. A question of rather more theoretical interest was the extent to which 'semantic' information might also be important during search. In the literature on search for alphanumeric characters and words it has become clear that the semantic relationship between target and distractor items can influence search time even though such effects appear paradoxical (for example on an assumption of search to a logically sufficient depth within a Pandemonium-type framework (see Henderson 1978)).

TABLE 2. AVERAGE TIMES (MILLISECONDS) TO REJECT EACH OF EIGHT TYPES OF DISTRACTOR FACE (BRUCE 1979)

(These latencies include a constant time taken by the slide projector to advance from one slide to the next.)

visually similar				*visually dissimilar*			
familiar politician	familiar actor	unfamiliar politician	unfamiliar actor	familiar politician	familiar actor	unfamiliar politician	unfamiliar actor
2221	2165	2160	2159	2049	1973	1966	2008

In the context of a series of experiments on visual search for faces (Bruce 1979) I set out to see whether the semantic relationship between distractor and target faces had any effect on the time taken to reject distractors in a serial classification task. Subjects were asked to respond positively if each of a series of faces was that of Alec Douglas-Home, Edward Heath, Harold Wilson or James Callaghan (at the time when the experiments were conducted, these were the four most recent Prime Ministers of the U.K.). The pictures shown of the target people were varied, so that the task could not be performed at the level of picture recognition. The distractor faces, which appeared in a ratio of 2:1 (distractors: targets) were varied on three bipolar dimensions. Distractors were either visually similar to one or more members of the target set or dissimilar (as rated independently by another group of subjects). They were either other politicians or actors (a convenient label for a general 'show-business' category of actors, singers, comedians, etc), and they were either familiar or unfamiliar (again as rated by an independent group of subjects). Examples of each of the eight distinct groups of distractors were intermingled with target items in the test series and the decision latency to each face recorded.

The average rejection latencies to the different kinds of distractor face are shown in table 2. Here we see that visually similar faces took longer to reject than dissimilar ones, and that familiar politicians took longer to reject than familiar actors. The observation that there is no difference in the time to reject *unfamiliar* politicians and actors suggests that the 'semantic' effect on the familiar faces is indeed due to their semantic category, rather than due perhaps to something about the photographic style of politicians' pictures compared with actors'.

Therefore this experiment has indicated that for faces, as with words, both visual and semantic similarity of distractor to target items slows their rejection. Perhaps the most interesting feature of the data obtained in this experiment is that the effects of visual and semantic similarity are independent. The semantic effect on the dissimilar faces is as great as that on the similar ones. This observation severely constrains the kind of model that can be constructed to account for these latency data, and this will be discussed in more detail below.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

At the time when these data were published there were no equivalent experiments with words in which visual and semantic similarity had been varied independently. Therefore it was not possible to compare explicitly the processing model proposed for faces with the data obtained in analogous experiments with words. However, in a further study (Bruce 1981) a similar task to that performed with faces was conducted with subjects responding positively if any one of four animal names (varied in type face and case) appeared in a series of other words and non-words. The distractors were varied in the visual similarity that they bore to the targets (in terms of word length and spelling), and in terms of their category. Distractors were either other animal names (cf. politicians' faces), unrelated words (cf. actors' faces) or pronounceable non-words (cf. unfamiliar faces). Again rejection latency was the dependent variable.

Table 3. Average times (milliseconds) to reject each type of distractor item when searching for target words (Bruce 1981)

	visual similarity			
	high	intermediate	low	mean
animal words	631	564	568	588
unrelated words	587	539	538	555
non-words	603	540	531	558
mean	607	548	546	

The results are shown in table 3. Here we see independent effects of visual and semantic similarity on the time taken to reject distractors. Visually similar items took longer to reject than items of intermediate or low visual similarity, and animal words took longer to reject than unrelated words or non-words. This pattern of results is so similar to that obtained for faces in the earlier experiment that it strongly suggests that similar processes are at work. Such processes can be described by a model in which the visual analysis of each item proceeds in parallel with its semantic analysis until some criterion for rejection, or acceptance of targets, is reached. (For an expansion of this argument see Bruce (1979).) This is not to suggest that semantic analysis does not require visual analysis. We must be careful to distinguish between the kind of task-directed visual feature analysis required of subjects in a visual search task, and any other kind of 'visual analysis' that may underlie the access of semantic representations in long-term memory. Thus for words we may suppose that while an active test of the features of the target items is being conducted, the encountered items are none the less analysed by some 'automatic' reading processes (direct visual access via orthographic features or phonological coding via GPC rules, or both). If the lexical units of the target set are 'primed' for acceptance, semantic effects can be explained by excitation spreading to the lexical units of categorically or otherwise related items by some associative priming mechanism. The presentation of any semantically related distractor may therefore result in 'positive' (i.e. target) evidence being sent to some hypothetical decision-making component, which will then require more 'negative' (disconfirming) evidence from the active visual feature checking process to compensate, and eventually to reject the distractor correctly.

A similar process can be suggested for faces. While an active check of the features of the target face is proceeding, faces are none the less analysed more directly for meaning. We may suppose that direct access to semantic codes in long-term memory is mediated by structural coding mechanisms. An interesting, though at the moment speculative, possibility, is that structural

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

codes for faces contact 'recognition units' for faces analogous to 'lexical units' for words (cf. Hay & Young (1982); discussed at more length below). Thus when the face of Harold Wilson is presented this contacts some recognition unit for Wilson, which will fire when sufficient evidence is obtained via the structural coding processes. We might therefore suggest that semantic effects in visual search for faces are obtained by priming these recognition units. In this context we may note that in both the faces and words experiment there was no difference in the time taken to reject familiar items unrelated to the targets (actors or unrelated words) compared with unfamiliar items (unfamiliar faces or non-words). Because only items related to the targets would be primed and hence send positive evidence to the decision process, this is not puzzling. Whether or not an item was represented in long-term memory would make no difference unless it happened to be an associate of the targets.

Contextual effects in face recognition: a case for face recognition units?

In the previous section there was some suggestion that associative priming effects could account for semantic effects in search and classification tasks involving both faces and words. Associative priming leading to altered levels of activation in recognition units has often been suggested as the mechanism whereby an appropriate context facilitates word recognition, reading time or lexical decision in a variety of tasks (see, for example, Morton 1969; Meyer & Schvaneveldt 1976). Indeed such contextual facilitation effects are a ubiquitous feature of word recognition studies. At least some of these effects now seem to be attributable to conscious, post-lexical interpretative processes, rather than to 'lexical' priming (see, for example Forster 1981). The evidence for automatic priming is still strong in tasks where the contexts for the 'primed' words are lexical, i.e. other words, as opposed to sentences. Passive priming effects can be said to be in operation when facilitation spreads to related, though not necessarily predictable, items (e.g. DOCTOR to OPERATION) as well as to predictable associates (e.g. DOCTOR to NURSE). I was interested to see whether similar effects occurred in the processing of faces in a task analogous to the lexical decision task.

In the lexical decision task the subject is asked to decide whether a letter string is or is not a word that is in their vocabulary. A comparable task for faces would be one in which subjects must decide whether or not each face of a presented series is familiar to them (i.e. whether or not it is in their 'vocabulary' of acquaintances). A task in which subjects had to discriminate faces from non-faces (e.g. other patterns or jumbled faces) would not be logically equivalent to the lexical decision task since jumbled faces could never be endowed with the same type of meaning as real faces. In the standard lexical decision task, non-words are often pronounceable, and some of them may even have obscure or obsolete meaning, unknown to the subjects or experimenter. Thus the lexical decision task is really one in which subjects say whether or not they 'know' a letter string.

The question of interest is whether subjects will be faster to decide that a face is familar if it is accompanied by a face related to it compared with a condition where a face is presented with an unrelated familiar face. In unpublished exploratory experiments, a student (Sue Heppell) and myself showed that subjects were faster to decide that a pair of faces were both familiar if both were members of the academic staff in the subjects' department, or both were television celebrities, than if the two faces were drawn from different categories (one academic

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

staff member, one t.v. celebrity). In a further experiment similar effects were observed if the faces were presented sequentially. Subjects were on average faster to respond positively to a staff member preceded by another staff member than if preceded by a t.v. celebrity, and so on. The pictures used in this study were carefully selected so that the academic staff's faces and celebrities' faces were of similar quality, contrast and range of expression, making it unlikely that these effects were due to similarities or differences in the pictorial quality of pairs of faces. However, it seemed worthwhile to explore such effects further in an experiment where all the familiar faces were those of celebrities, and in a design which allowed an analysis by items (faces) as well as by subjects. This experiment will be described here in more detail.

Table 4. Average times (milliseconds) to decide that celebrities' faces are familiar

preceding slide	related	unrelated
slide group 1	687	803
slide group 2	774	911
mean	731	857

Eighteen subjects were tested and each saw a series of 60 faces of which half were familiar and half unfamiliar (the assignment of faces to these categories was done subjectively, but subjects' responses in the experiment suggested that this had been successful: the 'error' rate was reasonably low, at about 8%). Embedded in this series of 60 faces there were 10 'critical' familiar faces, the latencies to which were of interest. Five of these were preceded by a 'related' face (e.g. Ernie Wise preceded by Eric Morecambe – a British comedy duo; Richard Baker preceded by Kenneth Kendall – both British t.v. newsreaders; Jon Pertwee preceded by Tom Baker – both actors who have played 'Dr Who' in a British t.v. series). The remaining five faces were preceded by unrelated faces (e.g. Ernie Wise preceded by a t.v. newsreader). Those faces that were preceded by related faces for half the subjects were preceded by unrelated faces for the rest, and vice versa. Each subject group saw only two or three 'predictable' pairs (such as Morecambe followed by Wise) in the entire series of 60 faces, reducing the likelihood of conscious anticipatory strategies.

The mean decision latencies to the critical faces in this experiment are given in table 4. Here we see that, on average, decisions to faces preceded by a related face are faster than to those preceded by an unrelated face. On the subjects analysis, this effect is qualified by an interaction with subject group. The semantic effect was very large for the first group of subjects (212 ms) and very small (22 ms, not statistically significant) for the second group. This is probably because the faces assigned to the 'related' and 'unrelated' conditions for the two subject groups were not of equal familiarity. If those faces in the related condition for the first group of subjects were more easily, and hence quickly, responded to than those in the unrelated condition, then when the assignment of faces to conditions was reversed for the second group of subjects this additional factor would work against any semantic facilitation effect. This interpretation of the interaction is supported by the statistical analysis performed with faces, rather than subjects, as the random factor. As the means in table 4 indicate, the average times to respond to the first group of slides (which were in the 'related' condition for the first group of subjects) are shorter than those to the second group of slides (which were in the 'related' condition for the second group of subjects). The fact that this second analysis still reveals a significant main effect of semantic relatedness suggests that the effect extends to the 'unpredictable' related pairs of faces (e.g. the newsreaders) as well as to the 'predictable' related pairs (e.g. Morecambe and Wise). A *post hoc* analysis comparing the effect of semantic relatedness for the total of five predictable pairs with that for the five unpredictable pairs included in the experiment again shows no interaction. The facilitation effect on the second faces in unpredictable pairs, of 118 ms, is almost as large as that on the predictable pairs, 135 ms.

This experiment therefore suggests that semantic facilitation effects occur in a task using faces that is analogous to the lexical decision task. The fact that semantic facilitation is observed from one face to another that is categorically related, but not necessarily a predictable associate (from one newsreader to another in the last experiment, and from one academic staff member to another in earlier exploratory work), suggests that this facilitation is of the 'automatic' rather than the 'conscious' variety. Further evidence on the nature of these facilitation effects could be gained by checking that no inhibition was present when, say, a newsreader's face was preceded by Eric Morecambe's face, compared with when it was preceded by a familiar face to which no obvious associate could be predicted. The design of the above experiment did not allow this comparison to be made. Experiments in which the faces themselves are presented very briefly, or disguised or otherwise degraded, and in which context is also manipulated, should help to answer the question of whether facilitation effects with faces should be attributed to priming of thresholds in 'face recognition units', or whether they should be attributed to other semantic memory processes. What can be said for now is that effects that have been attributed to the priming of lexial units for words are also observed for faces, thus strengthening the case for the inclusion of face recognition units as a distinguishable subcomponent in face recognition.

Face recognition compared with other recognition processes

The general theoretical framework that emerges from the above discussion is as follows. When a face is perceived, visual coding processes furnish a structural code, which, if the face is familiar, can be matched with a stored code that allows access to specific semantic information about the person, perhaps including a name. This process may be thought analogous to the access of semantic information by the activation of lexical units in word recognition (see, for example, Morton 1969). We might therefore suppose that face recognition has occurred when some recognition unit that captures the structural code for a face is sufficiently activated. The evidence presented in the preceding two sections supports this analogy with word recognition because semantic effects in search tasks, and in tasks requiring familiarity judgements, give similar effects for faces and for words, though further research is needed to pinpoint the locus of these facilitation effects in facial recognition. The hypothetical processes that allow us to identify an individual by his or her face thus involve stages of perceptual classification (via structural coding processes) and semantic classification (via the semantic memory system).

In addition to these processes, which may yield specific semantic information about the identity of a familiar face, both unfamiliar and familiar faces may be encoded pictorially (if photographs are presented) and associatively. The episodic memory experiments described in this paper suggest that pictorial detail may be preserved even for familiar faces. In addition, associative coding processes that allow us to describe a picture of a face, or a 'live' face, as 'happy', 'pretty' or 'sinister' are an essential aspect of facial processing because they allow us to identify from an individual's face information that may be important in social interaction. In this respect the processing of words and of faces are quite dissimilar, because the surface characteristics of a particular word do not convey meaning in the way that the surface characteristics of a face do. On the whole, words are arbitrarily formed with respect to their meaning. Faces are similarly arbitrarily related to their specific semantic reference (the faces of politicians may on the whole look rather older than those of rock stars, but such correlations are likely to be loose, and rather rare). However, the surface characteristics of a face – whether it is wrinkled or smooth, grey-haired or dark-haired – may be correlated with general associations and attributions that can be made, such as 'old' or 'young'. The fact that pictorial detail of faces is preserved in memory may be a consequence of the attention that we must pay to the surface detail of faces in everyday interactions.

The distinctions drawn between pictorial, structural, semantic and other associative coding processes may allow us to understand why some faces are easier to remember than others, and why faces as a class are easier to remember than other homogeneous materials. Faces will be well remembered if many codes are formed and if these codes are themselves 'distinctive' (used here in a neutral sense). Thus episodic recognition memory for familiar faces is better than that for unfamiliars because structural and semantic codes are activated for the former but not the latter. Faces rated as very attractive or unattractive are remembered better than faces given intermediate ratings (Ellis 1975). This may be because the pictorial codes for the faces given extreme ratings are more distinctive (some measure of inter-item physical similarity here would be of interest), or because more specific associations can be made to atypical faces than to typical ones, or both (cf. Goldstein & Chance 1981). Faces to which personality attributions have been made are better remembered than those to which simple physical feature judgements have been made (Bower & Karlin 1974; Patterson & Baddeley 1977), presumably because the former judgements involve more associative coding and inference by the subjects. As Winograd (1978) has shown, however, asking subjects to find the most distinctive physical feature of a face removes the advantage of personality attributions over physical feature judgements, perhaps by enhancing pictorial and structural coding processes in the latter physical feature condition. Faces as a class are better recognized than other materials such as pictures of buildings or canine faces because our social interactions with faces endow them with associative meaning, which is unlikely to be shared by many other materials.

This general framework for face recognition is entirely consistent with a functional model recently proposed by Hay & Young (1982). Their model is shown in figure 1. Though their terminology is rather different from my own, the model is essentially similar. In their model, representational processes (cf. structural coding) allow access of person information (cf. semantic codes) via the activation of face recognition units. In addition to these processes, however, other 'visual processes' can lead for example to the analysis of expression (cf. other associative coding). The formation of pictorial codes, in my scheme, could be seen as a by-product of these other 'visual processes' in theirs, when photographs of faces are used as materials. Hay & Young developed this framework to explain failures of the face recognition system, both in everyday life and in clinical syndromes. The experimental evidence presented here provides additional support for their model.

This kind of model for face recognition is also compatible with ideas recently put forward in the general area of picture and object recognition (Seymour 1973, 1979; Warren & Morton 1982). Warren & Morton (1982) and Seymour (1973, 1979) have proposed that the

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

recognition of pictures involves the activation of recognition units termed 'pictorial access-exit nodes' or 'pictogens', which allow access to a common semantic system (also accessed by words via logogens). Warren & Morton further suggest that although two different pictures of an object (say two clowns) may activate the same pictogen and hence the same semantic information, some literal representation of the particular picture presented may also be maintained. This literal representation would be analogous to a 'pictorial' code for a face.

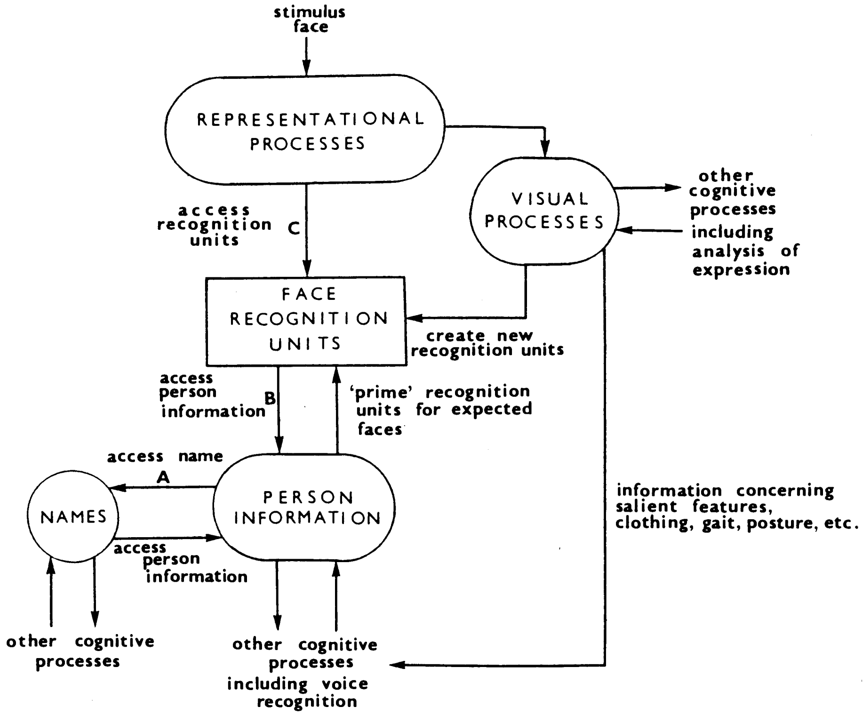

FIGURE 1. Hay & Young's (1982) model of the functional components involved in face recognition (reproduced with permission).

Thus the different functional subcomponents proposed for face recognition both here and by Hay & Young fit well with current ideas about picture recognition in general. Picture recognition, including face recognition, can be described by models similar to those proposed by some authors in the area of verbal recognition. The differences between episodic memory for pictures and for words (see, for example, Deffenbacher *et al.* 1981) may be explained by the differential availability and usefulness of different sorts of code. Words can easily be encoded phonologically; pictures must be named first. The surface structure of a word is probably highly subject to interference and is not on the whole related to its meaning. For pictures the surface structure may provide a distinctive code for later recognition, and is essential to its meaning. Just as pictures in general may be remembered rather differently from words, so we should acknowledge that some aspects of face processing and memory may be 'unique'. The type of structural coding system needed to distinguish one face from another may rarely be developed for any other materials, and the way in which we analyse faces for expression may likewise involve processing of a face-specific kind.

In conclusion, I hope to have demonstrated that by considering the perception and memory for faces together, and by conducting experiments with familiar as well as unfamiliar faces, it is possible to construct a functional theoretical framework for face recognition. The ideas I have presented are very similar to those of Hay & Young (1982), and together may provide a framework that can both guide and stimulate further research.

I thank Alan Baddeley and Pamela Briggs for their comments on an earlier draft of this paper.

References

Bower, G. H. & Karlin, M. B. 1974 Depth of processing pictures of faces and recognition memory. *J. exp. Psychol.* **103**, 751–757.

Bruce, V. 1979 Searching for politicians: an information-processing approach to face recognition. *Q. Jl exp. Psychol.* **31**, 373–395.

Bruce, V. 1981 Visual and semantic effects in a serial word classification task. *Curr. psychol. Res.* **1**, 153–162.

Bruce, V. 1982 Changing faces: visual and non-visual coding processes in face recognition. *Br. J. Psychol.* **73**, 105–116.

Craik, F. I. M. & Lockhart, R. S. 1972 Levels of processing: a framework for memory research. *J. verb. Learn. verb. Behav.* **11**, 671–684.

Davies, G. M. 1981 Face recall systems. In *Perceiving and remembering faces* (ed. G. Davies, H. Ellis & J. Shepherd), pp. 227–250. London: Academic Press.

Davies, G. M., Ellis, H. D. & Shepherd, J. W. 1978 Face recognition as a function of mode of representation. *J. appl. Psychol.* **63**, 180–187.

Davies, G., Ellis, H. & Shepherd, J. 1981 *Perceiving and remembering faces.* London: Academic Press.

Deffenbacher, K. A., Carr, T. H. & Leu, J. R. 1981 Memory for words, pictures and faces: retroactive interference, forgetting and reminiscence. *J. exp. Psychol.: hum. Learn. Memory* **7**, 299–305.

Devlin, Lord Patrick 1976 *Report to the Secretary of State for the Home Department of the Departmental Committee on Evidence of Identification in Criminal Cases.* London: H.M.S.O.

Ellis, H. D. 1975 Recognising faces. *Br. J. Psychol.* **66**, 409–426.

Ellis, H. 1981 Theoretical aspects of face recognition. In *Perceiving and remembering faces* (ed. G. Davies, H. Ellis & J. Shepherd), pp. 171–197. London: Academic Press.

Forster, K. 1981 Priming and the effects of sentence and lexical contexts on naming time: evidence for autonomous lexical processing. *Q. Jl exp. Psychol.* A **33**, 465–495.

Gibson, J. J. 1966 *The senses considered as perceptual systems.* Boston: Houghton-Mifflin.

Goldstein, A. G. & Chance, J. E. 1971 Visual recognition memory for complex configurations. *Percep. Psychophy.* **9**, 237–241.

Goldstein, A. G. & Chance, J. E. 1974 Some factors in picture recognition memory. *J. Gen. Psychol.* **90**, 69–85.

Goldstein, A. G. & Chance, J. E. 1981 Laboratory studies of face recognition. In *Perceiving and remembering faces* (ed. G. Davies, H. Ellis & J. Shepherd), pp. 81–104. London: Academic Press.

Hay, D. C. & Young, A. W. 1982 The human face. In *Normality and pathology in cognitive function* (ed. A. W. Ellis), pp. 173–202. London: Academic Press.

Henderson, L. 1978 Pandemonium and visual search. *Perception* **7**, 97–104.

Meyer, D. E. & Schvaneveldt, R. W. 1976 Meaning, memory structure and mental processes. *Science, Wash.* **192**, 27–33.

Morton, J. 1969 The interaction of information in word recognition. *Psychol. Rev.* **76**, 165–178.

Neisser, U. 1963 Decision time without reaction time: Experiments in visual scanning. *Am. J. Psychol.* **76**, 376–385.

Paivio, A. 1971 *Imagery and verbal processes.* New York: Holt, Rinehart & Winston.

Parkin, A. J. & Hayward, C. 1983 The influence of trait and physical feature based orienting strategies on aspects of facial memory. *Br. J. Psychol.* **74**, 71–82.

Patterson, K. E. & Baddeley, A. D. 1977 When face recognition fails. *J. exp. Psycol.: hum. Learn. Memory* **3**, 406–417.

Scapinello, K. F. & Yarmey, A. D. 1970 The role of familiarity and organisation in immediate and delayed recognition of pictorial stimuli. *Psychonom. Sci.* **21**, 329–330.

Seymour, P. H. K. 1973 A model for reading, naming and comparison. *Br. J. Psychol.* **64**, 35–49.

Seymour, P. H. K. 1979 *Human visual cognition.* London: Collier-Macmillan.

Shepard, R. N. 1967 Recognition memory for words, sentences and pictures. *J. verb. Learn. verb. Behav.* **6**, 156–163.

Warren, C. & Morton, J. 1982 The effects of priming on picture recognition. *Br. J. Psychol.* **73**, 117–129.

Winograd, E. 1978 Encoding operations which facilitate memory for faces across the life span. In *Practical aspects of memory* (ed. M. M. Gruneberg, P. E. Morris & R. N. Sykes), pp. 255–262. London: Academic Press.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

Discussion

M. E. Le Voi (*Human Cognition Research Laboratory, The Open University, Milton Keynes, U.K.*). Dr Bruce has made a strong case for an analogy between face and word perception. One problem that arose in the area of verbal experiments in the early 1970s was the problem of generalization from the sample of words used in an experiment, to the whole population of words: the 'language as fixed effect' fallacy. Dr Bruce's last experiment showed a significant effect due to two different randomly sampled lists of faces. In my experience, this often indicates considerable variance in the stimuli, which immediately produces problems of generalizing to the population of faces. Does Dr Bruce think this constitutes a serious problem?

As a corollary, most if not all face experiments have used male faces as stimuli. It is reasonable to expect that the results for female faces would inevitably be the same?

Vicki Bruce. There are a number of problems when doing these kinds of experiment with faces. The population of faces that one can assume are highly familiar to all one's subjects is fairly small: individual subjects have different television viewing habits and political interests. Over and above this, photographs differ in quality and in the extent to which they are considered 'good likenesses', and a particular picture of a celebrity may fail to be recognized by some subjects while being quickly recognized by others. In addition I am sure there are familiarity or 'frequency' effects that affect recognition latency for familiar faces. In the unpublished experiments I mention in this paper, for example, undergraduate subjects responded more quickly to academic staff faces than to television celebrities; and some celebrities, such as Morecambe or Wise, are always more quickly responded to than others. It therefore doesn't surprise me when I find a lot of variability within any sample of faces (it surprises me more if I don't!).

As far as the statistical issues are concerned, for at least some of my experiments one could argue that, within the U.K., I have exhausted the population of faces rather than randomly sampled it. How many famous politicians are there who also look like Harold Wilson, for example? Again, when working with predictably associated faces, such as Morecambe and Wise, there are not too many pairs one can use. This makes it difficult, if not impossible, to design good experiments, but may also make it less important to demonstrate that the effects are significant for faces as well as for subjects. (Nevertheless the effects I have described here for the search and 'lexical decision' type of experiment are significant on an items analysis.)

My feeling is that the only satisfactory way in this research to show that effects generalize to other faces would be to repeat experiments of a similar kind. It would not have been possible to repeat the experiment on searching for politicians with a new set of distractors in 1976 (when the experiment was conducted) because, as I have argued, some of the distractor types were depleted in that experiment. But it might be possible to repeat it now, when different political faces are prominent, or to repeat the experiment with a different target set (though it is difficult to think of a good one). With the 'lexical decision' type of experiment I think it is important that in three different studies I found similar main effects, even though in each study there were also interactions due to the variability in the materials.

As for the use of male faces, it is true that I have used male faces throughout my research (though other workers have used female faces in episodic memory tasks). Unfortunately there are simply not enough female academic faces in our department, or female politicians in the

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

U.K., to use in the kinds of experiment I have been doing. I should be very surprised if the effects did not hold for female faces, and it certainly should be possible to demonstrate context effects for female faces, for example. However, I should not rule out slight differences in the kinds of 'structural' and 'affective/associative' codes formed for female and male faces. It might be that hair style was a better feature to attend to for female than for male faces, so one might for example find interactions between sex of face and initial orienting instructions in a 'levels of processing' type of experiment. The kinds of attributions that subjects make may also differ for female and male faces, which could explain why there are sometimes interactions between sex of face and sex of subjects in memory experiments, though such effects are small.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES