

Section A: Human Experimental Psychology

Editor:

Glyn W. Humphreys (School of Psychology, University of Birmingham, Edgbaston, Birmingham, B15 2TT)

Associate Editors:

Vicki Bruce (University of Nottingham)

Ruth Campbell (Goldsmiths College, London)

Martin A. Conway (University of Lancaster)

John Duncan (MRC Applied Psychology Unit)

Jonathan St. B.T. Evans (Polytechnic South West)

Ken Flowers (University of Hull)

Alan Garnham (University of Sussex)

Dennis Norris (MRC Applied Psychology Unit)

Quentin Summerfield (MRC Institute of Hearing Research, Nottingham)

Book Reviews Editor:

Stuart Sutherland (University of Sussex)

Books for review should be sent to Professor N. S. Sutherland, Laboratory of Experimental Psychology, University of Sussex, Brighton, BN1 9QG, U.K.

Contributors should refer to the notes for authors which appear at the end of Volume 43, Number 1.

The Quarterly Journal of Experimental Psychology, Section A: Human Experimental Psychology is published for The Experimental Psychology Society by Lawrence Erlbaum Associates Ltd., London and Hove, U.K., and Hillsdale, New Jersey.

Subscriptions

Subscription rates for Volume 43A, 1991 (4 issues):

To individuals: £30.00; \$57.00 (USA);

To institutions: £95.00; \$180.00 (USA).

Rates for combined subscription (A and B, 4 issues each):

To individuals: £42.00; \$80.00 (USA);

To institutions: £135.00; \$256.00 (USA).

Undergraduates and graduates of less than three years' standing may subscribe to both volumes at the special reduced rate of £6.00. They should write to the Hon. Treasurer, Dr. Nigel Harvey, Department of Psychology, University College London, Gower Street, London WC1E 6BT, with a letter from a supervisor confirming their status.

All other subscription orders should be addressed to Lawrence Erlbaum Associates Ltd., The Distribution Centre, Blackhorse Road, Letchworth, Herts, SG6 1HN, U.K. Send notices of change of address to the offices of the publishers at least six weeks in advance. Please include both old and new addresses.

Permissions

The Experimental Psychology Society waives permission and fees for the photocopying of isolated articles for non-profit classroom or library reserve use by instructors and educational institutions. This consent does not extend to other kinds of copying, such as copying for general distribution, for advertising or for promotional purposes, for creating new collective works or for resale.

Correspondence

Correspondence for the Experimental Psychology Society should be sent to the Hon. Secretary, Professor A. W. Young, Department of Psychology, University of Durham, South Road, Durham, DH1 3LE, U.K.

Advertising

Advertising enquiries should be sent to the advertising sales office: REW Space Sales, 11 Bourne Way, West Ewell, Surrey KT19 9PP.

A Unified Account of the Effects of
Distinctiveness, Inversion, and Race in Face
Recognition

Tim Valentine

University of Manchester, U.K.

A framework is outlined in which individual faces are assumed to be encoded as a point in a multidimensional space, defined by dimensions that serve to discriminate faces. It is proposed that such a framework can account for the effects of distinctiveness, inversion, and race on recognition of faces. Two specific models within this framework are identified: a norm-based coding model, in which faces are encoded as vectors from a population norm or prototype; and a purely exemplar-based model. Both models make similar predictions, albeit in different ways, concerning the interactions between the effects of distinctiveness, inversion and race. These predictions were supported in five experiments in which photographs of faces served as stimuli. The norm-based coding version and the exemplar-based version of the framework cannot be distinguished on the basis of the experiments reported, but it is argued that a multidimensional space provides a useful heuristic framework to investigate recognition of faces. Finally, the relationship between the specific models is considered and an implementation in terms of parallel distributed processing is briefly discussed.

The rated distinctiveness of a face, the orientation in which it is seen, and the race of a face are all factors known to influence the ability of an observer to subsequently recognize the face. There has been a tendency for each of these factors to be investigated in isolation largely appealing to different theoretical explanations. In the present paper it is proposed that similarity between faces can account for the effects of all three factors. First the current literature on the effects of distinctiveness and inversion is reviewed briefly. In

Requests for reprints should be sent to Tim Valentine, Department of Psychology, University of Manchester, Oxford Road, Manchester M13 9PL, U.K.

The experiments reported here were carried out while the author was employed by the Medical Research Council at the MRC Applied Psychology Unit, 15 Chaucer Road, Cambridge. I would like to thank Gillian Rhodes, Glyn Humphreys, and anonymous reviewers for helpful comments on an earlier draft of the manuscript.

the second section a general framework is proposed in which it is assumed that faces can be represented by a point in a multidimensional space. Within this framework, two models can be identified. One model assumes a norm or prototype is abstracted, the other model is purely exemplar-based. Predictions based on these models concerning the interaction between the effects of distinctiveness and inversion in face processing tasks are tested in Experiments 1–4. The multidimensional space framework is then extended to account for the effect of race. Experiment 5 tests a prediction derived from the framework concerning the interaction between the effects of race and inversion. In the General Discussion the relationship between the norm-based and purely exemplar-based versions of the framework and other models of race recognition are discussed.

Any model of visual object recognition must specify how stored knowledge is used to facilitate recognition of visual stimuli (Palmer, 1975). Palmer cited faces as an example of a perceptual category that would include in its representation information about the prototypical values (or central tendency) of the relevant dimensions. Such information would be specific to face processing. Fodor (1983) proposes a similar view of face processing in his suggestion that faces are “favourite candidates” for an eccentric stimulus domain—that is, a domain “whose perceptual analysis requires information that is highly specific to the domain in question.” (Fodor, 1983, pp. 51–52).

The effect of distinctiveness on the recognition of faces provides some indication that category-specific knowledge, presumably acquired through experience of the population of faces previously seen, is used to facilitate recognition of faces encountered subsequently. Generally, previously unfamiliar faces that are rated as distinctive or unusual are more accurately recognized in a recognition memory paradigm. Conversely, typical faces are more likely to be incorrectly identified as having been seen before (Going & Read, 1974; Cohen & Carr, 1975; Light, Kayra-Stuart, & Hollander, 1979; Winograd, 1981; Bartlett, Hurry, & Thorley, 1984). An effect of distinctiveness has been found in three measures of recognition memory performance—hit rate, false positive rate, and combined measures of sensitivity based on signal detection theory. The exact nature of the effect varies slightly from study to study, but generally the advantage found for distinctive faces in false positive rate and sensitivity measures is robust. The effect on hit rate seems to be dependent upon the procedure and is found most often when the initial encoding conditions are good, involving trait judgement encoding activities and/or relatively long exposure times of 8 sec (Light et al., 1979).

Light et al. (1979) proposed a two-component theory of memory based on inter-stimulus similarity in order to account for the effect of distinctiveness. They suggested that a distinctive face is more likely to access a *specific* memory because it is less similar to specific memories of other faces. Thus

access to specific memory gives rise to an advantage in hit rate for distinctive faces. However, in the absence of a specific memory, Light and colleagues suggested that subjects base their recognition judgement on *schematic* memory of category structure (e.g. similarity to a prototype). Use of schematic memory gives rise to the greater false positive rate to typical faces. In contrast, Bartlett et al. (1984) interpreted the effect of distinctiveness in terms of familiarity information alone (cf. Mandler, 1980). Bartlett and colleagues suggested that presentation of a distinctive face resulted in a greater increment in familiarity than that resulting from presentation of a typical face. Valentine and Bruce (1986b) examined the relationship between the effects of distinctiveness and familiarity in processing highly familiar faces. It was found that both the rated distinctiveness and familiarity of a face affected the RT to decide it was familiar in a familiarity decision task. The more distinctive or the more familiar a face was rated, the faster it could be recognized. However, the two effects were found to be independent, suggesting that familiarity information does not form the basis of the effect of distinctiveness. (See Valentine, 1990, for further discussion of the relationship between distinctiveness and familiarity.)

Valentine and Bruce (1986c) argued that if the effect of distinctiveness on the latency to recognize familiar faces arises from the role of a facial prototype, then distinctive faces should take longer than typical faces to be classified as a face in a task in which faces must be distinguished from jumbled faces. This prediction was confirmed. Thus distinctive faces are recognized faster in a familiarity decision task but classified as faces more slowly than are typical faces. Valentine and Bruce suggest that a general face prototype is extracted from faces previously encountered and that individual faces are stored as a set of transformations required to match the face to the prototype. This proposal was termed “the prototype hypothesis”.

Further evidence of prototype abstraction in face recognition comes from the study of caricature in recognizing faces. Perkins (1975) describes caricature as a process of exaggerating the distinctive features that individuate a particular face. Rhodes, Brennan, and Carey (1987) studied the recognition of computer-generated caricatures. Veridical line drawings of faces were generated by joining the co-ordinates of 169 specified points on each face by lines. A norm was generated by averaging the position of these points across several faces. Caricatures were then generated by increasing the distance between the location at each point in an individual’s face and the norm by a fixed proportion. An anti-caricature could be generated by reducing the differences between a face and a norm. Rhodes and colleagues found that caricatures were recognized faster than veridical line drawings, which in turn were recognized faster than anti-caricatures; however, there were no differences in accuracy in recognizing the three classes of faces. They interpret

their result as evidence of norm-based coding, a holistic encoding process in which the distinctive aspects of faces are encoded by comparison to a norm or average face.

Both Valentine and Bruce (1986b, 1986c) and Rhodes et al. (1987) have, quite independently, proposed very similar accounts of the role of a prototype or norm in face recognition. There are clear parallels between these hypotheses and Goldstein and Chance's (1980) proposal that the role of a face schema can account for the development of the effect of race and inversion in face recognition. Adults recognize other-race faces less accurately than own-race faces. However, young children recognize other-race faces and own-race faces with equal accuracy (Goldstein & Chance, 1980; Chance, Turner, & Goldstein, 1982). Analogous results are also found for the effect of inversion on face recognition. Adults show a large effect of stimulus inversion on their ability to recognize faces. Indeed, the effect of inversion on face recognition is disproportionately large compared to the effect of inversion on recognition of other stimulus classes (Yin, 1969). However, young children show little effect of inversion upon recognition memory for faces (Goldstein, 1975). Goldstein and Chance account for the development of the effects of race and inversion by arguing that with increasing age children become more efficient in their use of a face schema, leading to better face recognition performance. However, this increase in efficiency is accompanied by an increase in "schema rigidity", so that as the schema develops, it becomes relatively less efficient at processing unusual stimuli such as inverted or other-race faces.

Goldstein and Chance's (1980) face schema theory has the basic idea that knowledge of the population of faces is acquired and used in face processing in common with the prototype hypothesis. However the evidence, from the development of the effects of race and inversion on face recognition, used to support schema theory is equivocal. Goldstein and Chance (1980) found considerable improvement in face recognition ability across the age range 6-12 years but did not find an interaction between the effect of race and age in this range. In a later study Chance et al. (1982) did find an Age of Subject \times Race of Stimulus Face interaction in this age range. However, other studies (Cross, Cross, & Daley, 1971; Kagan & Klein, 1973; Feinman & Entwisle, 1976) have failed to find an increase in the effect of race with age. Recent work on the development of the effect of inversion on face recognition has also shown that earlier work has overstated the effect. Although the ability to recognize upright faces improves faster than the ability to recognize inverted faces, young children do show an effect of inversion, and the ability to recognize inverted faces does improve with age (Flin, 1985; see Valentine, 1988, for a review of the effect of inversion on face recognition).

Diamond and Carey (1986) have demonstrated that recognition of another stimulus class is as adversely affected by inversion as is face recognition. It was found that dog experts showed a similarly large effect of inversion on their ability to recognize individual dogs. Diamond and Carey argue that a large effect of inversion will be found if: (1) the exemplars of a stimulus class have a common configuration but subtle differences in the spatial relations between the features (termed second-order relational information); and (2) the observers have sufficient expertise to distinguish the exemplars of the stimulus class on the basis of these differences in configural information. Thus the disproportionate effect of inversion is not specific to faces but likely to be found for recognition of any highly familiar and highly homogeneous stimulus class.

In summary, the prototype hypothesis (Valentine & Bruce, 1986b, 1986c), the norm-based coding model (Rhodes et al., 1987), schema theory (Goldstein & Chance, 1980), and Diamond and Carey's (1986) study of inversion, together, suggest that the effects of distinctiveness, inversion, and race on face recognition could all be explained by the role of knowledge of the population of faces previously encountered. Valentine and Bruce, Rhodes et al., and Goldstein and Chance have all proposed that a prototype, norm, or schema is abstracted. However, alternative accounts can also be formulated, which are based on inter-stimulus similarity but do not assume the existence of a stored face prototype. For instance, this would be in keeping with recent models in the concept learning literature (e.g. Medin & Schaffer, 1978; Nosofsky, 1986).

In the next section a framework is proposed in which faces are assumed to be encoded as points in a multidimensional space. Two specific models are identified. In one model faces are encoded by reference to an abstracted norm or prototype. In order to avoid possible confusion with the different ways in which the term "prototype" has been used in the concept learning and face recognition literature, I will adopt the term used by Rhodes and colleagues and refer to this model as the "norm-based coding model". The other model identified assumes that a norm is not abstracted and that only specific faces (or category exemplars) are stored. This will be referred to as the "exemplar-based model". The term "multidimensional space framework" will be used as a generic term to refer to both the norm-based and exemplar-based models.

The experiments reported do not aim to distinguish between the norm-based and exemplar-based models. Indeed, the two models make similar predictions, although based on some different assumptions. The aim is to show that knowledge of the population of faces can provide a parsimonious account of the effects of distinctiveness, inversion, and race.

A Framework for Coding Processes in Face Recognition

The main assumption of the proposed framework is that a location in a Euclidean multidimensional space provides an appropriate metaphor for the mental representation of a face.¹ The dimensions of the space represent the physiognomic features that are used to encode faces. No attempt will be made to identify the aspects of a face that the dimensions represent. However, previous work using multidimensional scaling techniques suggest that the principal dimensions needed would represent hair colour and length, face shape, and age (at least for the caucasian faces, which are "own-race" for the subjects in the experiments reported here: see Shepherd & Derogowski, 1981). It is assumed that the number of dimensions could be large enough to represent any aspect of a face that could serve to discriminate between faces. At this stage in the discussion it will be assumed that the overwhelming majority of faces encountered are own-race faces. Although the model to be described will be discussed in terms of dimensions rather than features, this does not mean to exclude the possibility that the processes involved may be based on frequency information of discrete feature values. However, dimensions appear to cope more naturally with discrimination within a stimulus class of which all exemplars share a common structure, such as faces. This view is in line with Garner's (1978) definition of a dimension as "an attribute that exists for each stimulus in the relevant set at some positive, mutually exclusive value" (p. 104).

The origin of the multidimensional space is defined as the central tendency of the dimensions. It is assumed that the values of the feature dimensions of the population of faces experienced will vary normally around the central tendency (at least for own-race faces). Therefore by definition, typical faces (close to the central tendency) will be seen more often than distinctive faces (distant from the central tendency). Thus the density of points (i.e. the number of previously seen faces) will decrease as the distance from the central tendency increases. The population of points will include familiar faces in addition to many points representing faces that have been seen previously, but would not necessarily be "familiar" in the sense that they could be identified. Thus, an implicit knowledge derived from a lifetime's experience with faces contributes to the normal distribution of faces within the multidimensional space.

The implicit knowledge of faces is something a subject brings to an experiment, and it is assumed that the effect of the set of faces used within an experiment will have an insignificant influence compared to a lifetime's

¹The assumption of a Euclidean metric is only made for simplicity in the absence of any evidence of the most appropriate metric. As the dimensions are not identified, the metric of the space is not an issue addressed in the current paper.

experience of faces. Therefore, the proposed framework is fundamentally different from the theoretical basis of studies in which sets of artificial faces have been manipulated within an experimental context (e.g. Reed, 1972; Goldman & Homa, 1977; Neumann, 1977; Das-Smaal & De Swart, 1984; Malpass & Hughes, 1986). As the distribution of faces within the multidimensional space is assumed to emerge as a consequence of a lifetime's experience, it is important to use photographs of real faces in order to investigate the influence of such implicit knowledge. Photographs of faces are stimuli that vary along the same dimensions and with the same central tendencies as the population of faces previously experienced. Of course this would not be true of an artificially constructed set of faces (e.g. schematic faces, photo-fit faces), therefore performance in tasks based upon artificial stimulus sets might be unaffected by knowledge of the population of real faces. It is impossible to distinguish distance-based and frequency-based models because the exemplars that contribute to the implicit knowledge of faces cannot be controlled experimentally.

At this point, two specific models based on the multidimensional space framework need to be distinguished. A norm-based coding model will be described first, in which it is assumed that faces are encoded in terms of their deviation from a single general face norm or prototype located at the origin of the space (i.e. representing the central tendency). It is assumed that there is only a single norm (Valentine & Bruce, 1986c). This is quite different from the notion of a different prototype to represent each type or "family" of faces (cf. Ellis & Christie cited in Ellis, 1981). If the number of dimensions is denoted by n , an n -dimensional vector from the norm (or origin) to the point representing the dimension values of a particular face can uniquely specify that face. Figure 1 illustrates the norm-based coding model using just a two-dimensional rather than a multidimensional space for the purposes of illustration only.

The recognition process can be viewed as involving two stages. First, a stimulus face is encoded as an n -dimensional vector. Second, some form of decision process is required to determine whether or not the stimulus matches a vector for a known face. The encoding process is assumed to have some associated error or noise, which will depend upon the encoding conditions. Therefore, under difficult viewing conditions, the vector derived will have a relatively large associated error estimate, which could be represented in Figure 1 as a region of uncertainty around the co-ordinates of the stimulus vector. It is assumed that the confidence signalled by the decision process depends on: (1) the error associated with the vector derived from the stimulus face; (2) a measure of similarity between the vector derived from the stimulus and the vector corresponding to the nearest known face; and (3) the similarity between the stimulus vector and the vector of the next nearest neighbour. The measure of similarity used is assumed to be some




FIG. 1. A two-dimensional representation of a norm-based coding model of face recognition. Each point plotted represents a previously seen face located in an n -dimensional space. The origin represents a general face norm. Two dimensions are used in this figure for the purposes of illustration only. The axes are unlabelled but could represent any dimension that could be used to discriminate faces.

function of a vector similarity measure such as the dot product between two vectors. (See Micko, 1970, for a discussion of vector-based measures of similarity in multidimensional scaling.)

The exemplar-based model assumes that there is no extracted norm or prototype. The model is similar to the norm-based coding model, except that it is more appropriate to consider faces as being encoded as points rather than vectors. The origin of the multidimensional space plays no part in encoding stimuli, it merely indicates the point of maximum exemplar density. The exemplar-based model assumes that similarity between two faces is a monotonic function of the distance separating the representations of the faces in the multidimensional space. The decision process is assumed to

depend on: (1) the estimate of error associated with encoding the stimulus; (2) the distance between the location of the stimulus and the nearest known face; and (3) the distance between the stimulus and the next nearest neighbour.

It should be noted that though both the norm-based coding and the exemplar-based models are nearest neighbour exemplar models, they differ in the role of an abstracted norm in coding and the use of a vector- or distance-based measure of similarity. (See Valentine, 1990, for further discussion of this point.)

Predictions Derived from the Framework

Both the norm-based coding model and the exemplar-based model account for the effects of distinctiveness found in recognition of faces by appealing to exemplar display. Distinctive faces are located in regions where the density of points is low. Therefore, when a distinctive familiar face is encountered in a recognition task, the location encoded in the multidimensional space will be much closer to the representation of the "target" face stored in memory than to the location of another face. (By definition, few faces will resemble a distinctive face.) In addition, the error associated with the encoded location is likely to be small compared to the distance to the nearest neighbour, especially under good encoding conditions. Therefore, the face can be identified accurately and rapidly. If the face presented is a typical familiar face, the location derived from this face is close to the central tendency, so it will fall in a region in which there is a high density of points. Although the location of the stimulus will be close to the location of the target face, it will also probably be close to other previously seen but "unfamiliar" faces. Therefore the decision as to whether the stimulus is closer to the location of the familiar face than to a nearby "unfamiliar" face will be more difficult, rendering it slower and more error-prone. The multidimensional space framework therefore predicts that distinctive familiar faces will be recognized more accurately or more quickly than typical familiar faces.

Consider the case of a face that has not been seen before being encountered. Of course there is no stored description for the face in memory. If the stimulus face is distinctive, the density of points around it will be low. Therefore, it is unlikely that the location of a previously seen face will be close enough to the stimulus with its associated error to give a false positive. If the stimulus face is a typical unfamiliar face, its derived location will be in a region of a high density of points. Therefore, it is more likely to be close to the representation of a previously seen face. If the similarity between the encoded face and a previously seen face is sufficiently high, a false positive will result. Thus the framework predicts that distinctive unfamiliar faces can be rejected more accurately or more quickly than typical unfamiliar faces.

Both the norm-based coding model and the exemplar-based model clearly predict an effect of distinctiveness in correctly rejecting unfamiliar faces in a familiarity decision task. However, Valentine and Bruce (1986c) found a trend in the opposite direction that did not reach statistical significance. It is possible that the failure to find the expected effect in rejecting unfamiliar faces was the result of an artifact of the stimulus sets used. This prediction of the framework is examined in Experiment 3.

The norm-based coding model and the exemplar-based model both account for the effect of distinctiveness on recognition in terms of exemplar density. However, the role of a vector similarity measure in the norm-based coding model produces an effect that works against the effect of distinctiveness. In the norm-based coding model the similarity between two faces that are equidistant in the space is dependent on the distance of the points from the norm. Consider two pairs of points (A, B and C, D) located such that the distance between A and B is equal to the distance between C and D. [A possible arrangement is shown in Figure 2. In this figure A and B are equidistant from the origin (O), as are C and D.] If a similarity measure that is based on the distance between exemplars is used (as in the exemplar-based model), A and B are as equally similar to each other as are C and D. However, the angle between the vectors OA and OB is greater than the angle between the vectors OC and OD, and so the faces represented by OA and OB will be more dissimilar than the faces represented by OC and OD. Thus in the norm-based coding model, faces separated by a given distance will be more difficult to discriminate, the further they are from the norm. The norm-based




FIG. 2. An illustration of the effect of distance from the norm on similarity in the norm-based coding model. The distance between A and B is the same as the distance between C and D, but the angle between vectors OA and OB is greater than the angle OC and OD. See text for explanation.

coding model can still support the prediction that distinctive faces can be more accurately recognized than typical faces because distinctive faces are less densely clustered (i.e. further apart from each other) than typical faces. However, it is necessary to assume that the effect due to the difference in exemplar density is greater than the opposite effect due to distance from the norm on the vector similarity measure.

It was noted above that the effect of distinctiveness found in a face classification task (in which faces must be distinguished from jumbled faces) was the opposite to that found in a recognition task. The classification task can be regarded as a judgement of how closely a stimulus resembles the central tendency of the population of faces. In terms of the norm-based coding model, the RT in this task will depend upon the length of the vector derived from the stimulus face (i.e. its distance from the norm in the n -dimensional space). The density of points in a region will be irrelevant to the task, because there is no requirement to distinguish between faces. Therefore typical faces should be correctly classified as a face faster than distinctive faces, because typical faces are closer to the norm. Thus the norm-based coding model successfully predicts the reversal of the effect of distinctiveness between a recognition and classification task.

According to the exemplar-based model, both recognition and classification of faces will be affected by exemplar density but in quite different ways. In a recognition task high exemplar density makes the nearest neighbour likely to be close to a "target" face, thus impeding the recognition decision. The classification task has no requirement to distinguish between individual faces. In this task exemplar-density in the region around the stimulus is used to judge category membership, the greater the density of faces in this region, the faster the stimulus can be classified as a face. Thus in a classification task it is assumed that RT is determined by the exemplar density in the multidimensional space around the location of the stimulus (Krumhansl, 1978). As distinctive faces are located in areas of lower exemplar density, typical faces will be classified faster than distinctive faces. Both the norm-based coding and the exemplar-based models make the same prediction of the effect of distinctiveness on classification as a face, but the norm-based coding model appeals to distance from the norm and the exemplar-based model appeals to exemplar density.

Some researchers have suggested that inversion disrupts the normal face recognition process to such an extent that different features are used to recognize upright and inverted faces (e.g. Carey & Diamond, 1977; Diamond & Carey, 1986). This view arose from the observation that recognition of faces is more disrupted by inversion than recognition of other stimulus classes (Yin, 1969). Inversion makes face recognition slower and less accurate, but there is no compelling evidence that upright and inverted faces are processed qualitatively differently (see Valentine, 1988, for a review). One

aim of the multidimensional space framework is to account for the effects of inversion on face recognition and classification. It is assumed that the magnitude of the error in deriving the location of a stimulus face is dependent upon the encoding conditions. Presenting a face upside-down is an example of one experimental manipulation that would make the encoding conditions difficult, leading to a large error associated with the location in the multidimensional space derived from a stimulus face. The effect of inversion will be represented in the framework in this way. Any experimental manipulation that impairs recognition could be represented by assuming an increase in the error associated with the location of the stimulus face. Thus the multidimensional space framework makes the strong prediction that other manipulations (e.g. adding visual noise or blurring faces) should interact with distinctiveness in the same way as inversion. However, in this paper the concern is to account for the effect of inversion because the inversion paradigm has had a considerable impact on the face recognition literature.

Representing inversion in this way, the norm-based coding and the exemplar-based models make similar predictions concerning the interaction between the effects of inversion and distinctiveness in a recognition task. A typical unfamiliar face seen inverted will be more likely to be mis-identified as familiar than the same face seen upright. The high density of nearby points and the relatively large error associated with the stimulus location derived from an inverted face means that it is more likely to be sufficiently close to a neighbouring point corresponding to a previously seen face to give a false positive. If a typical familiar face is seen inverted, it is more likely to be missed than when seen upright, because the increased error in deriving the location from the stimulus face means it is likely to be further from the location that forms its representation in memory. A miss is more likely because there is a greater chance that a previously seen but "unfamiliar" face will be closer than the "target". Of course the same arguments apply to inverted distinctive faces. Therefore, inversion should impair recognition of both typical and distinctive faces. However, because the locations derived from distinctive faces will be in regions with a low density of points, a larger error can be tolerated in recognition of distinctive faces than in recognition of typical faces. Therefore, inversion should cause greater disruption to recognition of typical faces than of distinctive faces. This prediction of an interaction between the effects of distinctiveness and inversion was tested using a recognition memory paradigm in Experiment 1.

There are a number of important and as yet unanswered issues raised by the application of a multidimensional space framework to face recognition. Most notably, the dimensions of the space have not been identified, nor has the dimensionality of the space required. The determinants of similarity have also not been precisely specified—for example, the metric of the space has not been determined. In view of our currently imprecise knowledge of the

dimensions underlying face perception, and therefore our impoverished understanding of similarity between faces, it would be premature to precisely specify the parameters of the framework. The aim is to demonstrate that the multidimensional space framework has predictive utility and can be used to guide future research.

A number of predictions can be derived from the multidimensional space framework. In Experiments 1–3, the prediction that recognition of typical faces is more impaired by inversion than recognition of distinctive faces is tested. Experiment 4 demonstrates that this interaction is not found in a task that does not require recognition of individual faces (an intact/jumbled face classification task). The effect of race on face recognition is then considered in relation to the multidimensional space framework. It is predicted that the effect of race will interact with the effects of inversion and task demands in the same way as distinctiveness. This analysis is consistent with data reported by Valentine and Bruce (1986a), in which an interaction between the effects of inversion and race on recognition memory for faces was found. In Experiment 5 it was demonstrated that the effect of race is additive with that of inversion in a face classification task.

EXPERIMENT 1

Method

The experiment was in two parts. First, faces were rated for distinctiveness. Different subjects then carried out a recognition memory task.

Distinctiveness Ratings

Subjects. Sixteen (2 male and 14 female) members of the Applied Psychology Unit subject panel made the distinctiveness ratings; their mean age was 44.9 years.

Materials. The stimuli were slides of 64 faces. All were male, photographed in a full-face pose with a neutral expression. Faces with a beard, moustache and/or glasses were included. All clothing was masked.

Apparatus. The slides were projected onto a wall using a Kodak Carousel projector. The subject was given a remote advance button. The faces subtended a visual angle of approximately 10° (horizontally) by 13° (vertically).

Procedure. The procedure for rating the faces was that used by Valentine and Bruce (1986c). Subjects were asked to rate each face on a 1–7 scale. They were instructed to imagine that they had to meet each person at a

railway station and to rate each face for how easy it would be to spot in a crowd. A face that was very distinctive (or unusual) and so would be relatively easy to spot in a crowd should be rated 7, a typical face that would be difficult to identify in a crowd should be rated 1. Subjects were allowed to proceed through the list of slides at their own pace. Four different random orders of the slides were used. Based on the ratings, two sets of 16 distinctive faces (mean rating 4.86 and 4.87) were selected to serve as distinctive targets and distractors, respectively, in a recognition memory task. Similarly, two sets of 16 typical faces (mean rating 3.12 and 3.13) were selected to serve as typical targets and distractors, respectively. A *t*-test confirmed that the distinctive and typical faces differed significantly in their distinctiveness ($t(30) = 14.39, p < 0.005$).

Recognition Memory Task

Subjects. Twenty (5 male and 15 female) members of the APU subject panel acted as subjects; their mean age was 39.6 years.

Materials. The stimuli included two full-face photographs of each target face, one with a neutral expression (which had been used for the ratings) plus one smiling. Only one photograph of each distractor was required, but for half of the faces the photograph with a neutral expression used for the ratings was substituted for one with a smiling expression.

Apparatus. The slides were projected onto a wall using a Kodak Carousel projector. A paddle fitted to the projector was controlled automatically to give a set presentation and interval time. When a slide was projected, a timer was started automatically; it was stopped by the subject pressing a button on the response box. The subject's response and response latency were recorded manually. The approximate visual angle subtended by the faces was the same as that quoted above.

Design. There were two within-subjects factors: the distinctiveness of the faces and the orientation of the faces at test. The experiment consisted of two separate recognition memory tasks, one comprising typical faces and the other distinctive faces. In order to avoid the possibility of ceiling and floor effects due to better recognition memory for distinctive faces, a difference in exposure time was used to equate performance on upright typical and distinctive faces. The difference in exposure time used was based on pilot work. Two study lists of 16 slides each were constructed, one of typical and one of distinctive faces. In both lists half of the faces were smiling and half had a neutral expression. Two test lists were constructed, one of 32 typical and one of 32 distinctive faces, each consisting of 16 target and 16 distractor

items. Different pictures (involving a change of expression) of the target faces were shown in the test lists. Half of both the targets and distractors were smiling, half had a neutral expression. The targets and distractors in both test lists were divided into two sets of 8 faces, matched on distinctiveness. One set was presented upright, and one set was presented inverted.

Procedure. All subjects were tested individually. They were told that they would be shown two lists of slides of faces, and that there would be a recognition test immediately after each list. In the study lists typical faces were presented for 6 sec, and distinctive faces were presented for 3 sec. There was an interval of 2 sec between all slides. A different random order of the study list was constructed for every two subjects. At test all slides were presented for 5 sec, with a 2-sec interval. A quasi-random order was constructed for the test lists, with the constraint that no more than three target or distractor faces were presented in sequence, and no more than three consecutive faces were presented in the same orientation. A second order was made by reversing the order in which the halves of the list were presented. Slide order at test, the assignment of set of slides to orientation, and the order in which typical and distinctive faces were presented were counter-balanced across subjects. Subjects were warned that target faces would have a different facial expression at test. They were also warned that half of the faces would be presented upside-down and instructed to keep their heads upright when looking at these items. Subjects were instructed to respond as quickly and as accurately as possible.

Results

For each subject hit and false positive rates in each condition were calculated and combined in *A'* scores (Rae, 1976). *A'* scores were calculated from hit and false positives to either upright or inverted stimuli alone. The maximum number of hits or false positives was 8. Mean *A'* scores, mean number of hits and false positives, and the mean latency of hits and correct rejections are shown in Table 1.

An F_{\max} test was carried out on the raw data before conducting all of the analyses reported in this paper. If F_{\max} indicated that the assumption of homogeneity of variance was violated, an appropriate transformation was made and the transformed data subjected to another F_{\max} test. In all analyses reported in which a transformation has been carried out, the raw data violated the assumption of homogeneity of variance and the transformed data did not. When an analysis of untransformed data is reported, the raw data did not violate the assumption of homogeneity of variance.

A' scores were subjected to a $\sin^{-1}\sqrt{A'}$ transformation prior to being subjected to an analysis of variance (McNicol, 1972, p. 117). A within-

TABLE 1
Mean A' Scores, Hits,^a False Positives^a and Mean RT of Hits^b and Correct Rejections^b as a Function of Orientation and Distinctiveness (Experiment 1)

	Upright		Inverted	
	Distinctive	Typical	Distinctive	Typical
A'	0.905	0.919	0.802	0.699
Hits	6.50	6.55	5.35	4.85
F.P.	1.00	0.70	1.75	2.70
RT of hits	1356	1595	1770	2050
RT of C.R.	1418	1723	1767	2208

^a Maximum 8.

^b In msec.

subjects ANOVA gave a significant main effect of orientation, $F(1, 19) = 43.07$, $p < 0.001$. Upright faces were recognized more accurately than were inverted faces. The main effect of distinctiveness was not significant. The interaction between orientation and distinctiveness was significant, $F(1, 19) = 12.32$, $p < 0.0025$. Inversion caused a greater impairment to recognition of typical faces than of distinctive faces. Tukey HSD tests revealed that there was a significant effect of inversion on recognition of both distinctive and typical faces ($p < 0.05$), and that distinctive faces were recognized more accurately than were typical faces only when tested upside-down ($p < 0.05$).

A within-subjects ANOVA of hit rate data showed only a significant main effect of orientation, $F(1, 19) = 25.19$, $p < 0.001$. More hits were made to upright faces than to inverted faces. (Both other F ratios were less than 1.) The false positive data were subjected to a $\sin^{-1} \sqrt{(x + \frac{3}{8}) / (n + \frac{3}{4})}$ transformation as many of the data points were, or were close to, zero (Johnson & Leone, 1964). An ANOVA of the transformed false positive scores revealed a significant main effect of orientation, $F(1, 19) = 17.63$, $p < 0.001$. Fewer false positives were made to upright faces than to inverted faces. The main effect of distinctiveness was not significant [$F(1, 19) = 1.22$, $p = 0.28$]. There was a significant interaction between orientation and distinctiveness, $F(1, 19) = 8.73$, $p < 0.01$. Inverted presentation caused a greater rise in false positive rate to typical faces than to distinctive faces. Tukey HSD tests revealed a significant effect of orientation for typical faces ($p < 0.05$), but not for distinctive faces. The effect of distinctiveness was significant when stimuli were tested inverted ($p < 0.05$) but not when tested upright.

Mean reaction times of correct responses were also analysed. RTs of hits and correct rejections were subjected to separate within-subjects ANOVAs. The RT of hits were subjected to a log transformation (Winer, 1962) before an ANOVA was carried out. Analysis of transformed RT of hits revealed a

significant main effect of distinctiveness, $F(1, 19) = 7.86$, $p < 0.02$. Distinctive faces were recognized faster than typical faces. There was also a main effect of inversion, $F(1, 19) = 23.74$, $p < 0.001$. Upright faces were recognized faster than inverted faces. The interaction between these factors was not significant ($F < 1$). An ANOVA of the mean RTs to reject distractor faces correctly revealed a significant main effect of distinctiveness, $F(1, 19) = 18.68$, $p < 0.001$. Distinctive distractor faces were rejected more rapidly than typical distractors. There was a significant main effect of orientation, $F(1, 19) = 17.72$, $p < 0.001$, but the interaction between distinctiveness and orientation was not significant ($F = 1.33$).

Discussion

The use of different exposure times was successful in approximately equating performance in recognition of upright typical and distinctive faces. Of course, the lack of a main effect of distinctiveness under these conditions does not represent a failure to replicate the effect of distinctiveness on recognition memory. Indeed, an advantage for distinctive faces was found in the latency of both hits and correct rejections.

The accuracy data supported the predictions of the multidimensional space framework; an interaction between distinctiveness and orientation was found in the analysis of the A' and false positive data. As predicted, inversion was more disruptive to recognition of typical than of distinctive faces. Although the predicted interaction was not found in the analysis of hits, this pattern of results is consistent with the results of Light et al. (1979). They found a significant effect of distinctiveness in the analysis of d' and false positives, but not hits in several experiments.

No interaction between the effects of distinctiveness and orientation was found in the analyses of RT of hits and correct rejections. The lack of an interaction in the RT of hits is consistent with the analysis of accuracy in which no interaction was found in the hit rate. The RT of correct rejections showed a trend in the direction of the predicted interaction that was not significant, possibly due to the high variability of the RT data in a recognition memory paradigm.

In order to avoid the possibility that any interaction obtained in Experiment 1 may be explicable in terms of ceiling or floor effects, exposure time was manipulated to equate performance. However, this procedure raises the possibility that the difference in exposure time rather than the distinctiveness of the faces can account for the interaction with orientation. The aim of Experiment 2 was to eliminate this possible account of the Distinctiveness \times Inversion interaction by using a different method to attempt to equate performance in recognition of distinctive and typical upright faces.

EXPERIMENT 2

Experiment 2 was a replication of Experiment 1, except that in the initial lists the same exposure time was used for all stimuli. In order to equate performance, the number of faces included in the study list of distinctive faces was increased. The number of stimuli included in the test list was unchanged.

Method

Subjects. Twenty-eight (23 female and 5 male) members of the APU subject panel acted as subjects. Their mean age was 31 years.

Design and Procedure. All aspects of the design and procedure were the same as Experiment 1, except for the following details. All stimuli in both the study and test lists were presented for 5 sec, with a 2-sec interstimulus interval. The study list of distinctive faces consisted of 24 faces: 16 were the same faces as used in Experiment 1, and 8 new faces were included to make the list longer. The extra 8 faces were not included in the test list. The initial list of typical faces was unchanged, consisting of 16 faces all of which were included in the test list. Each of the four stimulus orders was used for seven subjects.

Results

The hit rate and false positive rate were calculated for each subject and combined in A' scores. Mean A' scores, mean number of hits and false positives, and mean latency of hits and correct rejections are shown in Table 2.

TABLE 2
Mean A' Scores, Hits,^a False Positives^a and Mean RT of Hits^b and Correct Rejections^b as a Function of Orientation and Distinctiveness (Experiment 2)

	Upright		Inverted	
	Distinctive	Typical	Distinctive	Typical
A'	0.915	0.894	0.824	0.693
Hits	6.75	6.39	5.75	4.82
F.P.	1.0	1.14	1.93	2.75
RT of hits	1242	1471	1626	1851
RT of C.R.	1382	1647	1830	1910

^a Maximum 8.

^b In msec.

Transformed A' scores ($\sin^{-1}\sqrt{A'}$) were subjected to a 2×2 within-subjects ANOVA. There was a significant main effect of orientation, $F(1, 27) = 127.49, p < 0.001$. Upright faces were recognized more accurately than inverted faces. There was also a significant main effect of distinctiveness, $F(1, 27) = 11.92, p < 0.005$. Distinctive faces were recognized more accurately than typical faces. The interaction between inversion and distinctiveness was significant, $F(1, 27) = 4.41, p < 0.05$. Inversion was more disruptive to recognition of typical faces than to recognition of distinctive faces. Tukey HSD tests did not reveal a significant effect of distinctiveness for either upright or inverted faces alone. The effect of inversion was significant for both distinctive and typical faces ($p < 0.05$).

Hit rate data were subjected to a 2×2 within-subjects ANOVA. There was a significant main effect of inversion, $F(1, 27) = 18.73, p < 0.001$. More hits were made to upright faces than to inverted faces. The main effect of distinctiveness was also significant, $F(1, 27) = 4.85, p < 0.05$. More hits were made to distinctive faces than to typical faces. The interaction was not significant [$F(1, 27) = 1.44, p > 0.2$].

A similar analysis of the false positive data was carried out. There were significant main effects of inversion, $F(1, 27) = 28.43, p < 0.001$, and distinctiveness, $F(1, 27) = 8.27, p < 0.01$. More false positives were made to inverted than to upright faces and to typical than distinctive faces. The interaction between inversion and distinctiveness just failed to reach the conventional level of statistical significance [$F(1, 27) = 3.28, p = 0.08$]. There was a trend for the increase in false positive rate caused by inversion to be greater for typical than for distinctive faces. Using Tukey HSD tests, the effect of inversion was significant for both distinctive and typical faces ($p < 0.05$). The effect of distinctiveness was significant for faces tested inverted ($p < 0.05$) but was not significant for faces tested upright.

An ANOVA on the latency of hits was carried out. There was a main effect of inversion, $F(1, 27) = 28.79, p < 0.001$. Upright faces were correctly recognized faster than inverted faces. There was a main effect of distinctiveness, $F(1, 27) = 16.52, p < 0.001$. Distinctive faces were correctly recognized faster than typical faces. There was no interaction between distinctiveness and inversion ($F < 1$).

The latency data for correct rejections were subjected to an ANOVA. There was a main effect of inversion, $F(1, 27) = 48.19, p < 0.001$. Upright faces were correctly rejected faster than inverted faces. There was a main effect of distinctiveness, $F(1, 27) = 13.75, p = 0.001$. Distinctive faces were correctly rejected faster than typical faces. The Inversion \times Distinctiveness interaction approached significance [$F(1, 27) = 3.61, p = 0.07$]. There was a trend for inversion to slow correct rejections of distinctive faces more than correct rejections of typical faces. Tukey HSD tests showed that the effect of

inversion was significant for upright and inverted faces ($p < 0.05$), and that the effect of distinctiveness was significant for upright faces ($p < 0.05$) but not for inverted faces.

Discussion

The use of a different number of faces in the study lists of distinctive and typical faces failed to equate recognition memory performance for distinctive and typical faces. An advantage for distinctive faces over typical faces was found in all five measures of performance. Although the intention was to control for the effect of distinctiveness, this result replicates the effect of distinctiveness on recognition memory for unfamiliar faces (e.g. Light et al., 1979) and demonstrates that the effect is found for the stimulus set used in Experiment 1. The critical prediction of a Distinctiveness \times Inversion interaction in the A' data was supported. Inversion was more disruptive to recognition of typical faces than of distinctive faces. This replicates the result of Experiment 1 in a task in which exposure duration was not manipulated, therefore the source of the interaction cannot be attributed to the use of differential exposure duration. The interaction between distinctiveness and inversion was not found in the analysis of hits, but it approached statistical significance in the analysis of false positives. This pattern is similar to Experiment 1 in which there was no interaction in the hits data but a significant interaction in the analysis of false positives. The interaction between distinctiveness and inversion appears weaker in Experiment 2, but this could be due to floor and ceiling effects reducing the interaction due to the failure to match performance on upright faces.

The analysis of the latency data showed essentially the same results as Experiment 1. Responses were slower to typical faces than to distinctive faces, and were slower to inverted faces than to upright faces. The Distinctiveness \times Inversion interaction approached significance in the analysis of RT of correct rejections, but this was in the opposite direction to the expected interaction. Inversion slowed responses to distinctive faces more than responses to typical faces. There is no clear theoretical interpretation of this trend, and there was no evidence of such a trend in Experiment 1.

EXPERIMENT 3

Experiments 1 and 2 have demonstrated that an interaction between the effects of inversion and distinctiveness was obtained in a task requiring recognition memory for previously unfamiliar faces. If the proposed multidimensional space framework is an effective model for the representation of information about faces in memory, the same interaction should be obtained in a task that requires recognition of previously familiar faces. This prediction was tested in Experiment 3 using a familiarity decision task. This task

has been used extensively in recent face recognition research. The subject is shown a series of faces, some of which are of famous (or personally familiar) faces, others are unfamiliar faces. The dependent variable is the subject's RT to decide whether each face is familiar or unfamiliar. A problem found in pilot work using a familiarity decision task to explore the effect of inversion was the range of performance found for upright and inverted faces. Accuracy for inverted faces was too low to allow RT to be reliably used as the dependent variable, but accuracy was effectively 100% for upright faces, so accuracy was also an unreliable measure. In order to overcome this problem, accuracy was enhanced by giving subjects a list of the names of the famous people who appeared in the experiment and using RT as the dependent variable. Reading a familiar person's name does not produce a priming effect on the subsequent RT in a familiarity decision task (Bruce & Valentine, 1985; Ellis, Young, Flude, & Hay, 1987).

An additional aim of Experiment 3 was to investigate the effect of distinctiveness upon the RT to reject unfamiliar faces. The multidimensional space framework predicts that subjects should be faster to reject distinctive unfamiliar faces than to reject typical unfamiliar faces. In a previous study this effect was not found in a familiarity decision task, but the distinctiveness of familiar and unfamiliar faces had not been matched (Valentine & Bruce, 1986c). Therefore, in Experiment 3 distinctiveness of familiar and unfamiliar faces was matched.

Method

The experiment was in two parts. First faces were rated for distinctiveness and familiarity. Different subjects then carried out a familiarity decision task.

Distinctiveness and Familiarity Ratings

Subjects. Eighteen (14 female and 4 male) members of the APU subject panel rated the stimulus faces for distinctiveness and familiarity. Their mean age was 42.4 years.

Materials. Pictures of 54 famous faces and 39 unfamiliar faces were rated. All of the pictures were of males. Pictures of the famous faces had been collected from a variety of sources (picture libraries, political parties, magazines, etc.). The poses varied from full-face to three-quarters profile. All pictures were copied onto monochrome slides through a circular mask, which excluded the majority of the background. The photographs of unfamiliar faces were either copied from a directory of actors or were studio portraits of academics. They were prepared in an identical format to the famous faces and included a similar range of poses. It was known from

previous work that all the actors were likely to be unfamiliar to the subjects. The academics lived in another part of the country and so were very unlikely to be familiar. Some pictures of faces with beards and glasses were included in the stimulus sets.

Apparatus. As for the distinctiveness ratings collected in Experiment 1.

Procedure. Subjects rated the unfamiliar faces first. The procedure for these ratings was identical to that described for Experiment 1. The subjects then rated the famous faces. When making distinctiveness ratings of famous faces, subjects were instructed to treat the faces as if they were unfamiliar—that is, to base their judgement entirely on the information available in the picture shown and to ignore any distinctive feature which they knew the celebrity had, but which could not be seen in the photograph (e.g. a distinctive hair colour). The subjects also rated the famous faces for familiarity on a 1–7 scale. In all other aspects the procedure was the same as that described above.

The stimulus sets to be used in the second phase of the experiment were selected on the basis of these ratings. These consisted of 16 typical unfamiliar faces (mean distinctiveness rating = 3.21) and 16 distinctive unfamiliar faces (mean distinctiveness rating = 5.05). A *t*-test for independent samples confirmed that there was a significant difference in distinctiveness between these two sets, $t(30) = 9.41, p < 0.005$. Sets of famous faces that matched the group mean ratings of unfamiliar faces as closely as possible were selected. The 16 typical famous faces had a mean distinctiveness rating of 3.33, and the 16 distinctive famous faces had a mean rating of 5.07. Again a *t*-test confirmed there was a significant difference in distinctiveness between these two groups, $t(30) = 9.70, p < 0.005$. The mean familiarity ratings were 5.45 for the typical famous faces and 5.90 for the distinctive famous faces. A *t*-test showed that this difference was not significant [$t(30) = 1.59, 0.1 < p < 0.05$]. All *t*-tests reported were two-tailed.

Familiarity Decision Task

Subjects. Twenty-four (22 female, 2 male) members of the APU subject panel acted as subjects. Their mean age was 38.9 years.

Materials. The stimuli consisted of the 32 slides of unfamiliar faces and the 32 slides of famous faces selected on the basis of the rating task described above.

Apparatus. The apparatus described for Experiment 1 was used.

Design. A list of 64 slides of faces was presented. Half of the slides were of famous people, half were of unfamiliar faces. The subject was instructed to respond YES if a face was familiar and NO if it was unfamiliar. There were 16 faces in each of the following four categories of stimuli: distinctive famous faces, typical famous faces, distinctive unfamiliar faces, and typical unfamiliar faces. Eight faces in each category were presented upright and 8 were presented upside-down. Each face was presented upright to half of the subjects and inverted to the remainder.

Procedure. All subjects were tested individually. Subjects were given a list of the names of the famous people who appeared in the experiment. They were asked whether they thought they would be able to recognize all of the people on the list. Two subjects who were not familiar with many of the faces did not take part in the experiment. All of the remaining subjects were familiar with all but two or three faces at most. If subjects said they would not recognize a particular face, the experimenter provided some other semantic information about the individual (e.g. a film they had appeared in, a description of a part played in a TV series, etc.). Often subjects would then be more confident of recognizing the face.

Subjects were told that they would see a series of slides of faces. They were informed that about half of the faces of celebrities, and that the names of all of the celebrities in the series were included in the list they had just read. Subjects were instructed to press the YES button if a face was familiar and to respond NO if it was unfamiliar. They were informed that RT was being measured and were instructed to respond as quickly and as accurately as possible. Subjects were warned that half of the faces would be presented upside-down and were instructed to keep their head upright while looking at upside-down faces. The slides were presented for 5 sec, with a 2-sec interval between slides. A quasi-random slide order was constructed with the constraint that there were no more than three famous or unfamiliar faces in sequence and no more than three consecutive faces presented in the same orientation. A second slide order was generated by reversing the order in which the two halves of the first series were presented. The slide order and the assignment of slides to the orientation were counter-balanced across subjects.

Results

For each subject, hit and false positive rates were calculated and combined in *A'* scores. Mean *A'* scores, mean number of hits and false positives, and mean RT of hits and correct rejections are shown in Table 3.

Assessment of accuracy in a familiarity decision task is always slightly problematic as it is possible that a subject does not know a "famous" face or genuinely recognizes an "unfamiliar" face. The former "disagreement" is

TABLE 3

Mean Latency and "Accuracy" of Responses in a Familiarity Decision Task as a Function of Orientation and Distinctiveness (Experiment 3)

	Upright		Inverted	
	Distinctive	Typical	Distinctive	Typical
Mean RT of correct familiar responses	983	1076	1552	2312
Mean RT of correct unfamiliar responses	1526	1698	2197	2372
Mean no. of correct familiar responses*	7.5	7.3	5.0	3.0
Mean no. of familiar responses to "unfamiliar" faces*	0.3	0.9	0.7	1.3
Mean A' of familiarity decisions	0.975	0.944	0.865	0.715

* Max no. of responses in each cell = 8.

much more likely, but the occurrence of such disagreements in this experiment should have been very much reduced by checking at the start that subjects were familiar with the names of the celebrities. Despite this precaution, "misses" occurred on 28.75% of trials in which a famous face was presented, and "false positives" occurred on 10% of trials in which an unfamiliar face was presented. In view of the high error rate, which reflects the extreme difficulty of recognizing upside-down faces, an analysis of the accuracy data is reported below.

Analysis of RT of "correct" responses does not raise such problems. The RT data of correct responses to famous faces were subjected to a log transform before an ANOVA was carried out. There was a main effect of distinctiveness, $F(1, 23) = 50.19$, $p < 0.001$. Distinctive faces were recognized more rapidly than typical faces. There was a main effect of orientation, $F(1, 23) = 149.41$, $p < 0.001$. Upright faces were recognized more rapidly than inverted faces. There was a significant interaction between these factors, $F(1, 23) = 18.07$, $p < 0.001$. Inversion increased the latency of recognition of typical faces more than of distinctive faces. Analysis of the simple main effects using Tukey HSD tests showed that the effect of inversion was significant for both distinctive and typical faces ($p < 0.05$) and that there was a significant effect of distinctiveness for faces tested inverted ($p < 0.05$) but not for faces tested upright. Due to the high error rate the mean number of correct RTs to typical famous faces seen inverted was only 3 (see Table 3). Therefore, a non-parametric test was used to support the results of the parametric analysis. The significant interaction was the predicted effect. In

order to test this interaction using a non-parametric test, the difference between RT to upright and inverted faces was calculated separately for typical and distinctive faces. A Wilcoxon test carried out on these differences in RT for typical and distinctive faces was significant, $T(24) = 17$, $p < 0.01$, indicating that the effect of inversion was greater for typical than for distinctive faces.

The RT data of correct rejections were subjected to a within-subjects ANOVA. There was a main effect of distinctiveness, $F(1, 23) = 11.34$, $p < 0.005$. Distinctive unfamiliar faces were rejected more quickly than typical unfamiliar faces. The main effect of orientation was significant, $F(1, 23) = 70.06$, $p < 0.001$. Upright unfamiliar faces were rejected faster than inverted unfamiliar faces. The interaction between distinctiveness and orientation was not significant.

In view of the high error rate in some conditions accuracy data were also analysed. A' was calculated by treating "familiar" responses to unfamiliar faces as false positives. Transformed A' scores ($\sin^{-1}\sqrt{A'}$) were subjected to a within-subjects ANOVA. There was a main effect of distinctiveness, $F(1, 23) = 51.48$, $p < 0.001$. Distinctive faces were recognized more accurately than typical faces. There was also a main effect of orientation, $F(1, 23) = 160.92$, $p < 0.001$. Upright faces were recognized more accurately than inverted faces. The predicted interaction between these factors was statistically significant, $F(1, 23) = 5.16$, $p < 0.05$. Inversion caused a greater impairment to recognition of typical faces than of distinctive faces. Tukey HSD tests revealed a significant effect for all four pairwise comparisons of simple main effects ($p < 0.05$).

The proportion of correct "familiar" responses to famous faces (hit rate) were also subjected to a within-subjects ANOVA. There was a main effect of distinctiveness, $F(1, 23) = 43.29$, $p < 0.001$, a main effect of orientation, $F(1, 23) = 149.41$, $p < 0.001$, and a significant interaction between these factors, $F(1, 23) = 14.60$, $p < 0.01$. Inversion caused a greater impairment to recognition of typical faces than of distinctive faces. Tukey HSD tests showed a significant effect of inversion for both distinctive and typical faces ($p < 0.05$), and a significant effect of distinctiveness for faces tested inverted ($p < 0.05$) but not for faces tested upright.

There were many empty cells in the "false positive" data. These data were found to violate the assumption of homogeneity of variance even after an appropriate transformation had been carried out. Therefore, a parametric test would not be valid. A Wilcoxon test on the difference in error rate made to upright and inverted faces computed separately for distinctive and typical faces, as described above, was carried out to test the interaction between distinctiveness and inversion. No significant difference was found [$T(16) = 62.5$, n.s.].

Discussion

The RT data supported the predictions derived from the multidimensional space framework. Familiar distinctive faces were accepted faster than familiar typical faces. Thus the effect of distinctiveness in a familiarity decision task, found in previous work, was replicated. It was also predicted that an effect of distinctiveness should be found on the RT to reject unfamiliar faces. This prediction was supported: unfamiliar distinctive faces were rejected faster than unfamiliar typical faces.

The predicted Inversion \times Distinctiveness interaction was found in the RTs to familiar faces. RT to accept familiar typical faces showed a greater increase due to inversion than did the RT to accept familiar distinctive faces. There was no Distinctiveness \times Inversion interaction in the RT to reject unfamiliar faces. It is not clear why the interaction between distinctiveness and inversion should not be found in rejection latencies, but it is possible that when, in the context of looking for highly familiar faces, a match to an inverted face is not found, some form of a checking process may be initiated which obscures the Distinctiveness \times Inversion interaction. The RT data were supported by the accuracy data. The predicted main effect of distinctiveness and the Distinctiveness \times Inversion interaction were found in the analysis of the A' scores and correct "familiar" responses. This contrasts with the results of Experiments 1 and 2 in which the interaction was found in the A' and false positive rate data, but not in the hit rate data. The difference in the task demands between the two experiments may be a cause for this difference in the results. Experiments 1 and 2 required recognition of previously unfamiliar faces seen once a few minutes earlier, whereas as Experiment 3 required recognition of highly familiar faces.

EXPERIMENT 4

Experiments 1-3 have demonstrated an interaction between distinctiveness and inversion in two tasks that require *recognition* of individual faces. In Experiment 4 the interaction between these factors in a face *classification* task was investigated. As described above, the effect of distinctiveness reverses in a classification task. Typical faces can be classified as a face more rapidly than distinctive faces. In terms of the norm-based coding model, it is assumed that the length of the derived vector will determine the RT in a face classification task. Typical faces are closer to the norm and so can be classified faster than distinctive faces. In the exemplar-based model the exemplar density will determine the RT of classification. High exemplar density will lead to faster classification decisions than low exemplar density. As typical faces will be located in regions of higher exemplar density than distinctive faces, the exemplar-based model also predicts that typical faces will be classified faster than distinctive faces. If the stimuli in this task are

presented upside-down, the effect would be represented in the multidimensional space by increased random error associated with each encoded face. In terms of the norm-based coding model, an increased random error is as likely to move the vector encoded from the stimulus closer to the norm, as it is to move it further away. In terms of the exemplar-based model, an increased random error is as likely to move the encoded location of the stimulus to an area of greater exemplar density, as it is to move it to an area of lower exemplar density. Therefore, in both models the effects of the increased variability of encoding should tend to cancel each other out. Thus, it is predicted that the effect of distinctiveness should be found for both upright and inverted faces, with no interaction between distinctiveness and inversion.

The prediction of a Distinctiveness \times Inversion interaction in a *recognition* task arises because inversion causes a greater mismatch between the encoding of a face and its representation in memory. The effect of this greater mismatch will depend on the exemplar density in the region. No interaction between distinctiveness and inversion is predicted in a face *classification* task because there is *no* requirement to discriminate between the representation of individual faces in memory.

Method

The experiment was in two parts. First faces were rated for distinctiveness. Different subjects then carried out a face classification task.

Distinctive Ratings

Subjects. Fifteen (2 male and 13 female) of the sixteen subjects who rated the faces used in Experiment 1 also rated the faces used in this experiment. Their mean age was 44.9 years.

Materials. Thirty-two faces were prepared as monochrome prints approximately 130 mm \times 110 mm. The faces were all male, in a full-face pose with a neutral expression. None had a beard or a moustache and none wore glasses. They were copied from a directory of actors, and on the basis of previous work were not likely to be familiar to subjects.

Procedure. The procedure was the same as that described above, except that the prints rather than slides were used. Based on the ratings, the faces were split into a set of 16 distinctive faces (mean rating, 4.9) and 16 typical faces (mean rating, 3.3). A *t*-test confirmed that there was a significant difference in perceived distinctiveness between the two sets of faces, $t(30) = 7.88, p < 0.005$.

Face Classification Task

Subjects. Twenty-one members of the APU subject panel acted as subjects. Data from one subject were discarded because of an error rate of 10%. Thus usable data were obtained from 20 (4 male and 16 female). Their mean age was 41.4 years.

Materials. Stimuli were prepared from the 32 faces described above, together with 4 faces for use as practice items. An "intact" and a "jumbled" version of each face was made. A grid of lines was drawn on the intact faces before being copied onto slides, so that the intact faces had the same lines on them as the jumbled faces. To make the jumbled faces, a rectangle was drawn around the eyes-nose-mouth region of each face and was divided by horizontal lines into three regions corresponding to each of these features. The features were cut out and reassembled either in the order mouth-eyes-nose or nose-mouth-eyes (from top to bottom). Figure 3 shows an example of an intact and a jumble face.

Apparatus. The apparatus was the same as that used in Experiment 1.

Design. The intact and jumbled versions of the extra four faces comprised eight practice items at the beginning of the list. The experimental trials




FIG. 3. An example of an intact and a jumbled face used in the face classification task in Experiment 4.

consisted of 16 distinctive intact faces, 16 distinctive jumbled faces, 16 typical intact faces, and 16 typical jumbled faces. The faces were split into two sets, each consisting of half of each of these categories of stimuli. The mean distinctiveness ratings of faces in the two sets of distinctive and typical faces were matched. Set A was presented upright, and Set B was presented inverted to half of the subjects. The remaining half saw the stimuli presented in the opposite orientation. The subject was asked to classify each stimulus as it was presented as either a "face" or a "jumbled face". The design had two within-subjects factors, the distinctiveness and orientation of the face. The dependent variable was the RT of correct responses.

Procedure. Subjects were tested individually. They were shown an example print of an intact and a jumbled face, and the nature of the task was explained. They were informed that half of the faces would be presented upside-down and were instructed simply to classify each stimulus as a face or a jumbled face regardless of its orientation. The response box had one button labelled "face" and one labelled "jumbled face". Subjects were informed that response latency was being measured and instructed to respond as quickly and as accurately as possible. The stimuli were presented in a quasi-random order, with the following constraints: (a) no more than three consecutive intact or jumbled faces, (b) no more than three typical or three distinctive faces occurred in sequence, and (c) no more than three consecutive stimuli were presented in the same orientation. Two orders were used, one being generated by reversing the order of presentation of the two halves of the first list. The slide order and the assignment of slide set to orientation were counter-balanced across subjects. Slides were shown for 2.5 sec, with a 2.5-sec interval between slides. Response latencies greater than 2 sec were scored as errors.

Results

Errors were made on 2.3% of trials in which an intact face was presented. There were too few errors to analyse. Mean RTs of correct responses to the intact faces are shown in Table 4. These data were subjected to a 2×2 ANOVA with distinctiveness and orientation as within-subjects factors. There was a main effect of orientation, $F(1, 19) = 47.47, p < 0.001$. Upright faces were classified faster than inverted faces. There was also a main effect of distinctiveness, $F(1, 19) = 9.44, p < 0.01$. Typical faces were classified faster than distinctive faces. The interaction between orientation and distinctiveness was not significant ($F < 1$).

Errors were made on 3.3% of trials in which a jumbled face was presented. There were too few errors to analyse. Mean RTs of correct responses to jumbled faces are shown in Table 4. These data were subjected to a 2×2

Face Classification Task

Subjects. Twenty-one members of the APU subject panel acted as subjects. Data from one subject were discarded because of an error rate of 10%. Thus usable data were obtained from 20 (4 male and 16 female). Their mean age was 41.4 years.

Materials. Stimuli were prepared from the 32 faces described above, together with 4 faces for use as practice items. An "intact" and a "jumbled" version of each face was made. A grid of lines was drawn on the intact faces before being copied onto slides, so that the intact faces had the same lines on them as the jumbled faces. To make the jumbled faces, a rectangle was drawn around the eyes-nose-mouth region of each face and was divided by horizontal lines into three regions corresponding to each of these features. The features were cut out and reassembled either in the order mouth-eyes-nose or nose-mouth-eyes (from top to bottom). Figure 3 shows an example of an intact and a jumble face.

Apparatus. The apparatus was the same as that used in Experiment 1.

Design. The intact and jumbled versions of the extra four faces comprised eight practice items at the beginning of the list. The experimental trials




FIG. 3. An example of an intact and a jumbled face used in the face classification task in Experiment 4.

consisted of 16 distinctive intact faces, 16 distinctive jumbled faces, 16 typical intact faces, and 16 typical jumbled faces. The faces were split into two sets, each consisting of half of each of these categories of stimuli. The mean distinctiveness ratings of faces in the two sets of distinctive and typical faces were matched. Set A was presented upright, and Set B was presented inverted to half of the subjects. The remaining half saw the stimuli presented in the opposite orientation. The subject was asked to classify each stimulus as it was presented as either a "face" or a "jumbled face". The design had two within-subjects factors, the distinctiveness and orientation of the face. The dependent variable was the RT of correct responses.

Procedure. Subjects were tested individually. They were shown an example print of an intact and a jumbled face, and the nature of the task was explained. They were informed that half of the faces would be presented upside-down and were instructed simply to classify each stimulus as a face or a jumbled face regardless of its orientation. The response box had one button labelled "face" and one labelled "jumbled face". Subjects were informed that response latency was being measured and instructed to respond as quickly and as accurately as possible. The stimuli were presented in a quasi-random order, with the following constraints: (a) no more than three consecutive intact or jumbled faces, (b) no more than three typical or three distinctive faces occurred in sequence, and (c) no more than three consecutive stimuli were presented in the same orientation. Two orders were used, one being generated by reversing the order of presentation of the two halves of the first list. The slide order and the assignment of slide set to orientation were counter-balanced across subjects. Slides were shown for 2.5 sec, with a 2.5-sec interval between slides. Response latencies greater than 2 sec were scored as errors.

Results

Errors were made on 2.3% of trials in which an intact face was presented. There were too few errors to analyse. Mean RTs of correct responses to the intact faces are shown in Table 4. These data were subjected to a 2×2 ANOVA with distinctiveness and orientation as within-subjects factors. There was a main effect of orientation, $F(1, 19) = 47.47, p < 0.001$. Upright faces were classified faster than inverted faces. There was also a main effect of distinctiveness, $F(1, 19) = 9.44, p < 0.01$. Typical faces were classified faster than distinctive faces. The interaction between orientation and distinctiveness was not significant ($F < 1$).

Errors were made on 3.3% of trials in which a jumbled face was presented. There were too few errors to analyse. Mean RTs of correct responses to jumbled faces are shown in Table 4. These data were subjected to a 2×2

TABLE 4

Mean RTs to Classify Correctly Intact and Jumbled Faces, as a Function of Distinctiveness and Orientation (Experiment 4)

	Intact faces			Jumbled Faces		
	Typical	Distinctive	Mean	Typical	Distinctive	Mean
Upright	683 (0.6)	737 (1.9)	710 (1.3)	741 (0)	738 (0.6)	739 (0.3)
Inverted	817 (1.3)	860 (5.6)	839 (3.5)	866 (7.5)	841 (5.0)	855 (6.3)
Mean	750 (0.9)	798 (3.7)		803 (3.7)	791 (2.8)	

Note: * % errors are shown in parentheses

within-subjects ANOVA. There was a significant effect of orientation, $F(1, 19) = 14.23, p < 0.002$. Inverted jumbled faces were classified more slowly than upright jumbled faces. No other effects were significant (both $F_s < 1$).

Discussion

The results of this experiment supported the predictions derived from the multidimensional space framework. The finding that typical faces can be classified as faces faster than distinctive faces has replicated the effect reported by Valentine and Bruce (1986c).

As predicted by the multidimensional space framework, an effect of distinctiveness was also found in classification of inverted faces. The orientation of the stimulus did not affect the magnitude of the effect of distinctiveness.

For completeness, data on the classification of jumbled faces were reported. However, the multidimensional space framework does not make any clear prediction of the effects to be expected in this case. It is not clear that a distinctive face when jumbled will become a distinctive jumbled face. In fact no effect of distinctiveness was found.

The Effect of Race

So far the multidimensional space framework has been applied to accounts of the effects of distinctiveness and inversion. The framework can also provide a parsimonious account of the effect of race in face recognition. Assuming that the dimensions of the space are based on experience with faces of predominately one race, the feature dimensions underlying the multidimensional space will be those that are appropriate for discriminating one particular race of faces. However, different dimensions will be optimum for

discriminating another race of faces. For example, there is evidence that different facial features are used to describe black and white faces (Ellis, Deregowski, & Shepherd, 1975), and to judge similarity of a set of simultaneously presented black or white faces (Shepherd & Deregowski, 1981). Although subjects are capable of using appropriate cues to distinguish simultaneously presented other-race faces, the evidence available suggests that they resort to using the familiar but inappropriate own-race facial cues in a recognition memory paradigm (Shepherd, 1981). Therefore, other-race faces will be encoded on dimensions that do not serve to discriminate well amongst faces of that race. The dimension values that are salient will tend to be those that are characteristic of the other race, rather than those that are characteristic of the individual. Therefore, other-race faces will be distant from the central tendency but have many feature values in common and so will be densely clustered. Thus other-race faces form an exception to the assumption that the exemplar density decreases with increasing distance from the central tendency.

An illustration of a two-dimensional version of the norm-based coding model including own and other-race faces is shown in Figure 4.

A number of predictions can be derived from this interpretation of the race effect. The norm-based coding model will be considered first. Increased distance from the norm will result in the vectors of faces separated by a given distance being more similar to each other and therefore more difficult to discriminate (see Figure 2). Assuming that other-race faces will be more distant from the norm than own-race faces, the norm-based coding model predicts that other-race faces will be more difficult to recognize. This prediction holds even if other-race faces are equally densely clustered as own-race faces. The other-race effect will be enhanced if other-race faces are indeed more densely clustered because the dimensions defining the space are inappropriate. The difficulty of discrimination between similar vectors, due to the distance from the norm and/or higher density of other-race faces, means that recognition of other-race faces should be severely disrupted by an increase in the error of deriving a stimulus vector. By the same argument used above for typical and distinctive faces, this leads to the prediction that recognition of other-race faces will be more disrupted by inversion than recognition of own-race faces (i.e. in a recognition task own-race faces are analogous to distinctive faces and other-race faces are analogous to typical faces.)

The exemplar-based model assumes that similarity is determined by the distance separating faces and so is independent of the location of the norm. Therefore, the exemplar-based model can only account for better recognition of own-race faces than of other-race faces if it is assumed that exemplar density is greater for other-race faces than for own-race faces. The work reviewed above suggests that there is some evidence to support this view. As




FIG. 4. A two-dimensional representation of a norm-based coding model of face recognition, showing populations of own and other-race faces. As in Figure 1, faces are represented by points in an n -dimensional space. Two dimensions have been plotted for the purpose of illustration only. The axes are unlabelled but could represent any dimensions that serve to discriminate between faces and differ in the central tendency of own and other-race faces.

the exemplar-based model uses differences in exemplar density to account for the other-race effect, it also leads to the prediction that recognition of other-race faces will be more disrupted by inversion than recognition of own-race faces. This result has been reported by Valentine and Bruce (1986a). The results of this study are given in Table 5.

The design of this experiment was identical to Experiment 1, except that lists of white faces and black African faces were used instead of lists of typical and distinctive white faces. The only differences in the procedure were that an identical picture of the target faces was used at test and exposure times of 2 sec (white faces) and 5 sec (black faces) in the initial list were used to equate performance. Twenty white subjects took part in the experiment. Separate

TABLE 5
Mean A' Scores, Hits,^a False Positives^a and Mean RT of Hits^b and Correct Rejections^b as a Function of Orientation and Race in a Recognition Memory Task

	Upright		Inverted	
	White	Black	White	Black
A'	0.900	0.918	0.723	0.609
Hits	6.5	6.7	4.9	4.7
F.P.	1.2	0.9	2.3	3.6
RT of Hits	1222	1553	1685	1933
RT of C.R.	1502	1792	2002	2355

^a Maximum 8.

^b In msec.

Note: Accuracy data from: Valentine & Bruce (1986a).

analysis of the hit and false positive data, and analysis of the latency data were not reported by Valentine and Bruce (1986a), so they will be reported here briefly. The results of both the accuracy and latency measures were directly analogous to the results of Experiment 1 in all aspects. There was a significant interaction between the effects of race and orientation in the analyses of transformed A' scores ($\sin^{-1}\sqrt{A'}$), $F(1, 16) = 16.52$, $p < 0.001$, and false positive rate $F(1, 16) = 21.41$, $p < 0.001$, but not in the analyses of hit rate or latency of hits or correct rejections (all F ratios < 1). In the analyses in which an interaction was found, it reflected a greater effect of orientation upon recognition of other-race faces than of own-race faces. In all five analyses there was a strong main effect of orientation. In the accuracy data no effect of race was found, due to the use of differential exposure duration to equate performance. However, an own-race advantage was found in the latency of hits, $F(1, 16) = 6.09$, $p < 0.05$, and in the latency of correct rejections, $F(1, 16) = 14.40$, $p < 0.001$.

The effect of inversion is modelled in the multidimensional space framework as an increase in the error of encoding the stimulus. This assumption does not involve any aspect that is unique to the effect of inversion. In fact, it leads to the prediction that any manipulation that impairs face recognition will impair recognition of other-race faces more than recognition of own-race faces. Ellis and Derogowski (1981) found that a change in pose (full-face vs. three-quarters profile) between inspection and test in a recognition memory task impaired recognition of other-race faces more than recognition of own-race faces. Therefore, Ellis and Derogowski's data support the predictions of the multidimensional space framework. This study included a cross-cultural control in which European and African subjects were tested using black and white faces.

It is interesting to note that the multidimensional space framework and the distinction between component and configural information (see introduction) lead to different predictions of the interaction between inversion and race. Rhodes, Tan, Brake, and Taylor (1989) base their analysis of the effect of race on Diamond and Carey's (1986) work on the role of expertise in encoding second-order relational information. Rhodes and colleagues argue that lack of expertise in encoding other-race faces would mean that second-order relational information cannot be encoded from other-race faces. As it is assumed that disruption to encoding of second-order relational information in own-race faces causes the disproportionate effect of inversion on own-race face recognition, Rhodes and colleagues argue that recognition of other-race faces will be *less* disrupted by inversion than recognition of own-race faces. They found support for their hypothesis using a somewhat different procedure to that used by Valentine and Bruce (1986a). These apparently conflicting results will be discussed further below. In Experiment 5, the exploration of the interaction between the effects of distinctiveness and inversion is extended to a face classification task.

EXPERIMENT 5

The norm-based coding model makes a clear prediction that other-race faces should be classified as a face more slowly than own-race faces in an intact/jumbled face classification task. The length of the vector is assumed to affect RT in a face classification task. As other-race faces will, as a group, be located further from the norm than own-race faces, other-race faces will be classified more slowly than own-race faces. The exemplar-based model does not make a clear prediction. Classification RT is assumed to be determined by the exemplar density in a region of a given radius around the stimulus (Krumhansl, 1978). Higher exemplar density will lead to faster classification RTs. In order to account for the effect of race on recognition memory, the exemplar-based model must assume that other-race faces are more densely clustered than own-race faces. However, there will be fewer other-race faces in memory than own-race faces. Therefore, there will be two opposing factors that determine RT in the exemplar-based model. The effect of race on classification RT will depend on the value of unspecified parameters such as the radius of the region in question, the density of own and other-race faces, and the number of faces in memory. The exemplar-based model could account for faster classification of own-race faces if the exemplar density around a stimulus is greater for own-race faces, despite being less densely clustered, because there are many more own-race faces falling within the region concerned. The norm-based coding model and the exemplar-based model both predict that the effect of race on classification of faces should not be affected by inversion of the stimulus. An increase in the random encoding

error caused by stimulus inversion would not have a systematic effect on either the distance from the norm or the exemplar density around the location of the stimulus. The effect of race and inversion in a classification task were investigated in Experiment 5.

Method

Subjects. Sixteen (10 female and 6 male) members of the APU subject panel acted as subjects. Their mean age was 40.5 years. All were Caucasian.

Materials. Sets of 18 white faces and 18 black African faces were selected. None had a beard or wore glasses. The 18 white faces were a subset of those used in Experiment 4. For each face a slide of an "intact face" and a "jumbled face" were prepared as described above.

Apparatus. This was the same as for Experiment 1.

Design. Two of the 18 faces of each race were assigned to be practice items. The intact and jumbled versions of these faces comprised eight practice items at the beginning of the list. The experimental trials consisted of 16 black intact faces, 16 black jumbled faces, 16 white intact faces, and 16 white jumbled faces. As in Experiment 4, the stimuli were split into sets, and the assignment of stimulus set to orientation was counter-balanced across subjects. The design had two within-subjects factors, the race and orientation of the face. The dependent variable was the RT to classify intact faces correctly.

Procedure. This was as described for Experiment 4, except that subjects were informed that they would see black and white faces.

Results

Mean RTs to classify the intact faces correctly were calculated for each subject. These data and the error rates are shown in Table 6. As in Experiment 4, responses over 2 sec were scored as errors. Errors were made on 2.9% of trials.

The RT data were subjected to an ANOVA with race and orientation of the face as within-subject factors. This analysis gave a main effect of race, $F(1, 15) = 26.13, p < 0.001$, and orientation, $F(1, 15) = 94.66, p < 0.001$. Thus own-race (white) faces were classified as a face faster than other-race (black) faces and upright faces were classified faster than inverted faces. There was no interaction between race and orientation ($F = 1.04$).

Although no predictions were made concerning the effect of race and orientation on the RT to classify jumbled faces, an analysis of these data was

TABLE 6
Mean RTs to Classify Intact and Jumbled Faces, as a Function of Race and Orientation
(Experiment 5)

	Intact Faces			Jumbled Faces		
	White	Black	Mean	White	Black	Mean
Upright	611 (0.8)	765 (3.1)	688 (1.9)	739 (4.7)	707 (0)	723 (2.3)
Inverted	719 (1.6)	843 (6.3)	781 (3.9)	796 (7.0)	724 (0)	759 (3.5)
Mean	665 (1.2)	804 (4.7)		768 (5.9)	716 (0)	

Note: % errors are shown in parentheses

carried out. Errors were made on 2.9% of trials. The mean RT of correct responses are shown in Table 6.

A 2×2 ANOVA revealed a main effect of race, $F(1, 15) = 11.17, p < 0.005$. Other-race (black) faces were classified as jumbled faces faster than own-race (white) faces. Neither the main effect of inversion nor the interaction between race and inversion achieved statistical significance; however, the interaction term approached statistical significance, $F(1, 15) = 3.86, p < 0.07$. There was a trend for classification of own-race faces to show a larger effect of inversion.

Discussion

This experiment provides evidence that in a face classification task, other-race faces take longer to be classified as a face than own-race faces. As predicted, there was no interaction between the effects of race and inversion on correct classification of intact faces. These results are consistent with the multidimensional space framework and are analogous to the results of Experiment 4. According to the norm-based coding model, other-race faces take longer to classify as a face because they are relatively distant from the norm. Therefore, in this task other-race faces behave like distinctive faces. This is in contrast to a recognition task when other-race faces behave like typical faces showing a larger effect of inversion, because in a recognition task similarity between faces rather than distance from the own-race norm is the critical parameter. The exemplar-based model can also account for the effect of race in classification if it is assumed that RT is determined by exemplar density in a region around the stimulus and that the relatively small

number of other-race faces in memory makes the exemplar density of other-race faces relatively low.

One limitation of Experiment 5 and Valentine and Bruce's (1986a) experiment is that only subjects from one race took part. The interaction found by Valentine and Bruce was interpreted as evidence that recognition of other-race faces is more impaired by inversion than recognition of own-race faces. However, the possibility cannot be excluded that the black faces used in the experiment would show a greater effect of inversion even to black subjects. A cross-cultural study would be needed to eliminate this possibility. However, it should be noted that Ellis and Deregowski (1981) found that a change in pose affected recognition of other-race faces more than recognition of own-race faces using a cross-cultural design. Similarly a cross-cultural study would be needed to confirm that the effect of race on classification found in Experiment 5 was genuinely due to an other-race effect.

The results of Experiment 5 and Valentine and Bruce (1986a) provide at least tentative evidence that the multidimensional framework can be applied to the effect of race in face processing. Two studies of *recognition* tasks provide a note of caution, however Buckhout and Regan (1988) examined the interaction between the effect of race and inversion in recognition memory for unfamiliar faces. Both black and white subjects took part (presumably residents of New York). Buckhout and Regan found an own-race advantage but no Race \times Orientation interaction. They do not report separate analysis of false positives and hits but use a percent correct measure. It is possible that the use of this measure may obscure an effect on false positives alone as found by Valentine and Bruce (1986a).

Rhodes et al. (1989) studied the effect of race and inversion in a cross-cultural study using Chinese and European faces. They used a two-alternative forced choice test procedure and found a greater inversion effect for own-race faces than other-race faces (i.e. the opposite interaction to that found by Valentine & Bruce, 1986a). In their first experiment the significant Race \times Orientation interaction was found in forced choice RT but not accuracy. In a second experiment the interaction was found only in forced choice accuracy and only when the test face pairs were not highly similar. The use of a forced choice test procedure in Rhodes and colleagues' study may have obscured the effect in false positives found by Valentine and Bruce, but this does not explain why the opposite interaction should be found. A number of parameters may account for the different results of the three studies that have investigated the interaction of race and inversion, e.g. the difficulty of tasks, the test procedures and measures used, the subjects' experience of other-race faces, homogeneity of stimulus sets, etc. Further research is needed to resolve the issues raised by these studies.

GENERAL DISCUSSION

An account of structural encoding of faces has been outlined in which the mental representation of faces is considered in terms of locations in a multidimensional space. Within this general framework, two specific models were identified, a norm-based coding model and a purely exemplar-based model. Although the detailed assumptions of the models differ, both models predict a greater effect of inversion for typical faces than for distinctive faces in recognition tasks but not in a classification task. These predictions were supported. It was found that inversion was more disruptive to recognition of typical faces than to recognition of distinctive faces in a recognition memory task using previously unfamiliar faces (Experiments 1 and 2), and in a familiarity decision task using famous faces (Experiment 3). In Experiment 4, it was found that the effects of inversion and distinctiveness were additive in a face classification task. These results are consistent with inversion causing an increased error of encoding that only has a differential effect on distinctive and typical faces when the tasks require individual faces to be discriminated.

The multidimensional space framework was also extended to account for the effects of race. It was proposed that other-race faces form a class of faces that violate the statistics of the own-race face population. Other-race faces have been found to behave like typical faces in a recognition task (Valentine & Bruce, 1986a) and like distinctive faces in a face classification task (Experiment 5). The extension of the multidimensional space framework to account for the effects of race is, at this stage, somewhat tentative as data from studies that include full cross-cultural controls are conflicting. However, the multidimensional space framework potentially provides a unified account of the effects of distinctiveness, inversion, and race. A further advantage is that it is applicable to recognition of both familiar (e.g. famous) faces and previously unfamiliar faces.

The multidimensional framework (in either its norm-based coding or purely exemplar-based forms) provides a principled framework that can account for the effects found in the experiments reported here. Can some other theory account for these results rendering the multidimensional space framework unnecessary?

It is possible that the interaction between distinctiveness and orientation found in Experiments 1–3 could be explained in terms of disruption to processing of configural information in inverted faces (Carey & Diamond, 1977; Diamond & Carey, 1986). Recognition of distinctive faces may be less affected by inversion because distinctive faces can be more easily encoded in terms of component information (i.e. information based on isolated features rather than global properties). This interpretation of the data is not mutually exclusive to the multidimensional space framework, but it begs the question of the basis of distinctive information in faces. The multidimensional space

framework is neutral in respect to the issue of whether distinctive information is conveyed predominately by component or configural information. It should be possible to separate an account based solely on disruption to processing of configural information from the more general distinctiveness account offered by the multidimensional space framework by exploring the interaction between distinctiveness and transformations other than inversion (e.g. masking with visual noise). The multidimensional framework predicts that any experimental manipulation that impairs encoding of faces should interact with distinctiveness in the same way as inversion. The account in terms of disruption to processing of configural information is specific to the effect of inversion; different accounts would be required for each transformation found to interact with distinctiveness. Further research is required to explore the effects of other transformations on recognition of distinctive and typical faces.

The predictions of the multidimensional space framework and the component/configural information differ in regard to the interaction between the effects of race and orientation on recognition memory for faces. Based on Diamond and Carey's (1986) distinction between isolated features and second-order relational information, Rhodes et al. (1989) predict that recognition of other-race faces will be less disrupted by inversion than recognition of own-race faces. Although the component/configural information distinction and the multidimensional space framework make different predictions concerning the Race \times Orientation interaction in recognition memory, the available data are conflicting (see discussion of Experiment 5).

In order to provide a full account for the present data in terms of a component/configural information distinction, it is also necessary to explain why distinctiveness affects a face classification task and to account for the additivity of the effects of distinctiveness and inversion in such a task. It is plausible to argue that second-order relational information is irrelevant to the task because only discrimination of first-order relational properties is required (i.e. Is the configuration face-like?). This argument could account for the lack of an interaction between distinctiveness and inversion. However, it is then not clear why the presence of the more distinctive isolated features of a distinctive face should affect the classification decision at all if their arrangement conforms to the first order relational properties of a face.

A third possible account would be to argue that the data reported reflect response bias. The data in recognition tasks could arise if subjects are more likely to respond positively: (1) to typical faces than to distinctive faces; (2) to inverted faces than to upright faces; and (3) to other-race than to own-race faces. However, a simple response bias model cannot account for the data in recognition tasks (Experiments 1–3) because in all experiments the critical orientation by distinctiveness interaction was found in A' , a

criterion-free measure of sensitivity. Therefore, I suggest that the multidimensional space framework provides a more parsimonious account of the data reported than the other possible accounts considered here.

So far the norm-based coding model has been described as if an explicit process of comparing each stimulus face to a single stored norm was involved. Rhodes et al. (1987) and Ellis (1981) proposed that multiple norms might exist for different categories of faces. For example, separate norms might be stored for male and female faces and for own-race and other-race faces. The norm-based coding model outlined here assumes that other-race faces are encoded by reference to the own-race norm, only because it is assumed that other-race faces have rarely been experienced. With increasing familiarity of faces of another race, there must be a point at which a new norm is abstracted and stored. Storage of multiple norms raises the problem of how the appropriate norm is selected to encode a face. If the nearest norm to a stimulus in the multidimensional space is selected, a distance-based similarity measure, as proposed in the purely exemplar-based model, is used to select the norm. If such a process is postulated, it is unparsimonious to suggest that a different basis of similarity is used to recognize individual faces. Therefore, the exemplar-based model should be preferred on the grounds of parsimony in the absence of any evidence compelling the norm-based coding model to be preferred.

The experiments reported here do not discriminate between the norm-based coding model and the exemplar-based model. This begs the question of how, if at all, it could be possible to distinguish between the two models. The critical difference between the models is that distance from the norm and exemplar-density are both important parameters in the norm-based coding model, but exemplar-density alone determines performance in the exemplar-based model. The difficulty in discriminating between the models arises because exemplar density is assumed to be correlated to distance from the norm. However, other-race faces provide a violation of this assumption. Therefore, exploration of the interaction between the effects of distinctiveness and race provides a means to distinguish the models. A cross-cultural study of these factors is currently underway (Valentine & Endo, 1990).

An alternative to the dichotomy between the norm-based coding and the exemplar-based model is to view "prototype abstraction" as an emergent property of distributed storage of faces. The description of faces encoded as vectors in a multidimensional space and the use of a vector similarity measure is strikingly similar to some parallel distributed processing models. Therefore, a distributed model of memory provides an alternative conceptual framework which allows the vector-based similarity measure of the norm-based coding model to be retained without raising the conceptual problems caused by storage of multiple norms. McClelland and Rumelhart (1985) demonstrated how "prototype" effects emerge from a distributed memory

model using an auto-associative network. According to such a model instances are overlaid in memory and stored by adjusting weights of connections between elements of a neural network. Patterns stored in a network in this manner can be regarded as vectors in a multidimensional space. Although an auto-associator is not a direct implementation of a norm-based coding model (Valentine, 1990), it implements the properties of vector-based similarity within a multidimensional space. McClelland and Rumelhart showed that when specific instances of similar patterns are learned, the network can respond to the prototype of the patterns, although the prototype itself has not been experienced. The network can appear to have abstracted multiple prototypes. Such an implementation of the norm-based coding model implies that the prototype effects emerge from the manner in which the memory traces are stored rather than the nature of the encoding processes per se. A PDP implementation has a number of advantages. It does not postulate an explicit process of extraction of a prototypical face or norm, nor does it require a "face-specific" process. Distinctiveness effects emerge simply because faces form a relatively homogeneous category of which many exemplars are experienced (cf. Diamond & Carey's, 1986, account of the inversion effect). As specific exemplars are overlaid in memory, it is possible for a large number of "old" exemplars to influence the norm but not to be individually retrievable any longer by any key as a specific instance. This provides a mechanism for previously seen but unfamiliar faces to influence recognition of familiar faces. Simulations of the effects of distinctiveness, inversion, and race in face recognition using a distributed memory model would be a useful line of future research.

In summary, the multidimensional framework is intended as a heuristic framework that could provide a potential link between many aspects of face recognition research that have proceeded somewhat independently. There are several aspects of the framework that have not been clearly defined. For example, the dimensions underlying the space have not been identified. The assumption of a Euclidean metric will almost certainly be an over-simplification. The exact nature of the decision process has not been made explicit. These aspects have not been specified simply because there is as yet insufficient evidence available to enable informed choices to be made. Two issues appear to be of particular theoretical interest: norm-based versus purely exemplar-based coding and the use of vector-based or inter-exemplar distance as a similarity measure. These theoretical issues may prove difficult to distinguish in experiments using photographs of faces as stimuli, but the use of "realistic" stimuli is essential to discover the processes by which people recognize faces. Further research on the effect of race is a potentially theoretically rich area as it could allow the effects of exemplar density and distance from a norm to be separated using natural stimulus classes.

REFERENCES

- Bartlett, J. C., Hurry, S., & Thorley, W. (1984). Typicality and familiarity for faces. *Memory and Cognition*, 12, 219-228.
- Bruce, V. & Valentine, T. (1985). Identity priming in the recognition of familiar faces. *British Journal of Psychology*, 76, 373-383.
- Buckhout, R. & Regan, S. (1988). Explorations in research on the other-race effect in face recognition. In M. M. Gruneberg, P. E. Morris, & R. N. Sykes (Eds.), *Practical aspects of memory: Current research and issues* (Vol. 1, pp. 40-46). Chichester: Wiley.
- Carey, S. & Diamond, R. (1977). From piecemeal to configurational representation of faces. *Science*, 195, 312-314.
- Chance, J. E., Turner, A. L., & Goldstein, A. G. (1982). Development of differential recognition for own and other-race faces. *Journal of Psychology*, 112, 29-37.
- Cohen, M. E. & Carr, W. J. (1975). Facial recognition and the von Restoff effect. *Bulletin of the Psychonomic Society*, 6, 383-384.
- Cross, J. F., Cross, J., & Daley, J. (1971). Sex, race, age and beauty as factors in recognition of faces. *Perception and Psychophysics*, 10, 393-396.
- Das-Smaal, E. A. & De Swart, J. H. (1984). Variation within categories. *Acta Psychologica*, 57, 165-192.
- Diamond, R. & Carey, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General*, 115, 107-117.
- Ellis, A. W., Young, A. W., Flude, B. M., & Hay, D. C. (1987). Repetition priming of face recognition. *Quarterly Journal of Experimental Psychology*, 38A, 193-210.
- Ellis, H. D. (1981). Theoretical aspects of face recognition. In G. M. Davies, H. D. Ellis, & J. Shepherd (Eds.), *Perceiving and remembering faces* (pp. 171-197). London: Academic Press.
- Ellis, H. D. & Deregowski, J. B. (1981). Within-race and between-race recognition of transformed and untransformed faces. *American Journal of Psychology*, 94, 27-35.
- Ellis, H. D., Deregowski, J. B., & Shepherd, J. W. (1975). Descriptions of white and black faces by white and black subjects. *International Journal of Psychology*, 10, 119-123.
- Feinman, S. & Entwisle, D. R. (1976). Children's ability to recognize other children's faces. *Child Development*, 47, 506-510.
- Flin, R. H. (1985). Development of face recognition: An encoding switch? *British Journal of Psychology*, 76, 123-134.
- Fodor, J. A. (1983). *Modularity of mind*. Cambridge, MA: MIT Press.
- Garner, W. R. (1978). Aspects of a stimulus: Features, dimensions and configurations. In E. H. Rosch & B. B. Lloyd, *Cognition and categorization* (pp. 99-131). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Going, M. & Read, J. D. (1974). The effect of uniqueness, sex of subject and sex of photograph on facial recognition. *Perceptual and Motor Skills*, 39, 109-110.
- Goldman, D. & Homa, D. (1977). Interogative and metric properties of abstracted information as a function of category discriminability, instance variability and experience. *Journal of Experimental Psychology: Human Learning and Memory*, 3, 375-385.
- Goldstein, A. G. (1975). Recognition of inverted photographs of faces in children and adults. *Journal of Genetic Psychology*, 127, 109-123.
- Goldstein, A. G. & Chance, J. E. (1980). Memory for faces and schema theory. *Journal of Psychology*, 105, 47-59.
- Johnson, N. L. & Leone, F. C. (1964). *Statistics and experimental design: In engineering and the physical sciences*. New York: Wiley.
- Kagan, J. & Klein, R. F. (1973). Cross-cultural perspectives on early development. *American Psychologist*, 28, 947-961.
- Krumhansl, C. L. (1978). Concerning the applicability of geometric models to similarity data: The interrelationship between similarity and spatial density. *Psychological Review*, 85, 445-463.
- Light, L. L., Kayra-Stuart, F., & Hollander, S. (1979). Recognition memory for typical and unusual faces. *Journal of Experimental Psychology: Human Learning and Memory*, 5, 212-228.
- Malpass, R. S. & Hughes, K. D. (1986). Formation of facial prototypes. In H. D. Ellis, M. A. Jeeves, F. Newcombe, & A. W. Young (Eds.), *Aspects of face processing* (pp. 154-162). NATO ASI series. Dordrecht: Martinus Nijhoff.
- Mandler, G. (1980). Recognizing: The judgement of previous occurrence. *Psychological Review*, 87, 252-271.
- McClelland, J. L. & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, 114, 159-188.
- McNicol, D. (1972). *A primer in signal detection theory*. Sydney: Allen & Unwin.
- Medin, D. L. & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85, 207-238.
- Micko, H. C. (1970). A "halo" model for multidimensional ratio scaling. *Psychometrika*, 35, 199-227.
- Neumann, P. G. (1977). Visual prototype information with discontinuous representation of dimension variability. *Memory and Cognition*, 5, 187-197.
- Nosofsky, R. M. (1986). Attention, similarity and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 1, 39-57.
- Palmer, S. E. (1975). Visual perception and world knowledge: Notes on a model of sensory-cognitive interaction. In D. A. Norman, D. E. Rumelhart, & LNR research group, *Explorations in cognition* (pp. 279-307). San Francisco: W. H. Freeman.
- Perkins, D. (1975). A definition of caricature, and caricature and recognition. *Studies in the Anthropology of Visual Communication*, 2, 1-24.
- Rae, G. (1976). A non-parametric measure of recognition performance. *Perceptual and Motor Skills*, 42, 98.
- Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, 3, 382-407.
- Rhodes, G., Brennan, S., & Carey, S. (1987). Identification and ratings of caricatures: Implications for mental representations of faces. *Cognitive Psychology*, 19, 473-497.
- Rhodes, G., Tan, S., Brake, S., & Taylor, K. (1989). Expertise and configural coding in face recognition. *British Journal of Psychology*, 80, 273-283.
- Shepherd, J. W. (1981). Social factors in face recognition. In G. M. Davies, H. D. Ellis, & J. Shepherd (Eds.), *Perceiving and remembering faces* (pp. 55-79). London: Academic Press.
- Shepherd, J. W. & Deregowski, J. B. (1981). Races and faces—a comparison of the responses of Africans and Europeans to faces of the same and different races. *British Journal of Psychology*, 20, 125-133.
- Valentine, T. (1988). Upside-down faces. A review of the effect of inversion upon face recognition. *British Journal of Psychology*, 79, 471-491.
- Valentine, T. (1990). Representation and process in face recognition. In R. Watt (Ed.), *Vision and visual dysfunction Vol. 14: Pattern recognition by man and machine* (Series ed. J. Cronley-Dillon). London: Macmillan.
- Valentine, T. & Bruce, V. (1986a). The effect of race, inversion and encoding activity upon face recognition. *Acta Psychologica*, 61, 259-273.
- Valentine, T. & Bruce, V. (1986b). Recognizing familiar faces: The role of distinctiveness and familiarity. *Canadian Journal of Psychology*, 40, 300-305.
- Valentine, T. & Bruce, V. (1986c). The effect of distinctiveness in recognizing and classifying faces. *Perception*, 15, 525-535.

- Valentine, T. & Endo, M. (1990). *The effects of distinctiveness and race on face recognition*. Paper presented to the Experimental Psychology Society, London, January.
- Winer, B. J. (1962). *Statistical principles in experimental design*. New York: McGraw-Hill.
- Winograd, E. (1981). Elaboration and distinctiveness in memory for faces. *Journal of Experimental Psychology: Human Learning and Memory*, 7, 181-190.
- Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, 81, 141-145.

Manuscript received 24 May 1990

The Coordination of Bimanual Aiming Movements: Evidence for Progressive Desynchronization

Barry Fowler, Tom Duck, Merv Mosher, and
B. Mathieson

York University, Downsview, Ontario, Canada

It is known that when simultaneous bimanual aiming movements are made to targets with different IDs (Index of Difficulty), Fitts' Law is violated. There is massive slowing of the easy target hand, but a debate has arisen over the degree of synchronization between the hands and whether this effect represents a coordinative structure or interference due to neural cross-talk. This issue was investigated in an experiment with 12 subjects who moved styli forward in the sagittal plane to pairs of targets that differed in difficulty (0.77/3.73 ID and 0.77/5.17 ID). Reaction time, movement time, and kinematic measures of resultant velocity and acceleration were analysed. The results showed clear-cut timing differences between the hands that depended on both the ID difference between target pairs and elapsed time of the movement. The violation of Fitts' Law was confined to the easy target hand. Pronounced individual differences in both timing differences and left-right asymmetry were also noted. Neither the coordinative structure nor the neural cross-talk models can fully account for these data, and it is possible that the initial constraints on movement are moderated by visually driven corrective movements.

Skilled performers appear capable of using their limbs virtually independently of each other to execute complex sequences of movements (Shaffer, 1981). On the other hand, untrained individuals exhibit constraints on tasks involving timing (for a review, see Keele, 1986) and aiming (Peterson, 1965; Robinson & Kravinsky, 1976). Understanding these constraints and how they are modified by practice is crucial to the development of a theory of

Requests for reprints should be sent to: B. Fowler, 338 Bethune College, York University, 4700 Keele Street, Downsview, Ontario, Canada, M3J 1P3.

The authors are grateful for the assistance provided by Susan Stapleton and Barry Kelso at various stages of this research and to Otmar Bock and an anonymous referee for comments on an earlier version of this paper. This work was in part funded by DCIEM Research Contract W-7-7011 to the first author.