PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

Cognitive Mechanisms of Face Processing [and Discussion]

Andrew W. Ellis and E. T. Rolls

Phil. Trans. R. Soc. Lond. B 1992 **335**, 113-119
doi: 10.1098/rstb.1992.0014

Cognitive mechanisms of face processing

ANDREW W. ELLIS

Department of Psychology, University of York, York YO1 5DD, U.K.

SUMMARY

Evidence from natural and induced errors of face recognition, from the effects of different cues on resolving errors, and from the latencies to make different decisions about seen faces, all suggest that familiar face recognition involves a fixed, invariant sequence of stages. To recognize a familiar face, a perceptual description of a seen face must first activate a long-standing representation of the appearance of the face of the familiar person. 'Semantic' knowledge about such things as the person's occupation and personality are accessed next, followed, in the final stage, by the name.

Certain factors affect the ease of familiar face recognition. Faces seen in the recent past are recognized more readily (repetition priming), as are distinctive faces, and faces preceded by those of related individuals (associative priming). Our knowledge of these phenomena is reviewed for the light it can shed upon the mechanisms of face recognition. Four aspects of face recognition – graded similarity effects and part-to-whole completion in repetition priming, prototype extraction with simultaneous retention of information about individual exemplars, and distinctiveness effects in classification and identification – are proposed as being compatible with distributed memory accounts of cognitive representations.

1. OBLIGATORY STAGES IN THE RECOGNITION OF FAMILIAR FACES

Over the course of the past ten years, psychologists studying the recognition of familiar faces have become increasingly convinced that this process involves an invariant series of stages which must be traversed if successful recognition is to be achieved. This conviction grew in response to converging evidence from both cognitive psychology and neuropsychology (Young, this symposium).

When a face is seen, a perceptual description must be formed which characterizes the face's appearance. In the influential model of Bruce & Young (1986) this is the stage of 'structural encoding'. To know that a face is familiar it must activate a long-standing visual representation of the appearance of the familiar person. Within the Bruce & Young model, those representations are the 'face recognition units'. If a match with a long-standing visual representation is achieved, current accounts would have it that the perceiver is aware only that the face is familiar. No further information is available until the next stage has been successfully negotiated. This involves accessing 'semantic' information stored in long-term memory which indicates such things as the familiar person's occupation, likes and dislikes, where he or she is known from, and so on. The last stage in familiar face recognition, according to most theories, is name retrieval. Names are abstract phonological representations capable of being translated into speech output (or orthographic representations capable of being translated into written output), and are held to be stored separately from semantic representations but accessed via them.

If something like these stages exist in something like this sequence, then problems in face recognition might arise if the processing of a seen face 'blocks' at a certain point. It is trivial to observe that inadequate structural encoding may result in a face not being recognized (e.g. the face is seen too briefly or in poor light). It is more interesting to contemplate the consequences of the inability of a structural description to activate the appropriate long-standing visual representation. Were this to happen, the face of a familiar person would be seen clearly, yet that person would not be recognized as familiar. Errors of precisely this type were documented by Young *et al.* (1985) in an analysis of naturally occurring recognition difficulties. Other naturally occurring problems included occasions when all that the perceiver knew was that the face inducing the error was familiar: for a while at least, the perceiver was unable to 'place' the person or recall their name. This is precisely the experiential state one would predict if a block occurred between the representations of the appearance of a familiar face and the semantic information regarding that person. A temporary blockage between semantic information and name retrieval is held responsible for those errors in which the perceiver can remember everything about a seen person except their name.

The errors the theory forbids are as interesting as the errors it predicts. For example, because names are accessed via semantic information, the theory does not allow errors to occur in which a familiar person's name is known in the absence of semantic information such as where the person is known from or what he or she does. Parallel access to semantics and spoken word-forms is widely believed to occur for the recogni-

Phil. Trans. R. Soc. Lond. B (1992) **335,** 113–119
Printed in Great Britain

tion of familiar written words, including names, but seems not to happen for faces (Ellis *et al.* 1987*a*). No name-without-semantics errors were recorded by Young *et al.* (1985), and none have been observed in subsequent studies which have induced large numbers of recognition errors in the laboratory (Hanley & Cowell 1988; Hay *et al.* 1991).

If different errors or blockages arise at different points in the sequence of stages, then different forms of additional information (cues or prompts) may help to resolve them. Hanley & Cowell (1988) induced errors in the laboratory by using photographs of celebrities. If a given photograph went completely unrecognized, a different photograph helped the subject attain recognition (in terms of the theory, helped to activate a long-standing visual representation), but a second photograph was unhelpful in resolving familiarity-only or name block states (because in such states a visual representation has already been activated, so the second photograph is redundant). Semantic information about the celebrity, provided verbally, helped resolve familiarity-only errors (in which a visual representation has been activated but semantic information has not), but did not help resolve name blocks (because semantics are already activated in such states). Finally, providing the celebrity's initials helped resolve naming problems but was relatively unhelpful in overcoming other types of difficulty. In a similar study, Brennen *et al.* (1990) asked subjects to name celebrities on the basis of brief semantic descriptions. When subjects were blocked on the name, showing the celebrity's face had no beneficial effect (because semantics have already been accessed and there is no direct link from long-standing visual representations to names), but supplying the celebrity's initials was beneficial.

It seems reasonable to propose that the further one has to progress through a sequence of obligatory stages, the longer it will take to access desired information. The simple feeling of familiarity that comes with activating a long-standing visual representation should be attained faster than semantic knowledge about a seen face which should, in turn, be attained faster than knowledge of the name. Young *et al.* (1986*a*) reported four experiments which consistently showed that subjects were faster to decide that a famous face was familiar than they were to assign it to a semantic category (politician or television personality). Young *et al.* (1986*b*) reported a further four experiments showing that subjects can make semantic category decisions faster to famous faces than they can name those faces. It might be argued that naming takes longer than semantic categorization for a variety of reasons other than the number of information processing stages needing to be traversed (e.g. the differing number of response alternatives). Young *et al.* (1988*a*) used matching tasks in which subjects had to decide whether two simultaneously presented faces came from the same semantic category or shared the same first name. Once again, semantic decisions were faster than name decisions. Johnston & Bruce (1990), noting that occupation decisions involve broad categories whereas names are properties of individuals, argued that it is better to contrast name decisions with semantic property decisions. Their tasks involved matching pairs of famous faces according to whether they shared semantic properties (e.g. American or British; dead or alive) or shared first names (James or John). Even in this fairest of comparisons, semantic information was accessed faster than name information, as predicted by the sequential stage theory.

2. REPETITION (IDENTITY) PRIMING

A face may be recognized with different degrees of fluency on different occasions. Several factors have been shown to influence the speed and accuracy of familiar face recognition. Three will be considered here – (i) repetition priming, (ii) distinctiveness effects, and (iii) associative priming – with a view to discovering what they can teach us about the processes involved in recognizing familiar faces.

A face will be recognized more readily if the same face has been encountered some time earlier. This is the phenomenon of repetition (or identity) priming. Repetition priming has most commonly been shown in studies using the familiarity decision task in which subjects are shown a mixture of familiar faces (usually celebrities) and unfamiliar faces and are required to indicate as rapidly as possible (usually by pressing one of two response buttons) whether each face is familiar or unfamiliar. In such a task, responses are faster when the face being responded to in the test phase has been seen earlier.

Repetition priming can be remarkably long-lasting (occurring over several weeks in the experiments of Flude *et al.* 1991), but there are several restrictions on its occurrence. Familiarity decisions to faces are not facilitated by earlier encounters with the names of the celebrities (Bruce & Valentine 1985; Ellis *et al.* 1987*b*). In addition, Ellis *et al.* (1990) found that some decisions made to familiar faces benefit from repetition, but that others do not. In that study, repetition priming was observed for familiarity decisions (Is this face familiar or unfamiliar?) but not for expression decisions (Is this face smiling or unsmiling?) or sex decisions (Is this face male or female?). Thus deciding that a famous face is smiling does not facilitate making the same decision to the same photograph some minutes later. Processing the same familiar face twice is not a sufficient condition for repetition priming to be observed, even when the same decision is taken on both occasions. These findings create problems for explanations of repetition priming couched in terms of enhanced structural encoding or episodic memory for particular faces or face-decision combinations (see, for example, Jacoby (1983); Roediger *et al.* (1989)).

Neuropsychological (Young, this symposium) and neurophysiological (Rolls, this symposium) evidence suggests that expression processing is to some degree independent of identity processing. Cognitive psychological studies support this view. Bruce (1986*a*) and Young *et al.* (1986*c*) found that no difference between familiar and unfamiliar faces on the time taken to make various expression decisions. The results of Ellis *et al.* (1990), along with demonstrations of repetition priming for faces in tasks involving semantic occupation decisions (Young *et al.* 1986*a*) and naming (Ellis *et al.* 1989), indicate that repetition priming only occurs in tasks which require that subjects respond to the identity of a seen face.

Burton *et al.* (1990) developed a computer simulation of face recognition based on 'interactive activation' principles in which individual processing units within one pool corresponded to the face recognition units of the Bruce & Young (1986) model. Repetition priming was implemented in terms of changes in the strengths of connections between 'face recognition units' (FRUS) and units called 'person identity nodes' (PINS). The PINS receive inputs from names and voices as well as faces and serve as a common route from different inputs to a pool of semantic units (see Bruce *et al.*, this symposium). Locating repetition priming at that locus allows an explanation to be given of why tasks which do not demand recognition of faces do not show priming, and why priming does not cross from names to faces.

It is worth noting in passing that Ellis *et al.* (1990) found repetition priming of familiarity decisions when subjects had previously made expression or sex decisions to the same faces. This suggests that although expression and sex decisions do not require the activation of FRUS, such activation is 'unstoppable' (Fodor 1983) and generates priming when the later decision requires identification of the face.

Can repetition priming tell us anything about the nature of the long-standing visual representations that mediate the recognition of familiar faces? Bruce (1986*b*) and Bruce & Valentine (1985) found maximum repetition priming between two identical photographs of a famous person, with a lesser, but still significant, degree of priming between two different photos of the same person. Warren & Morton (1982) obtained a similar pattern for recognition of object drawings. Ellis *et al.* (1987*b*) compared repetition priming of familiarity decisions under conditions when subjects had previously seen the same photographs of celebrity faces or photos of the same celebrities which independent raters had judged to be similar to or dissimilar from the photographs used in the familiarity decision task. Repetition priming was found to be maximal between identical photographs, less between similar photographs and least, but still present, between dissimilar photographs.

As Ellis *et al.* (1987*b*) observed, these results are hard to explain in terms of models in which the representations mediating familiar face recognition are abstractive, localized face recognition units of the sort envisaged by Bruce & Young (1986) and Burton *et al.* (1990). They are more readily assimilated to a model in which the representations within the face recognition system are of the sort proposed in parallel distributed processing, constraint satisfaction ('connectionist') models (Young & Bruce 1991). In terms of such models, a familiar face would be represented as a pattern of activity spread across units coding representational primitives or 'microfeatures' which may or may not correspond to the commonsense notion of a facial feature. McClelland & Rumelhart (1985) show that graded priming effects that are a function of congruity between prime and target are a property of such systems.

Much of the early progress in studying the cognitive basis of familiar face recognition was made by explicitly comparing face recognition with word recognition (Ellis *et al.* 1987*a*). There seems indeed to be many similarities between the two, but sensitivity of repetition priming to changes in surface form may be one important respect in which they differ. Several papers have reported as much repetition priming between, say, two different written forms of a word as between identical versions (e.g. Clarke & Morton 1983; Scarborough *et al.* 1977; Feustel *et al.* 1983; Carr *et al.* 1989). This finding, along with converging evidence from neuropsychology (Caramazza & Hillis 1991) and eye movement studies (Rayner *et al.* 1980) supports the view that the representations that mediate visual word recognition are more abstract in nature (Coltheart 1981) than the representations that mediate the recognition of familiar faces. Bruce & Valentine (1985) considered the possibility that the same-photograph advantage in repetition priming for faces may reflect the contribution of an episodic, pictorial memory of the actual stimulus in addition to a more abstractive priming effect responsible for the priming between different views. The lack of an effect of changes in surface form on repetition priming for words casts doubt on that notion because it is hard to see why there would be no pictorial memory available to contribute an identical-prime advantage to word recognition when there is one available to contribute to face recognition.

Distributed memory systems manifest a property called 'part-to-whole-completion' whereby part of a learned pattern will activate the internal representation of the whole by a process known as autocorrelation (McClelland *et al.* 1986). Brunas *et al.* (1990) and Brunas-Wagstaff *et al.* (1991) found that recognizing a celebrity from part of that person's face (the internal or external features only) yielded as much repetition priming when the full face was later seen in a familiarity decision task as initially recognizing the celebrity from the full-face view. Generalization was also observed from a fragment of one photograph of a celebrity to a different photograph of the same person. A limitation on part-to-whole completion was documented by Ellis *et al.* (1987*b*) who found no repetition priming from recognizing familiar individuals from their clothed but headless bodies to later recognizing the same individuals' faces. This could be construed as possible evidence for a face-specific recognition system, but might also be attributed to a lack of correlation between the patterns representing bodies and faces within the same system.

3. PROTOTYPICALITY, TYPICALITY AND DISTINCTIVENESS

Distributed memory systems are capable of extracting the prototype from a series of exemplars while retaining information regarding the particular exemplars shown (McClelland & Rumelhart 1985; Hintzman

1986). Bruce *et al.* (1991) created different versions of the same faces by changing the relative placement of the internal features. Memory for those different faces operated in a way which enhanced responses to 'prototypical' configurations, even when they had not been studied. Sensitivity to the individual exemplars studied was retained at the same time.

Given a series of faces, subjects will rate some as being relatively typical while others are more distinctive. Valentine & Bruce (1986*a*, *b*; Valentine 1991) have shown that distinctiveness can have different effects on face processing depending on the nature of the task. If the task is to distinguish real from jumbled faces, then typical faces are classified as faces faster than are distinctive faces. If, however, the task is to indicate whether each seen face is familiar or unfamiliar, then the effect is reversed and distinctive faces are recognized faster than typical faces. Valentine & Ferrara (1991) have shown that distributed memory models can give an account of both these effects. They report results of three simulations, two using a single-layer, auto-associative network, the third a multi-layer network using backward error propagation. For all three simulations, general category labels (equivalent to face–non face classification) were learned more readily for typical than distinctive patterns while unique labels (equivalent to the recognition of an individual) were learned more readily for distinctive than typical patterns.

Burton *et al.* (1990) offer a simulation of distinctiveness effects within their interactive activation model. Briefly, FRUs receive inputs from sets of 'feature units'. Typical faces share more features in common than distinctive faces. Hence, when the units representing the features of a typical face are activated, several FRUs receive inputs and try to rise in activation. This leads to mutual inhibition between the FRUs, with the result that the winner takes longer to emerge.

Before we grow too sanguine, it is perhaps worth noting that distinctiveness effects have also been studied in the domain of visual word recognition and have been considered to be compatible with connectionist accounts. The problem is that distinctiveness appears to exert the opposite effect in word recognition as compared with face recognition! Andrews (1989) defined distinctiveness in terms of orthographic 'neighbourhood size'; the number of different words that can be created by changing single letters in a target word. Faster lexical decision responses were found for high neighbourhood (typical) words than to low neighbourhood (distinctive) words. Andrews' account for this is couched in interactive activation terminology, but in this case the advantage for typical words is explained (though not simulated) by arguing that 'partial activation of the representations of neighbours of the target results in reactivation of the constituent letters of activated word units and consequently boosts the activation of the target node'. Heads I win, words you lose.

Before leaving distinctiveness, it is perhaps worth noting Grainger's (1990) claim that, as far as word recognition is concerned, what matters is not so much the size of the orthographic neighbourhood as whether that neighbourhood contains a powerful local attractor in the form of a high frequency word that is similar in appearance to the target word. As applied to face recognition, this would be a matter of whether there existed a very familiar face which resembled the face currently being processed. This is more likely to be true of a typical face than a distinctive face, so it may not be possible to assert at present exactly what underlies the distinctiveness effects reported.

Despite these difficulties, four properties of human face processing – (i) graded similarity effects and (ii) part-to-whole completion in repetition priming, (iii) prototype extraction with simultaneous retention of information about individual exemplars, and (iv) distinctiveness effects in classification and identification – all suggest that the next generation of simulations of face processing should look towards distributed memory systems for their characterizations of the long-standing visual representations that mediate familiar face recognition.

4. ASSOCIATIVE (SEMANTIC) PRIMING

Leaving now the properties of long-standing visual representations, but continuing the associationist theme, this section is concerned with studies of what used to be called 'semantic priming' but may be better referred to as 'associative priming'. Bruce (1983) and Bruce & Valentine (1986) reported the basic effect whereby the recognition of a famous face (e.g. Prince Charles) is faster if that face is preceded by the face of a related person (Princess Diana) than if preceded by the face of an unrelated or unknown person. Bruce (1986*b*) found that one unrelated famous face interposed between the prime and the target was sufficient to abolish the associative priming effect, whereas repetition priming could still be observed after 11 intervening faces. Comparable results for priming of word recognition were reported by Dannenbring & Briand (1982). Associative priming can be observed when the interval between prime and target is very short (250 ms in the study by Bruce & Valentine 1986). This, plus the fact that reaction times are not slowed by an unrelated famous prime as compared with an unfamiliar face prime, is taken as evidence that the effect does not depend (entirely) on conscious anticipation by the subject of likely upcoming faces.

Unlike repetition priming, associative priming crosses between faces and names (Young *et al.* 1988*b*). This has been taken to indicate a more central locus than repetition priming, and associative priming is commonly held to be mediated by connections between central semantic representations for familiar people (see Bruce *et al.*, this symposium).

An unresolved issue is whether this form of priming is best regarded as semantic (e.g. based on common category membership) or associative (based on spatial or temporal co-occurrence). Does Prince Charles prime Princess Diana by virtue of the fact that they are both royals (a semantic basis) or because we see them together (an associative basis)? The nature of semantic–associative priming has been debated in the literature on word recognition, where 'semantic' priming was first documented. Fischler (1977) and Seidenberg *et al.* (1984) obtained priming in a lexical decision task between primes and targets which are semantically related but do not occur together in word association norms (e.g. BREAD–CAKE). This result appears to favour a genuinely semantic basis, but the associatively minded theorist might adopt one of several arguments in an effort to discredit this finding. First, association norms are typically produced by asking large numbers of subjects to generate small numbers of associations to target words. Although CAKE might not be among anyone's first three associations to BREAD, it could well emerge if subjects were required to produce 10 or 20 associations to each target. Second, association norms may just be imprecise guides to patterns of co-occurrence in language, and co-occurrence should be what really matters. Third, if associations are formed in a modality-free semantic system which mediates comprehension of the physical world as well as of language, the co-occurrence of bread and cake in the world of nonverbal experience may be sufficient to cause the words BREAD and CAKE to prime one another associatively. Lupker (1985) reviewed and reanalysed experiments on priming of both word naming (i.e. reading aloud single words) and naming of object pictures. He argued that in both cases there is strong evidence for associative priming but not for semantic priming based on common category membership.

By using a lexical decision task, Napps (1989) found priming between pairs of associated words (e.g. NURSE–DOCTOR) but not between pairs of synonyms (e.g. FOWL–POULTRY). For all their closeness of meaning, synonyms tend to be rivals for particular slots in sentences and may co-occur fairly infrequently. Hence, lack of priming between them argues for an associative basis for priming. Goodman *et al.* (1981), Lukatela *et al.* (1983) and Seidenberg *et al.* (1984) found priming between pairs of words which co-occur in sentences but not in association norms (e.g. HE–SENT), indicating that co-occurrence in sentences may be what is needed for priming to be observed. Finally, Durgunuglu & Neely (1987) obtained priming between unrelated words which had earlier been paired in a paired-association learning task, suggesting again that co-occurrence without semantic relatedness may be sufficient to generate priming.

The question of whether the type of priming under consideration in this section is properly regarded as semantic or associative remains to be finally resolved (or rephrased). Given the apparent differences highlighted above between faces and words we should be cautious about generalizing results from one domain to the other. It is certainly possible, though, to defend the position that all apparently semantic priming is in fact associative. If that proves to be the case, one option would be to argue, with Fodor (1983), that a non-associative semantic system exists at a deeper level than the associative networks that generate priming. Another option would be to argue that semantic knowledge is fundamentally associative in nature (Hinton 1981), and that all psychological phenomena of an apparent semantic nature can be captured by associative systems (Hinton & Shallice 1991; Farah & McClelland 1991).

5. CONCLUDING REMARKS

Commentators on cognitive psychology are prone to argue that the subject all too easily degenerates into the self-perpetuating pursuit of laboratory effects. One antidote to that tendency is to investigate phenomena with obvious real-world relevance, and to keep in the forefront of one's mind the question of what the adaptive value of the phenomena under investigation might be. What, then, are the possible adaptive benefits conferred by distinctiveness effects, repetition priming and associative priming?

It is not obvious what the adaptive value of distinctiveness effects might be: what benefit could be attached to recognizing a distinctive face faster than a typical one? Distinctiveness effects may be just an emergent property of the particular structure that has evolved for recognition systems. If so, it does not invalidate their study. Far from it: the analysis of distinctiveness effects can be expected to tell us much about the internal organization of recognition processes. On the other hand, Valentine (1991) has argued that a common mechanism may underlie distinctiveness effects and the same-race advantage in recognition that has been reported on several occasions. It is at least conceivable that at some point in our evolutionary history, circumstances may have been such that selection pressures favoured the rapid recognition of a face as belonging to a member of a group or type different from one's own. In monkeys, apes or early hominids the mechanisms proposed to underlie distinctiveness effects would also lead to the rapid identification of an individual as belonging to a species, subspecies or group with different facial characteristics. One can imagine circumstances in which that might have been adaptive.

It is easier to give at least a superficially plausible evolutionary account of associative and repetition priming as ways of using past experience to tune our perceptual systems to be biased towards recognizing faces that are likely to occur on a particular occasion (Ellis *et al.* 1987*a*). Associative priming exploits past co-occurrences between familiar stimuli to enable the recognition of one to predict the likely recognition of the other. The absence of inhibition for unrelated faces means that while there is an advantage to being likely, there is no penalty attached to being unexpected. The cumulative effect of repetitions will be to tune the perceptual system to be biased towards recognizing stimuli which have been frequently encountered in the past. That is, repetition will result in frequency effects in recognition whereby commonly encountered stimuli are recognized faster than uncommon ones (cf. Morton 1979).

One of the puzzling things about repetition priming, though, is why just one more encounter with an already highly familiar stimulus may be sufficient to cause that stimulus to be recognized more rapidly several weeks later. One possible explanation is that the word, face or object has not been encountered in that place before, and that repetition priming reflects, at least in part, the learning of an association between a face and a context. I may see Princess Diana's face regularly through the media, but I still have no reason to expect to meet her in the cognitive psychology laboratory of my local Department of Psychology. When I do, albeit in vestigial form as a black-and-white slide, I may form an association between her face and that place, an association which will facilitate recognition of Princess Diana in the same location at any future time. Flude *et al.* (1991) have shown that if the prime face is perceived in one context and target in another, repetition priming is reduced, although not eliminated.

The various processes discussed may serve an adaptive purpose (or may have at some time in our evolutionary history). However, one thing is certain: large communities, increased mobility, the invention of photography, film, television, and so on, have resulted in our face processing apparatus being stretched far beyond anything for which it might have evolved. Seen in that light, it is perhaps easier to understand why, after millions of years of evolution, our face recognition apparatus still seems prone to the sorts of errors and lapses with which this paper began.

The ideas reported here have evolved in the course of discussions with many people, notably Andy Young, Brenda Flude, Dennis Hay, Vicki Bruce, Mike Burton and Hadyn Ellis. Some of the work reviewed in this paper has been funded by an ESRC programme award for a multi-centred investigation into face recognition (grant XC15250001 to Vicki Bruce at Nottingham University; XC15250002 to Ian Craw at Aberdeen University; XC15250003 to Hadyn Ellis at University of Wales College of Cardiff; XC15250003 to Andy Ellis and Andy Young at Lancaster University and XC15250005 to David Perrett at St. Andrews University). I am grateful to Vicki Bruce, Philip Quinlan and Andy Young for comments on an earlier draft of this paper.

REFERENCES

Andrews, S. 1989 Frequency and neighborhood effects on lexical access: Activation or search? *J. Exp. Psychol. Learn. Mem. Cogn.* **15**, 802–814.

Brennen, T., Baguley, T., Bright, J. & Bruce, V. 1990 Resolving semantically induced tip-of-the-tongue states for proper nouns. *Mem. Cogn.* **18**, 339–347.

Bruce, V. 1983 Recognising faces. *Phil. Trans. R. Soc. Lond.* B **302**, 423–436.

Bruce, V. 1986 Influences of familiarity on the processing of faces. *Perception* **15**, 387–397.

Bruce, V. 1986*b* Recognising familiar faces. In *Aspects of face processing* (ed. H. D. Ellis, M. A. Jeeves, F. Newcombe & A. W. Young), pp. 107–117. Dordrecht: Martinus Nijhoff.

Bruce, V. & Valentine T. 1985 Identity priming in the recognition of familiar faces. *Br. J. Psychol.* **76,** 363–383.

Bruce, V. & Valentine, T. 1986 Semantic priming of familiar faces. *Q. Jl exp. Psychol.* **38A**, 297–318.

Bruce, V. & Young, A.W. 1986 Understanding face recognition. *Br. J. Psychol.* **77**, 305–327.

Bruce, V., Doyle, A., Dench, N. & Burton, M. 1991 Remembering facial configurations. *Cognition* **38**, 109–144.

Brunas, J., Young, A.W. & Ellis, A.W. 1990 Repetition priming from incomplete faces: evidence for part-to-whole completion. *Br. J. Psychol.* **81**, 43–56.

Brunas-Wagstaff, J., Young, A.W. & Ellis, A.W. 1991 Repetition priming follows spontaneous but not prompted recognition of familiar faces. *Q. Jl exp. Psychol.* (In the press.)

Burton, A.M., Bruce, V. & Johnston, R.A. 1990 Understanding face recognition with an interactive activation model. *Br. J. Psychol.* **81**, 361–380.

Caramazza, A. & Hillis, A.E. 1991 Levels of representation, co-ordinate frames, and unilateral neglect. *Cogn. Neuropsychol.* **7**, 391–445.

Carr, T.H., Brown, J.S. & Charalambous, A. 1989 Repetition and reading: perceptual encoding mechanisms are very abstractive but not very interactive. *J. exp. Psychol. Learn. Mem. Cogn.* **15**, 763–778.

Clarke, R. & Morton, J. 1983 Cross modality facilitation in tachistoscopic word recognition. *Q. Jl exp. Psychol.* **35A**, 79–96.

Coltheart, M. 1981 Disorders of reading and their implications for models of normal reading. *Vis. Lang.* **15**, 245–286.

Dannenbring, G.L. & Briand, K. 1982 Semantic priming and the word repetition effect in a lexical decision task. *Can. J. Psychol.* **36**, 435–444.

Durgunuglu, A.Y. & Neely, J.H. 1987 On obtaining episodic priming in a lexical decision task following paired-associate learning. *J. exp. Psychol. Learn. Mem. Cogn.* **13**, 206–222.

Ellis, A.W., Young, A.W. & Hay, D.C. 1987*a* Modelling the recognition of faces and words. In *Modelling cognition* (ed. P. E. Morris), pp. 269–297. Chichester: J. Wiley.

Ellis, A.W., Young, A.W., Flude, B.M. & Hay, D.C. 1987*b* Repetition priming of face recognition. *Q. Jl exp. Psychol.* **39A**, 193–210.

Ellis, A.W., Young, A.W. & Flude, B.M. 1989 Two loci for repetition effects in the recognition of familiar faces. Paper presented to Experimental Psychology Society, Cambridge, 5–7 July, 1989.

Ellis, A.W., Young, A.W. & Flude, B.M. 1990 Repetition priming and face processing: Priming occurs within the system that responds to the identity of a face. *Q. Jl exp. Psychol.* **39A**, 193–210.

Farah, M.J. & McClelland, J.L. 1991 A computational model of semantic memory impairment: modality-specificity and emergent category-specificity. *J. exp. Psychol. Gen.* (In the press.)

Feustel, T.C., Shiffrin, R.M. & Salasoo, A. 1983 Episodic and lexical contributions to the repetition effect in word identification. *J. exp. Psychol. Gen.* **112**, 309–346.

Fischler, I. 1977 Semantic facilitation without association in a lexical decision task. *Mem. Cogn.* **5**, 335–339.

Flude, B.M., Young, A.W. & Ellis, A.W. 1991 Long-term perceptual memory for faces. Paper presented to *International Conference on Memory*, Lancaster, 15–19 July, 1991.

Fodor, J.A. 1983 *The modulatory of mind.* Cambridge, Massachusetts: MIT Press.

Goodman, G.O., McClelland, J.L. & Gibbs, R.W. 1981 The role of syntactic context in word recognition. *Mem. Cogn.* **9**, 580–586.

Grainger, J. 1990 Word frequency and neighborhood frequency effects in lexical decision and naming. *J. Mem. Lang.* **29**, 228–244.

Hanley, J.R. & Cowell, E.S. 1988 The effects of different types of retrieval cues on the recall of names of famous faces. *Mem. Cogn.* **16**, 545–555.

Hay, D.C., Young, A.W. & Ellis, A.W. 1991 Routes through the face recognition system. *Q. Jl exp. Psychol.* **43A,** 761–791.

Hinton, G.E. 1981 Implementing semantic networks in parallel hardware. In *Parallel models of associative memory* (ed. G. E. Hinton & J. A. Anderson), pp. 161–188. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

Hinton, G.E. & Shallice, T. 1991 Lesioning an attractor network: investigations of acquired dyslexia. *Psychol. Rev.* **98**, 74–95.

Hitzman, D.L. 1986 "Schema abstraction" in a multiple-trace memory model. *Psychol. Rev.* **93**, 411–428.

Jacoby, L.L. 1983 Perceptual enhancement: persistent effects of an experience. *J. exp. Psychol. Learn. Mem. Cogn.* **9**, 21–38.

Johnston, R.A. & Bruce, V. 1990 Lost properties? Retrieval differences between name codes and semantic codes for familiar people. *Psychol. Res.* **52**, 62–67.

Lukatela, G., Kostic, A., Feldman, L.B. & Turvey, M.T. 1983 Grammatical priming of inflected nouns. *Mem. Cogn.* **11**, 59–63.

Lupker, S.J. 1985 Context effects in word and picture recognition: a reevaluation of structural models. In *Progress in the psychology of language*, vol. 1 (ed. A. W. Ellis), pp. 109–142. London: Lawrence Erlbaum Associates.

McClelland, J.L. & Rumelhart, D.E. 1985 Distributed memory and the representation of general and specific information. *J. exp. Psychol. Gen.* **114**, 159–188.

McClelland, J.L., Rumelhart, D.E. & Hinton, G.E. 1986 The appeal of parallel distributed processing. In *Parallel distributed processing: explorations in the microstructure of cognition. Volume 1: Foundations* (ed. J. L. McClelland, D. E. Rumelhart & PDP Research Group), pp. 3–44. Cambridge, Massachusetts: MIT Press.

Morton, J. 1979 Word recognition. In *Psycholinguistics Series*, vol. 2 (ed. J. Morton & J. C. Marshall), pp. 107–156. London: Elek.

Napps, S.E. 1989 Morphemic relations in the lexicon: are they distinct from semantic and formal relationships? *Mem. Cogn.* **17**, 729–739.

Rayner, K., McConkie, G.W. & Zola, D. 1980 Integrating information across fixations. *Cogn. Psychol.* **12**, 206–226.

Roediger, H.L., Weldon, M.S. & Challis, B.H. 1989 Explaining dissociations between implicit and explicit measures of retention: a processing account. In *Varieties of memory and consciousness: essays in honour of Endel Tulving* (ed. H. L. Roediger & F. I. M. Craik), pp. 3–41. New Jersey: Lawrence Erlbaum Associates.

Scarborough, D.L., Cortese, C. & Scarborough, H.S. 1977 Frequency and repetition effects in lexical memory. *J. exp. psychol. Hum. Percept. Perform.* **3**, 1–17.

Seidenberg, M.S., Waters, G.S., Sanders, M. & Langer, P. 1984 Pre-and postlexical loci of contextual effects on word recognition. *Mem. Cogn.* **12**, 315–328.

Valentine, T. 1991 A unified account of the effects of distinctiveness, inversion and race in face recognition. *Q. Jl exp. Psychol.* **43A**, 161–204.

Valentine, T. & Bruce, V. 1986*a* The effects of distinctiveness on recognising and classifying faces. *Perception* **15**, 525–535.

Valentine, T. & Bruce, V. 1986*b* Recognising familiar faces: the role of distinctiveness and familiarity. *Can. J. Psychol.* **40**, 300–305.

Valentine, T. & Ferrara, A. 1991 Typicality in categorization, recognition and identification: evidence from face recognition. *Br. J. Psychol.* **82**, 87–102.

Warren, C. & Morton, J. 1982 The effects of priming on picture recognition. *Br. J. Psychol.* **73**, 117–130.

Young, A.W. & Bruce, V. 1991 Perceptual categories and the computation of "grandmother". *Eur. J. cogn. Psychol.* **3**, 5–49.

Young, A.W., Hay, D.C. & Ellis, A.W. 1985 The faces that launched a thousand slips. *Br. J. Psychol.* **76**, 495–523.

Young, A.W., McWeeny, K.H., Hay, D.C. & Ellis, A.W. 1986 Access to identity-specific semantic codes from familiar faces. *Q. Jl exp. Psychol.* **38A**, 271–295.

Young, A.W., McWeeny, K.H., Ellis, A.W. & Hay, D.C. 1986 Naming and categorising faces and written names. *Q. Jl exp. Psychol.* **38A**, 297–318.

Young, A.W., McWeeny, K.H., Hay, D.C. & Ellis, A.W. 1986 Matching familiar and unfamiliar faces on identity and expression. *Psychol. Res.* **48**, 63–68.

Young, A.W., Ellis, A.W. & Flude, B.M. 1988*a* Accessing stored information about familiar people. *Psychol. Res.* **50**, 111–115.

Young, A.W., Hellawell, D. & de Haan, E.H.F. 1988*b* Cross-domain semantic priming in normal subjects and a prosopagnosic patient. *Q. Jl exp. Psychol.* **40A**, 561–580.

Discussion

E. T. Rolls (*Department of Experimental Psychology, University of Oxford, U.K.*). It would be particularly impressive if priming produced by presentation of part of a face primed another part of the same individual's face. Does this occur? If so, is it fully effective as a priming stimulus?

A. W. Ellis. Experiments by Brunas *et al.* (1990) and Brunas-Wagstaff *et al.* (1991) reviewed in the paper show that recognizing part of a celebrity's face primes subsequent recognition of the whole face as much as when a whole face is initially recognized. In as yet unpublished work we have gone on to look at the effect of recognizing part of a famous face on subsequent recognition of either the same or a different, non-overlapping part. We have used left and right half-faces and faces in which horizontal strips have been covered up. In both cases we find priming from one part to a different part, but less than from one part to itself. Thus, there appears to be full part-to-whole completion when the whole includes the priming part, but less than full completion when the whole is required to mediate the transfer of priming from one part to a different part.