PERGAMON

Vision Research 38 (1998) 2307–2321

Vision Research

Coding spatial variations in faces and simple shapes: a test of two models

Gillian Rhodes [a,b,*], Susan Carey [c], Graham Byatt [a,b], Fiona Proffitt [a]

[a] *University of Canterbury, Christchurch, New Zealand*
[b] *University of Western Australia, Nedlands, Perth WA 6907, Australia*
[c] *Massachusetts Institute of Technology, MA, USA*

Abstract

Faces all have the same basic elements in the same overall arrangement, and must be discriminated using variations in this shared configuration. An efficient way to represent these variations would be to code how each configuration differs from an average face (norm-based coding model). Alternatively, configurations could be represented simply by coding their absolute values in some coordinate system (absolute coding model). The two models differ in the variables they predict will influence an image's recognizability. Absolute coding predicts that recognizability will depend on an image's distinctiveness and degree of distortion from its veridical target. Norm-based coding predicts that recognizability will also depend on the way the image differs from a norm or average face, namely its distance from the norm and/or its degree of displacement from the norm-deviation vector for the target. We determined the effects of these four critical variables on recognition of undistorted (veridical) images, caricatures, anticaricatures and 'lateral' distortions of famous faces (Experiment 1), newly learned faces (Experiment 2), and simple shapes that also share a configuration (Experiment 2). The results favored absolute coding of both faces and shapes, and indicate that caricatures derive their power from their distinctiveness.

Keywords: Face recognition; Shape recognition; Caricature effects; Absolute coding; Norm-based coding

1. Introduction

We can each recognize thousands of individual faces, despite their considerable similarity[1]. This ability poses a challenge for theories of object recognition; in which objects are recognized by their parts and the overall arrangement of those parts [1,2]. Such theories cannot account for our ability to recognize faces, which all have the same basic parts in the same overall arrangement. Instead, faces must be discriminated using variations in this shared configuration [3]. In this paper we consider how such information is mentally represented.

Many have suggested that variations in a shared configuration could be represented using a norm-based coding scheme [6,14–21]. Each face to be encoded would be compared to a representation of an average face the norm and represented in terms of its deviations from that norm. The norm might function as a measuring device that controls the selection and/or weighting of features coded for a face, and the way each face deviates from the norm, either its distance or direction from the norm, or both, would capture important individuating information about that face [9]. Different norms could be used for different categories (sex, race, etc.) of faces, but for ease of exposition we will talk in terms of a single face norm.

Several lines of evidence seem consistent with some form of norm-based coding for faces. First, consider the effects of typicality or distinctiveness on face perception. Subjects can make stable judgements of typicality [22], and typical faces are judged to be faces (rather than jumbled non-faces) more quickly than atypical faces [17,20]. The attractiveness of a face also depends on how closely its configuration conforms to that of a norm or average face [23–25]. Given that distinctive faces generally lie further from the norm than more typical faces [26–28], these results suggest

* Corresponding author.

[1] For example, people can recognize the faces of over 90% of their old high-school classmates, even for classes with as many as 800 students [43], and we all know many more faces than those of our old schoolmates.

0042-6989/98/$19.00
PII: S0042-6989(97)00470-7

that a face's proximity to the norm has powerful effects on the perception of that face, and attest to the psychological relevance of the 'average face' or norm.

A second line of evidence suggests that newly encountered faces may be encoded in terms of features that deviate from the norm. In recognition memory tasks using unfamiliar faces, atypical faces and faces with a distinctive feature are remembered better than typical faces [22,20,28–33]. This advantage is predicted by the norm-based coding hypothesis in that, because atypical faces differ more from the norm than do typical faces [26–28], their features are more easily and effectively encoded.

A third line of evidence comes from studies showing that caricatures, which exaggerate how faces deviate from a norm, are recognized as well as, or better than, veridical representations, and better than anticaricatures, which reduce the differences between each face and the norm [6–9,19,34–36]. These results are readily explained by norm-based coding, because caricatures exaggerate, and anticaricatures reduce, crucial information about how each face deviates from the norm.

So far, we have suggested that caricature and typicality effects are consistent with norm-based coding of faces. However, it may also be possible to explain them without reference to norms. Suppose that faces are simply represented using their absolute values on some set of shared dimensions, so that each face occupies a discrete location in some multidimensional face space, as in Valentine's [20] absolute coding or exemplar-only model. Assuming that the density of faces in the space is highest in the center of the space (where the average face would lie if it existed), we see that typical or non-distinctive faces would lie in relatively high density regions of the space, and so would afford many more partial matches with non-targets than would atypical or distinctive faces. This difference in distracter activation would make typical faces harder to recognize than more distinctive faces. It would also contribute to the difficulty of recognizing anticaricatures, which are moved towards the norm, and therefore into more densely populated regions of face space. In contrast, caricatures are moved into less densely populated regions of face space, and the resulting increase in distinctiveness could offset the negative effects of the mismatch between a caricature and its target face. On this view, the power of caricatures comes simply from their distinctiveness, with no need to appeal to norm-based coding. Several theorists have proposed such an account of the power of caricatures [6,7,34,37].

In the present research we attempted to distinguish these two models by examining the variables each predicts will influence recognition performance. The absolute coding model predicts that the distinctiveness and degree of distortion (mismatch) of a stimulus image from the veridical target will affect the recognizability of that image [10,38]. The alternative norm-based coding account does not deny these effects of distortion and distinctiveness, but proposes that recognition will also be affected by the way each face differs from the norm or average face. For example, distance from the norm might play a role in recognition (and classification), with faces being easier to recognize (and harder to classify) if they are further from the norm [20]. Of course, distance from the norm is correlated with distinctiveness [25], but it may be possible to distinguish the unique effects of these two factors (see below).

Another possibility is that vectors originating from the norm define 'privileged directions' in face space, so that distortions which displace images off the norm-deviation vector for their corresponding target faces (i.e. which change the direction of the face from the norm), would disrupt recognition more than those which do not[2]. This conjecture is motivated by the observation that all images lying on the same norm-deviation vector have the same ordering of features in terms of relative distinctiveness. Consider a face whose most distinctive feature is its nose, whose next most distinctive feature is its mouth, and so on. This internal ordering is preserved if we move the face out along the vector (in a caricature), or back along the vector (in an anticaricature), but not if we move the face off that vector. If norm-deviation vectors define privileged directions in face space, as conjectured, then the degree of displacement of an image from the norm-deviation vector for its target would be another variable that influences recognition performance.

To summarize, the face space framework can support two distinct models of face recognition: the absolute coding model and the norm-based coding model. In both cases, faces are conceived of as locations in a multidimensional face space, but the models can be distinguished by the variables they predict will influence recognition. On the absolute coding view, degree of distortion from the veridical target image (distortion) and distinctiveness of the stimulus image (distinctiveness) should influence recognition performance. On the norm-based coding view, additional variables will be needed that reflect the relationship of the stimulus to

[2] One way of formalizing this notion is to say that the similarity of faces depends not only on the Euclidean distance between their point representations in face space (proximity in face space represents perceived similarity or confusability), as in the absolute coding model, but also on the angle between their norm-deviation vector representations. For example, the similarity of two faces might be a weighted function of angle and Euclidean distance, such as, similarity $=$ $k\cos(q)$/distance where q is the angle between the two vector representations in face space, and k is a weighting constant representing the relative contributions of vector angle and Euclidean distance to perceived similarity [10]. According to this equation, similarity will increase as the angle between the two vectors decreases, and as the distance between their point representations decreases.

the norm, such as distance from the norm (distance) and/or displacement from the norm-deviation vector for the target (norm-deviation). We attempted to distinguish these two classes of model by determining the unique effects of each variable on recognition performance. Distinctiveness was assessed using subject ratings, and other variables were experimentally manipulated by distorting images of faces.

In order to distort facial images we used a system of landmark points, as in previous studies [5–8,10,11,13]. In using this method, we assume that the distribution of faces in this image space parallels their distribution in a psychologically real face space whose dimensions correspond to the actual features used to encode faces. This assumption seems justified given that the landmark points capture information about the identity [11], age [12], sex and race [13] of a face, and that the distribution of faces in this image space systematically influences recognition performance [9].

Three kinds of distortions were made for the current study caricatures, anticaricatures and laterals. For each type of distortion, the landmark points on a face were moved the same Euclidean distance, but in a different direction, as illustrated in Fig. 1[3]. This procedure results in images that differ in distinctiveness, but which are equally distorted from the original. Caricatures are more distinctive, and anticaricatures are less distinctive, than the respective veridical images [25], because the density of representations in face space decreases as we move out from the norm [26–28]. The lateral distortions were designed to produce images that were intermediate in distinctiveness to caricatures and anticaricatures (an assumption that was confirmed by ratings of distinctiveness). As shown in Fig. 1, each landmark point in a lateral was moved off the norm-deviation vector for that point. According to the absolute coding model, laterals should be recognized better than anticaricatures and worse than caricatures, because they are of intermediate distinctiveness to caricatures and anticaricatures.

Preliminary data reported by Carey [39,40] showed that laterals were recognized more poorly than anticaricatures of famous faces, contrary to the predictions of absolute coding. This result is, however, consistent with

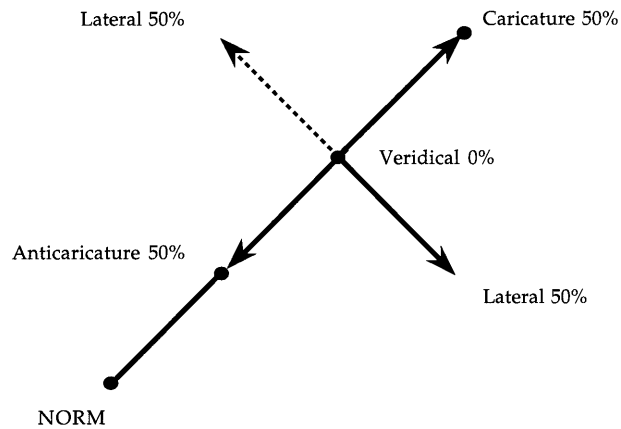

Fig. 1. Shows how a point on a face (veridical 0%) moves in a caricature, anticaricature and lateral. The corresponding point on the norm (NORM) is found (e.g. the tip of the nose on the two faces would be corresponding points) and the point on the face is moved relative to that norm point. In a 50% caricature the point on the face is moved 50% further away from the corresponding point on the norm, in the direction of the vector joining the two points. In a 50% anticaricature the point is moved 50% back along that vector towards the corresponding point on the norm. In the lateral caricature the point is moved orthogonally to the vector in one of the two possible directions shown.

the idea that norm-deviation vectors define privileged directions in face-space, so that moving an image off that vector (as in a lateral) is particularly disruptive to recognition. Poor performance on laterals is, therefore, prima facie evidence for norm-based coding. In line with this view, poor performance on laterals (Carey's stimuli) in a patient with impaired object recognition and intact face recognition has been taken as evidence for norm-based coding in the face system [21]. Unfortunately, the faces in Carey's study were not properly scaled to match the interpupil distance of the norm face before the laterals were produced, and the laterals were more distorted than they should have been. This extra distortion could, therefore, have contributed to poor performance on the laterals. In the experiments reported here we attempted to replicate this finding using properly scaled images.

We also examined another possible role for norms in encoding, one suggested by Attneave [41]. He conjectured that exposure to the norm for a class of stimuli would facilitate encoding by directing attention to those parts of the stimulus that differ most from the average, and from other members of the class. In support of this view, he found that people made fewer errors when learning the exemplars of an artificial category (polygons) if they first saw the norm from which those exemplars had been generated. In Experiment 2, we investigated whether prior exposure to an appropriate norm would facilitate learning (or recognition) of simple shapes. If it does, then we would have direct evi-

[3] In Fig. 1 we show how an individual landmark point moves in the various kinds of distortions. These transformations also generalize to the whole set of points for a face, i.e. to the whole-face vector consisting of the x and y coordinates of all the points. Linear algebra ensures that transformations have the same effect whether applied to each pairs of coordinates separately or to the whole-face vector. Therefore, if each landmark point on a face is equally distorted in two kinds of transformation (e.g. in an anticaricature and caricature transformation) then the whole-face vectors will also be equally distorted after these two kinds of transformation. Similarly, if the individual point vectors are moved orthogonally to the norm-deviation vector, as in a lateral transformation, then the whole-face vector is also moved orthogonally to the norm-deviation vector for that face.

dence of a functional role for norms in the mental representation of homogeneous objects. We also examined the effect of norm-exposure for faces, although it seems less likely that there would be any effect with a familiar class, for which subjects may already have norms.

Another goal was to investigate the generality of caricature effects in recognition. Based on neuropsychological and cognitive evidence, Farah and her colleagues have argued that faces are recognized using a special system or module [42]. One might ask, therefore, whether caricature effects reflect some special feature of this face recognition system. Initial evidence suggests that they do not, with caricature effects observed for recognition of birds [19], and inverted faces, which do not tap face-specific expertise [8]. Rather, caricature effects may be a very basic feature of perceptual systems in humans, other animals and connectionist networks [9]. We investigated the generality of caricature effects by including simple geometric shapes, as well as faces, in Experiment 2.

Finally, we investigated whether the same factors influence recognition of shapes and faces. If they do, then whichever model is supported (absolute or norm-based coding) would have to be considered a general model of how variations in shared configurations are coded, rather than a model that is specific to faces.

In summary, our main goal was to distinguish between absolute and norm-based coding models of face recognition, by testing which model better predicts the variables that influence recognition performance, and by directly comparing performance on anticaricatures and laterals. We also tested Attneave's conjecture that exposure to the norm might facilitate learning or recognition of a homogeneous class. Additional goals were to determine whether caricature effects generalize to homogeneous classes other than faces, and to determine whether or not the domain of the successful model is restricted to faces.

2. Experiment 1

We tested recognition of veridical (undistorted) drawings of famous faces, and 50% distortions (caricatures, anticaricatures, and lateral distortions) of those faces (Fig. 1). To assess the effects of the variables of interest on recognition performance, we coded each image on the four variables, and carried out (in section 2.2.4) multiple regression analyses. The absolute coding model predicts that only the degree of distortion from the target and distinctiveness will be significant predictors, whereas norm-based coding predicts that distance from the norm, and/ or displacement from the norm-deviation vectors will also be significant predictors, even when the effects of distortion and distinctiveness have been partialled out.

We also compared recognition performance directly for laterals and anticaricatures. As noted above, poorer recognition of laterals than anticaricatures, as found by Carey [39], would support norm-based coding. Better recognition of laterals than anticaricatures is predicted by absolute coding. However, this result could also occur in a norm-based coding system, if the greater distinctiveness of laterals outweighed the impact of their displacement from the norm-deviation vectors. The latter result would not, therefore, distinguish the two models.

A number of other results were expected (on either the model). First, we expected to replicate the usual finding that caricatures are recognized as well as, or better than, undistorted images, and better than anticaricatures. These caricature effects have been found for famous faces, personally known faces, and faces that were unfamiliar to the subjects prior to the experiment, so they do not require a high degree of familiarity with the faces or prior experience with caricatures of the faces [9]. They also occur with photographic quality images, as well as line drawings, and are not an artefact of impoverished stimuli [9,37]. Evidence that an agnostic with a damaged object recognition system, but intact face recognition, shows normal recognition of simple line drawing caricatures also indicates that these images activate the face recognition system, despite their superficial differences from real faces [21].

Second, we expected caricatures to be recognized better than laterals, because caricatures are more distinctive than laterals, and both are equally distorted from the veridical target (in the sense that their landmark points have been moved the same Euclidean distance in both cases). Third, we expected veridicals to be recognized better than laterals, because the two are similar in distinctiveness and only the laterals are distorted from the target. We tested our assumptions about the relative distinctiveness of the different kinds of distortion by collecting distinctiveness ratings for all the images.

2.1. Method

2.1.1. Subjects

Twenty-four students (12 male, 12 female) from the University of Canterbury were paid $15 each for participating.

2.1.2. Stimuli

Photographs of 14 female and 14 male famous faces were used to make the experimental stimuli. An additional five (one female, four male) famous faces were used to make practice stimuli. All of the faces had been used in a previous caricature recognition study [8]. For each face, we created seven drawings: A veridical (undistorted) drawing (V), an anticaricature (A), a cari-

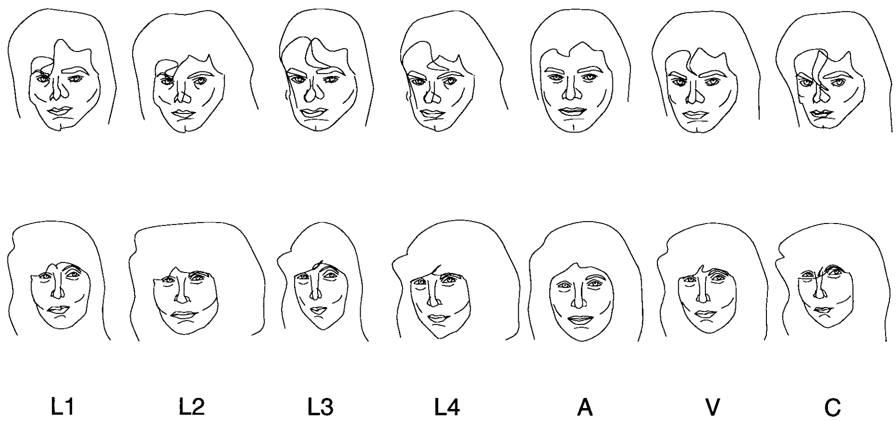

Fig. 2. A full set of distortions for two famous faces (Michael Jackson and Cher). L1, L2, L3, L4 = laterals (see text for details), A = anticaricature, V = veridical, C = caricature. All distortions are shown at the 50% level.

cature (C), and four laterals (Ls) (see below). All distortions were at the 50% level. Different norms were used to distort the male and female faces, with male and female norms created by averaging all 18 male faces and all 15 female faces, respectively.

Caricatures and anticaricatures were made in the normal way, by exaggerating or reducing the differences between each face and the appropriate norm. Specifically, the caricature generator computes a difference vector between each landmark point ($N = 169$) [8] on a veridical line drawing and the corresponding point on the norm (after matching the images on interpupil distance and horizontally aligning the pupils). Each of these norm-deviation vectors has an x-component and a y-component. A 50% caricature is created by multiplying each component by 0.5 and adding those values to the coordinates of the original point. To make a − 50% anticaricature, these values are subtracted from the coordinates of the original point.

To make a lateral, the points on the face are moved perpendicular to the direction of the norm-deviation vectors joining corresponding points on the faces and the norm (using a modified version of Brennan's caricature generator) [4,5]. There are two possible moves for each point, 'anti-clockwise' from the norm-deviation vector or 'clockwise' from the norm-deviation vector (Fig. 1). To move a point 'clockwise', the y-coordinate of the norm-deviation vector for that point is multiplied by 0.5 and added to the corresponding x-coordinate of the veridical, and the x-coordinate of the norm-deviation vector is multiplied by 0.5 and subtracted from the y-coordinate of the veridical. To move a point 'anticlockwise', the y-coordinate of the norm-deviation vector for that point is multiplied by 0.5 and subtracted from the corresponding x-coordinate of the veridical, and the x-coordinate of the norm-deviation vector is multiplied by 0.5 and added to the y-coordinate of the veridical.

To systematize the process of determining which of the two lateral vectors for each point on the face would be used, all points were classified as lying either to the right or left of center[4]. The program then moved all the points on the same side of the face in the same direction. This procedure produced four laterals for each face: L1 (all points moved anticlockwise), L2 (points on left side of face moved anticlockwise, points on right side moved clockwise), L3 (points on left side of face moved clockwise, points on right side moved anticlockwise), L4 (all points moved clockwise). Fig. 2 shows all seven distortions for two faces.

2.1.3. Procedure

Recognition was tested in a single session, lasting about 45 min. At the beginning of the session subjects were shown a list of the names of famous faces and told that they would be shown pictures of these faces in the experiment. Stimuli in the recognition test were blocked by sex of face. Half the subjects saw male faces first and half saw female faces first. They were told that on each trial a face would appear, which they should try to name as accurately as they could. A list of names for the faces in each block was present in front of the subject throughout the block. Subjects were asked not to say 'umm' or anything else that might trigger the voice key before they said the name. They initiated each trial by pressing the spacebar. A face then appeared in center of screen, where it remained until they responded. The experimenter recorded the subject's responses and noted any problem with the trial (e.g. voice-key prematurely activated by 'umm', or their voice failed to trigger the voice-key—the data from such trials were excluded from the reaction time analysis, but included for the accuracy analysis), and the computer (Macintosh LCIII) recorded reaction times.

[4] The point marking the vertex at the center of the upper lip fell on the midline and was not distorted.

The seven distortion types were shown once for each face. The two experimental blocks were preceded by a practice block of 35 trials, in which the seven distortions were shown for each of the five practice faces. Stimuli were randomized within the blocks.

Distinctiveness ratings were obtained in a separate session 1 week later, using the same procedure as for recognition. For each face, subjects were asked, "How distinctive is this face, that is how easily could you pick it out of a crowd?". Ratings were made on a 7-point scale (1 = low, 7 = high).

2.2. Results and discussion

2.2.1. Accuracy

A one-way ANOVA was carried out on the percent correct scores, with type of distortion as a repeated measures factor. Planned two-tailed t-tests were carried out to compare performance on laterals and anticarictures. Planned one-tailed t-tests were carried out to test for the usual pattern of performance on caricatures, veridicals and anticaricatures ($C \geq V > A$), and for the predicted (by both accounts) advantage of caricatures and veridicals over laterals.

A preliminary analysis showed that performance did not differ significantly for the four kinds of lateral (means = 44.5%, L1; 46.2%, L2; 47.1%, L3; 48.5%, L4), $F(3, 69) = 1.09$, ns. We therefore averaged the four scores to give one lateral score and four distortion levels (A, V, C, L) for each subject. There was a significant main effect of type of distortion, $F(3, 69) = 54.33$, $P < 0.00001$ (Fig. 3; top). Carey's [39] finding of worse performance on laterals than anticaricatures, which cannot be explained by absolute coding, was not replicated. Instead, laterals ($M = 46.7\%$) were recognized significantly more accurately than anticaricatures ($M = 32.7\%$), $t(69) = 6.89$, $P < 0.001$. As noted previously, this result does not distinguish the two models. As expected, caricatures ($M = 56.2\%$) were recognized as accurately as veridicals ($M = 53.8\%$), $t(69) = 1.18$, ns, and more accurately than anticaricatures, $t(69) = 11.56$, $P < 0.0005$ (one-tailed), replicating the usual caricature effect. Also as expected, laterals were recognized less accurately than veridicals, $t(69) = 3.49$, $P < 0.0005$ (one-tailed), and caricatures, $t(69) = 4.67$, $P < 0.0005$ (one-tailed).

No feedback was given during the recognition test, but it is still possible that some learning occurred. If it did then we may have overestimated performance on laterals, which were over-represented in the recognition test. We therefore carried out an additional analysis, in which only the first version of each face seen was included. The exact same pattern of results was obtained. Laterals ($M = 36.0\%$) were recognized significantly better than anticaricatures ($M = 22.1\%$) and worse than caricatures ($M = 49.8\%$), both $t > 2.17$, $P < 0.04$. Laterals were also recognized more poorly than veridicals ($M = 44.4\%$), although this difference was not significant, $t = 1.33$, ns.

2.2.2. Reaction times

RTs more than 2 SDs above the cell mean for each subject were discarded ($M = 1.3$ per subject). A one-way ANOVA was carried out on mean RTs with type of distortion as a repeated measures factor. A preliminary analysis showed that performance did not differ significantly for the four laterals (means = 3683 ms, L1; 3775 ms, L2; 3465 ms, L3; 3742 ms, L4), $F < 1$, so these were averaged to give one lateral score and four distortion levels (A, V, C, L) for each subject. There was a significant main effect of type of distortion, $F(3, 69) = 4.35$, $P < 0.008$ (Fig. 3; bottom). Laterals ($M = 3605$ ms) did not differ from anticaricatures ($M = 4066$ ms),

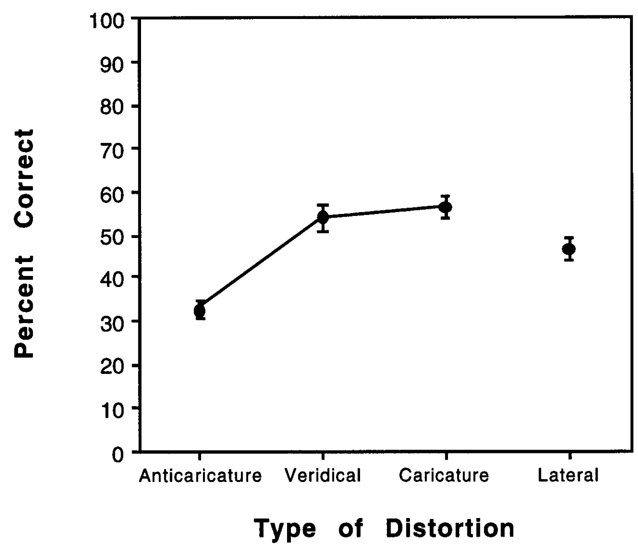

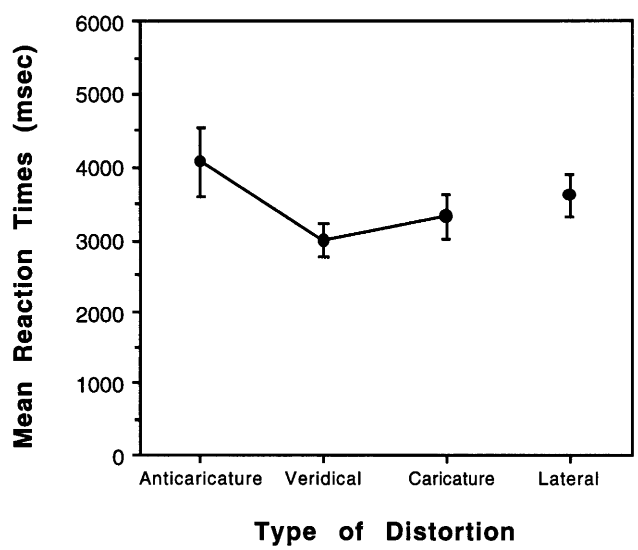

Fig. 3. Accuracy (top) and mean reaction times (bottom) as a function of type of distortion in Experiment 1. Standard error bars are shown.

Table 1
Coding of images on independent variables in Experiments 1 and 2

Type of image	Independent variable			
	Distortion	Distinctiveness	Distance	Norm-deviation
Experiment 1				
Anticaricature	1	Mean rating	0.5	0
Veridical	0	Mean rating	1.0	0
Caricature	1	Mean rating	1.5	0
Lateral	1	Mean rating	1.12	1
Experiment 2				
−50% Anticaricature	0.50	Mean rating	0.50	0
−25% Anticaricature	0.25	Mean rating	0.75	0
Veridical	0	Mean rating	1.00	0
25% Caricature	0.25	Mean rating	1.25	0
50% Caricature	0.50	Mean rating	1.50	0
25% Lateral	0.25	Mean rating	1.03[a]	0.25
50% Lateral	0.50	Mean rating	1.12[a]	0.50

[a] Distance of landmark points from the norm estimated using Pythagoras' theorem.

$t(69) = -1.50$, ns. Caricatures ($M = 3315$ ms) were recognized as quickly as veridicals ($M = 2999$ ms), $t(69) = -1.03$, ns (one-tailed), and more quickly than anticaricatures, $t(69) = 2.44$, $P < 0.01$ (one-tailed). Laterals ($M = 3605$ ms) were recognized more slowly than undistorted faces, $t(69) = 1.97$, $P < 0.05$ (one-tailed), but did not differ from caricatures, $t < 1$ (one-tailed). These results generally mirror the pattern of performance found for accuracy, and show no evidence of a speed-accuracy trade-off. In the follow-up analysis, with RTs for only the first-seen version of each face, there was no main effect of distortion type, $F(3, 72) = 1.48$, ns, and again no evidence of a speed-accuracy tradeoff.

2.2.3. Distinctiveness

Distinctiveness ratings showed good reliability with a coefficient alpha of 0.83. There was a significant main effect of type of distortion, $F(3, 69)$, $P < 0.0004$, with distinctiveness increasing from anticaricatures ($M = 3.8$), to laterals ($M = 4.2$), to caricatures ($M = 4.6$), both $t = 1.91$, $P < 0.05$, one-tailed, corroborating our assumptions about the relative distinctiveness of these different kinds of distortion. Caricatures ($M = 4.6$) were not significantly more distinctive than veridicals ($M = 4.5$), $t < 1$. These differences in distinctiveness between the different kinds of distortion mirrored the differences in recognition performance. This parallel is expected given that distinctiveness is an important determinant of recognition (predicted by both models)[5].

[5] Alternatively, subjects may simply have confused distinctiveness with ease of recognition. This seems plausible given that they had already been asked to recognize the faces. We eliminated this problem in Experiment 2, by using different subjects for the recognition and distinctiveness rating tasks.

2.2.4. Regression analyses

Each image was coded on the four variables of interest, as shown in Table 1. Distortion from the target was coded dichotomously, with 0 for veridical images and 1 for all distortions. Distinctiveness ratings, averaged across subjects, were used to code the distinctiveness of each image. Distance from the norm was coded 0.5, 1.0 and 1.5, respectively, for −50% anticaricatures, veridicals and 50% caricatures. These codings reflect the fact that, on average, caricatures are further from the norm than veridicals, and anticaricatures are closer to the norm than veridicals. For laterals, distance was coded as 1.12, using Pythagoras' theorem to give the relative distance of landmark points on the lateral from the corresponding points on the norm (compared with a distance of 1.0 for veridicals). The norm-deviation variable codes whether or not an image lies on the norm-deviation vector for its corresponding target image. Caricatures, anticaricatures and veridicals all fall on this vector and were coded 0, whereas laterals were displaced from this vector and were coded 1.

Multiple regression analyses were carried out to determine the contributions of the four variables of interest to variance in accuracy and RTs for correct responses. Table 2 summarizes the zero-order and partial correlations of accuracy and reaction times with each of the four independent variables. We will focus on the partial correlations, which indicate the unique contribution of each variable to recognition performance when the effects of the other variables are partialled out. A conservative significance level of $P < 0.01$ was adopted because we examined multiple correlations.

Table 2
Partial (and zero-order) correlations of each independent variable with accuracy and mean reaction times for famous faces (Experiment 1), initially unfamiliar faces (Experiment 2), and initially unfamiliar quads (Experiment 2)

	Distortion	Distinctiveness	Distance	Norm-deviation
Experiment 1				
Famous faces				
Accuracy ($N=188$)	−0.04 (−0.12)	**0.72 (0.75)**	0.07 **(0.27)**	0.06 (−0.01)
RT ($N=187$)	−0.00 (0.10)	**−0.57 (−0.55)**	**0.19** (−0.00)	0.00 (0.06)
Experiment 2				
Unfamiliar faces				
Accuracy ($N=187$)	**−0.32 (−0.29)**	**0.27 (0.27)**	−0.03 **(0.19)**	−0.08 (−0.08)
RT ($N=186$)	**0.35 (0.34)**	−0.14 (−0.14)	−0.03 (−0.16)	0.03 (0.09)
Unfamiliar quads				
Accuracy ($N=186$)	**−0.50 (−0.40)**	**0.50 (0.53)**	0.04 **(0.32)**	0.16 (0.02)
RT ($N=183$)	**0.41 (0.37)**	**−0.36 (−0.31)**	0.14 (−0.07)	−0.12 (0.01)

Correlations in bold type are significant at $P<0.01$.
Numbers vary due to removal of outlayers (images for whom residuals fall outside ± 2.5 S.D.).

Distinctiveness was the strongest predictor, and showed substantial partial correlations with both accuracy (positive correlation) and RTs (negative correlation). Surprisingly, distortion from veridicality made no unique contribution to performance. Note, however, that distortion was coded dichotomously and did not capture differences in the absolute amount of distortion associated with veridical faces at different distances from the norm (the absolute amount of distortion in a caricature, anticaricature or lateral is greater for faces that are further from the norm to begin with). Given that more distinctive faces generally lie further from the norm than less distinctive faces [26,27], we decided to recode distortion, giving higher distortion codings (coded as 2) to faces that fell above the median on distinctiveness than for those that fell below the median (coded as 1). However, this recoding still did not yield significant partial correlations with either accuracy ($r=0.04$) or reaction time ($r=-0.13$), although the zero-order correlations were significant, $r=0.27$, $P<0.001$, accuracy; $r=-0.22$, $P<0.01$, reaction time. Before considering other possible explanations, we will see whether this result replicates in Experiment 2.

There was a significant partial correlation between distance from the norm and RTs. Responses were faster to test images that were closer to the norm. Apparently, when distinctiveness is held constant, faces can be recognized more rapidly when they conform more closely to the average facial configuration. Note, however, that this effect of distance from the norm does not support norm-based coding, which predicts that performance should improve with distance from the norm. Rather, it may reflect a perceptual learning effect, where more frequent exposure to relatively average faces enables faster processing of those faces. There was also a significant zero-order correlation between distance from the norm and accuracy, which presumably reflects shared variance between distance and distinctiveness ($r=0.29$, $P<0.0001$). Neither the zero-order nor the partial correlations were significant for displacement from the target's norm-deviation vector. Overall, the partial correlation results are consistent with absolute, but not norm-based, coding accounts of recognition performance.

3. Experiment 2

We had several aims in Experiment 2. First, we wanted to replicate the results of Experiment 1 with a new set of faces. Second, we wanted to test Attneave's conjecture that prior exposure to the norm would facilitate learning (or recognition) of the members of a homogeneous class. Third, we wanted to compare caricature effects for faces and non-face objects (simple shapes) in a single experiment, to test our conjecture that caricature effects are not unique to faces. Finally, we wanted to compare the contributions of the four critical variables to recognition of faces and shapes, to determine whether the domain of the successful model is restricted to faces, or whether it also includes other homogeneous object classes.

We taught subjects to identify novel, randomly shaped quadrilaterals (quads), and subsequently tested their recognition of caricatures, anticaricatures, veridicals and lateral distortions of those shapes. Anticaricatures, laterals and caricatures were created by distorting the quads against a square, which was the average shape of the set[6]. To make conditions equivalent for

[6] The square also appears to be a psychological norm for quadrilaterals, being judged the best example of a quadrilateral [44], and showing the pattern of asymmetric similarity ratings associated with a prototype [45], i.e. quadrilaterals are judged more similar to the square than vice versa (unpublished data).

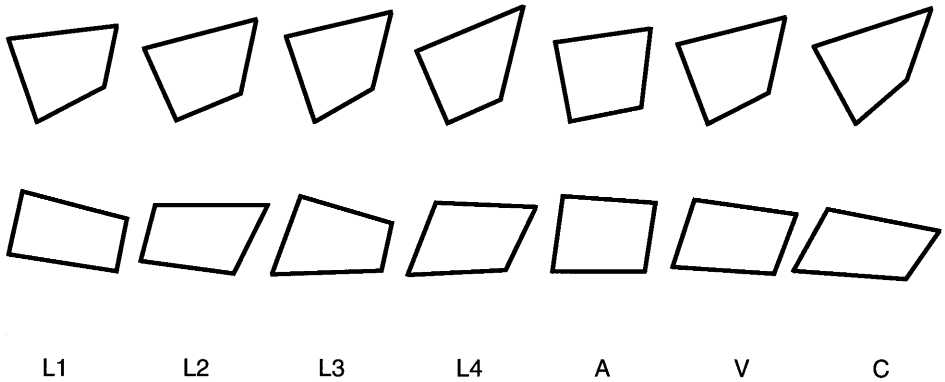

Fig. 4. A full set of 50% distortions for two quadrilaterals. L1, L2, L3, L4 are laterals, A is anticaricature, V is veridical, C is caricature.

faces and quads, we also trained subjects to identify unfamiliar faces and then tested recognition of anticaricatures, laterals, veridicals and caricatures of those faces. As noted earlier, caricature effects do not require prior familiarity with faces, so we expected to find similar patterns of results for these faces as for the famous faces used in Experiment 1 [9].

3.1. Method

3.1.1. Subjects

Forty-eight students (24 males, 24 females) from the University of Canterbury were paid \$20 each for participating in the recognition test. A total of 24 additional subjects (nine males, 15 females) from the University of Western Australia received either \$5 ($N = 20$) or course credit ($N = 4$) for rating the faces on distinctiveness.

3.1.2. Stimuli

A total of 28 unfamiliar faces (Riccarton high school faces from [8]) and 28 quads were used. A quad was made by randomly selecting one point from each of the four quadrants of a square, and joining these points with straight lines. The resulting shapes were then normalized for position (same center of mass) and size (area = 25 cm^2). For each face and quad, we created an anticaricature, a caricature, and four laterals (Fig. 4). Quad distortions were made using a square (5 × 5 cm), of equal area to the quads, as the norm. Male and female face distortions were made using the averages of all the male and female faces, respectively. Laterals for quads were made in the same way as the face laterals, with all points on each side of the midline moving in the same direction, to produce four different laterals for each quad. Distortions were created at the 25 and 50% levels. An additional five faces (three male, two female) and quads, and their associated distortions were used in practice trials. Each undistorted quad was assigned the name of one of the faces (e.g. 'Fred'), so that naming requirements were identical for the two types of object.

The final set of face images had the hair, brows and irises colored in, and impossible lines, and features obscured by the hair, erased (using superpaint). Such 'enhanced' drawings are normally more recognizable than the plain line drawings used in Experiment 1 (although they generally give similar caricature effects) [8].

3.1.3. Procedure

Subjects were tested in two 45 min sessions, held 2–7 days apart. Faces and quads were learned and tested in separate sessions. Training/testing sex of face. Analogously, sex of names was blocked by stimulous sex (names assigned to quads determined their 'sex'). Object order was counterbalanced with sex order

3.1.4. Training

Subjects learned to name (undistorted) faces and quads in blocks of seven. In a block each training image was shown together with its correct name, one after the other. These named images were then shown again. Next, the images were presented without names and the subject tried to name each one. The experimenter supplied the name if necessary. This cycle was repeated until all seven images were named correctly twice in a row. Recognition was tested immediately after each set of seven had been learned.

Photographs of faces were used in the face training, to facilitate comparisons of the recognition performance obtained in this study with that obtained for faces encoded under more naturalistic conditions, such as the famous faces used in Experiment 1. Drawings of the undistorted quads were used in the quads training. Therefore, the undistorted test images were identical to the training images for quads, but not for faces. This difference in training conditions means that recognition of veridicals may be disproportionately high (compared with recognition of the other distortions) for quads compared with faces. This point will be relevant when comparing the caricature effects obtained for faces and quads.

Each session began with a practice block, in which subjects were trained and tested on five practice faces or quads, depending on the session. Data from the practice trials were not analyzed.

3.1.5. Recognition test

Each type of distortion (A, V, C, L) was shown once for each exemplar. Only one of the four possible laterals was seen for each face, and the assignment of laterals to faces was counterbalanced across subjects, so that each kind of lateral was seen equally often in the experiment as a whole, and each subject saw the four kinds of laterals equally often. Subjects were told that, "the faces (quads) may be distorted but they will always be people (quads) that you have just learned", and that each person (quad) might appear more than once. They were told to try and name the test images as quickly and accurately as possible, and to avoid saying 'umm' or anything else before the name. The relevant list of names remained visible throughout the recognition test and subjects were told to consult the list if they couldn't remember a name.

Half the subjects saw 25% distortions and half saw 50% distortions. Distortion level (25 or 50%) was counterbalanced with training order (faces first or quads first) and sex order (male or female faces/names first). We also varied exposure to the norm. Half the subjects saw the appropriate norm (square or average same-sex face) prior to each training session and half did not. Exposure to the norm was counterbalanced with distortion level, training order and sex order.

3.1.6. Distinctiveness ratings

Subjects rated all the images from the recognition test on distinctiveness, using a 7-point scale. For faces, the same distinctiveness instructions were given as in Experiment 1. For quads, subjects were asked, 'How distinctive is this quad, that is, how easily could you discriminate it from other quads'? As in the recognition test, faces and quads were blocked, and object order was counterbalanced across subjects. Images were randomized within blocks. Each block began with the five practice exemplars and their distortions. Data from the practice trials were discarded.

3.2. Results and discussion

3.2.1. Training

A two-way ANOVA was carried out on the mean number of cycles required to learn each set of exemplars, with norm-exposure as a between-subjects factor and type of object as a repeated measures factor. Faces were learned significantly more quickly ($M = 3.7$ cycles) than quads ($M = 5.8$ cycles), $F(1, 46) = 44.21$, $P < 0.0001$, as might be expected given that people have more experience individuating faces than quads, and may be better able to detect subtle differences between individual faces than individual quads. Alternatively, faces might simply be less densely clustered than quads, independent of experience. Contrary to Attneave's suggestion, exposure to the norm did not facilitate learning, $F < 1$. Nor was there any interaction between norm-exposure and type of object, $F < 1$.

3.2.2. Recognition

Four-way ANOVAs were initially carried out on percent correct scores, and mean RTs for correct responses, with distortion level (25, 50%) and norm-exposure (shown, not shown) as between-subjects factors, and type of object (faces or quads) and type of distortion (anticaricature, veridical, caricature, lateral) as repeated measures factors. Reaction times more than 2 SDs above the cell mean for each subject were discarded ($M = 0.9$ per subject). Planned two-tailed t-tests were carried out to compare performance on laterals and anticaricatures for faces and quads. Planned one-tailed t-tests were carried out to test for the usual pattern of performance on caricatures, veridicals and anticaricatures ($C \geq V > A$), and for the predicted (by both accounts) advantage of caricatures and veridicals over laterals. Mean percent correct and RTs are shown in Table 3.

The initial ANOVAs showed no significant difference in accuracy or RTs between caricatures and anticaricatures for 25% distortions. As there distortions were too subtle to generate this basic caricature effect, we reanalyzed the data for the 50% distortion level only. It is those ANOVAs that we report in detail here.

Fig. 5 shows recognition performance for faces and quads, together with the famous face results from Experiment 1 (for ease of comparison). Note first that performance was qualitatively similar for newly learned and famous faces, confirming previous findings of similar caricature effects for these two sorts of faces [9]. Caricature effects do not therefore require a great deal of familiarity with the faces used, or any previous exposure to caricatures of those faces . Second, we did not replicate Carey's [39] finding that laterals were recognized more poorly than anticaricatures for either faces or quads (a finding that would have supported norm-based coding). Third, faces and quads showed similar patterns of performance. In particular, the usual pattern of superior performance on caricatures compared with anticaricatures was observed for quads as well as faces. The main difference between faces and quads was disproportionately high performance on veridical (compared with distorted) quads, which was expected given that the veridical quads had been used as the training images. For faces, where all the test images (enhanced drawings) differed qualitatively from the training images (photographs), there was no such peak in performance for veridicals.

Table 3
Mean percent correct and reaction times (ms) for each type of object, type of image (combines distortion type and distortion level) in Experiment 2

Type of image	Anticaricature −50%	Anticaricature −25%	Veridical[a] 0%	Caricature 25%	Caricature 50%	Lateral 25%	Lateral 50%
Accuracy							
Faces	38.4	54.6	53.2	57.4	51.8	53.6	47.0
Quads	41.7	66.8	69.0	69.3	61.9	69.2	59.2
Mean	40.1	60.7	61.1	63.4	56.9	61.4	53.1
Mean RTs							
Faces	3703	2247	2720	2070	2645	2049	3085
Quads	2597	2031	2185	1983	2553	1879	2522
Mean	3150	2139	2453	2027	2599	1964	2804

[a] Figures shown for 0% V images are from subjects in the 50% distortion condition.

There were significant main effects of type of object and distortion, for both accuracy, $F(1, 22) = 18.86$, $P < 0.0003$ (type of object), $F(3, 66) = 50.16$, $P < 0.0001$ (distortion), and RTs, $F(1, 22) = 6.41$, $P < 0.02$ (type of object), $F(3, 66) = 7.53$, $P < 0.0002$ (distortion), and these two factors interacted for both accuracy, $F(3, 66) = 4.59$, $P < 0.006$, and RTs, $F(3, 66) = 3.78$, $P < 0.02$ (Fig. 5). The planned comparisons between different types of distortion were therefore carried out for each type of object separately.

As noted above we did not find poorer performance on laterals than anticaricatures. For faces, laterals were recognized more quickly and accurately than anticaricatures, both $t > 2.89$, $P < 0.01$, as in Experiment 1. For quads, laterals were recognized more accurately than (and as quickly as, $t < 1$) anticaricatures, $t(66) = 7.14$, $P < 0.001$. As noted earlier, this result can be accommodated by either model, and does not serve to distinguish the models. Exposure to the norm did not facilitate recognition performance, for either faces or quads. There were no significant main effects or interactions involving exposure to the norm, all $F < 2.13$. These results provide no support for norm-based coding.

A similar caricature effect was observed for faces and quads, with caricatures recognized better than anticaricatures, for both faces and quads. This advantage was apparent in accuracy in both cases, both $t > 5.48$, $P < 0.0005$ (one-tailed), and in speed for faces, $t(66) = 4.95$, $P < 0.0005$ (one-tailed), but not for quads, $t < 1$. There was no advantage for caricatures over veridicals for either type of object. For faces, caricatures were recognized as quickly and accurately as veridicals, both $t < 1$. For quads, veridicals were actually recognized more accurately than (and as quickly as) caricatures, $t(66) = -2.90$, $P < 0.01$, no doubt because the veridical test quads were used in training. Not surprisingly, veridicals were recognized more accurately than laterals for both faces and quads, both $t > 2.53$, $P < 0.01$ (one-tailed), and more quickly than laterals for faces, $t(66) = 1.71$, $P < 0.05$ (one-tailed), and (marginally for) quads, $t(66) = 1.58$, $P < 0.10$ (one-tailed). Caricatures were also recognized more accurately than laterals for faces, $t(66) = 1.96$, $P < 0.05$ (one-tailed), as in Experiment 1, and for quads, although the latter difference was not significant, $t(66) = 1.10$, ns. Caricatures were recognized more quickly than laterals for faces, $t(66) = 2.05$, $P < 0.05$ (one-tailed), but not for quads, $t < 1$. The absence of a clear advantage for caricatures over laterals, for quads, raises the possibility that the effect of displacement from the norm-deviation direction may not be as great for quads as for faces. However, this suggestion was not supported by the results of the regression analyses, which indicated that the norm-deviation variable did not correlate with performance for either quads or faces (see section 3.2.4).

3.2.3. Distinctiveness

The distinctiveness ratings were highly reliable, with coefficient alphas of 0.94 for faces and 0.96 for quads. The main purpose of the distinctiveness ratings was to allow the test images to be coded on distinctiveness for the regression analyses. However, an ANOVA did confirm the expected effect of type of distortion on distinctiveness, $F(1, 23) = 27.60$, $P < 0.0001$, with distinctiveness increasing from anticaricatures ($M = 2.8$), to laterals ($M = 4.1$) to caricatures ($M = 4.6$). Caricatures were also more distinctive than veridicals ($M = 3.6$), as expected. These differences were found for both faces and quads (all $P < 0.0001$).

3.2.4. Regression analyses

The image codings used for each variable are shown in Table 1. By using two distortion levels, we created greater variation in the distortion variable than in Experiment 1. Therefore, we expected to find an effect of distortion as well as an effect of distinctiveness. Increasing the variation in the norm-based coding variables also gave greater power to detect any effect of those variables than in Experiment 1.

Table 2 shows the partial (and zero-order) correlations of accuracy and RTs with each of the four independent variables. For faces, both distortion and distinctiveness were significant predictors of performance, as expected on both models. Neither of the norm-based coding variables, distance from the norm or displacement from the norm-deviation vector of the target, contributed uniquely to the variance in recognition performance (although distance from the norm was highly correlated with distinctiveness, as expected, $r = 0.74$, $P < 0.0001$, accuracy; $r = 0.72$, $P < 0.0001$, analyses). For faces, therefore, the results support absolute coding, not norm-based coding.

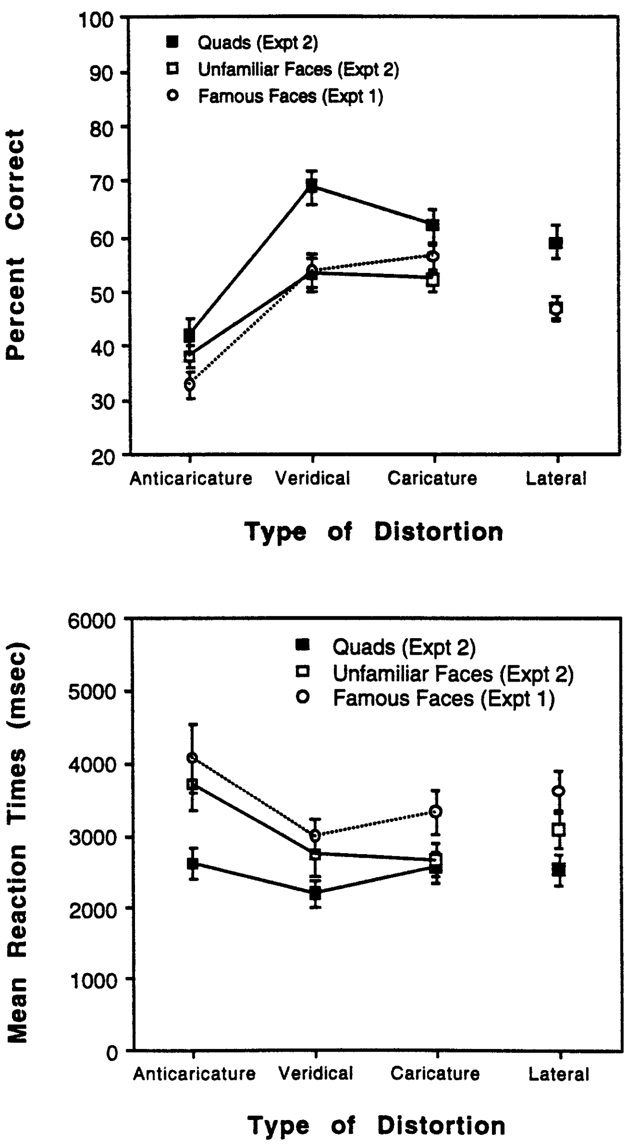

Fig. 5. Accuracy (top) and mean reaction times for correct responses (bottom) as a function of type of distortion and type of object (face or quad) in Experiment 2. Only 50% distortions are shown. The data for famous faces from Experiment 1 are superimposed for ease of comparison. Standard error bars are shown.

For quads, as for faces, distortion and distinctiveness both contributed significantly to performance, when the effects of the other variables were partialled out. Also as for faces, distinctiveness and distance from the norm were highly correlated, $r = 0.59$, $P < 0.0001$, accuracy; $r = 0.59$, $P < 0.0001$, RTs. Neither of the norm-based coding variables contributed uniquely to recognition performance, although there was a significant (positive) zero-order correlation between distance from the norm and RTs. These results therefore support absolute coding for quads as well as faces.

In summary, the recognizability of faces and simple shapes depended primarily on their distinctiveness, with more distinctive images recognized more readily than less distinctive images. Images were also easier to recognize if they matched their respective veridical targets more closely (other things being equal). There was no evidence from the partial correlations that moving a face or shape off its characteristic norm-deviation vector disrupted recognition any more than would be expected from the resulting changes in distinctiveness and distortion from the veridical target. Therefore, our hypothesis that norm-deviation vectors might define ‘privileged directions in face-space’ was not supported. Nor was there any evidence that images lying further from the norm were easier to recognize than would be expected given their greater distinctiveness.

Overall, these results support absolute coding of homogeneous classes and suggest that caricatures derive their power primarily from their distinctiveness.

4. General discussion

We began by noting that all faces have same basic elements in the same overall arrangement (i.e. faces form a homogeneous class), so that recognition requires us to discriminate variations in this shared configuration. We suggested that such variations could be effectively represented by coding how each face varies from the average configuration or norm, and reviewed several typicality and caricature effects that are consistent with such a view. We noted, however, that members of a homogeneous class could also be represented as absolute values on a shared set of dimensions, as in Valentine’s exemplar model [20], in which case a norm would play no special role in the encoding or recognition of faces.

We argued that these two models make different predictions about which variables will influence recognition performance. Absolute coding predicts that the recognizability of a stimulus image will depend on two variables: Its degree of distortion from the veridical target image and its distinctiveness. Norm-based coding predicts that additional variables will be needed, which reflect the relationship of the stimulus to the norm. One such variable is the distance of the stimulus from the norm and another is whether or not the stimulus falls on the norm-deviation vector for the target.

We sought to distinguish the two models by examining the effects of these four variables on recognition performance. Recognition was tested for undistorted faces, as well as for three types of distortion: caricatures, anticaricatures and laterals. Together, this set of stimuli provided images that differed on the four variables of interest. We also attempted to replicate Carey's [39] finding that laterals were harder to recognize than anticaricatures (a result that could only be accounted for by norm-based coding), and investigated whether exposure to an appropriate norm during training would facilitate learning and/or recognition of simple geometric shapes and faces. Finally, we investigated whether or not caricature effects were specific to face recognition, and whether or not the domain of the successful model is restricted to faces.

There was clear support for the predictions of absolute coding, and no direct support for norm-based coding. Distinctiveness was a clear predictor of recognition performance (for famous faces, initially unfamiliar faces and quads), with degree of distortion from veridicality also playing a role, at least for unfamiliar faces and quads. The norm-based coding variable of displacement from the target's norm-deviation vector had no effect on performance. The other norm-based coding variable, distance from the norm, influenced recognition speed for famous faces, but the effect was in the opposite direction to that expected, with faster responses to images that were closer to the norm. This result may indicate a perceptual learning effect resulting from greater experience with more average faces.

We did not replicate Carey's [39] finding of poorer performance on laterals than anticaricatures, which would have ruled out an absolute coding model. Laterals were recognized better than anticaricatures, which is consistent with either model. We also found no support for Attneave's [41] conjecture that exposure to an appropriate norm would facilitate learning or recognition of homogeneous objects. Overall, the results favored a model in which faces are mentally represented as absolute values on a set of common dimensions.

Byatt and Rhodes [10] also failed to find evidence for norm-based coding, in a cross-race recognition study. They reasoned that if people code faces as deviations from a norm, then a caricature should only be effective if it exaggerates how a face deviates from the norm people actually use to code that face. Faces from an unfamiliar race would be coded as deviations from an own-race norm (assuming that norms come from experience), and so caricatures of other-race faces should be more effective if they exaggerate deviations from the own-race norm than from an objectively more appropriate (but unavailable) other-race norm. Contrary to this prediction, European subjects recognized caricatures of Chinese faces better when they were made using a Chinese norm than a European norm.

In the present experiments, similar variables influenced recognition of both quads and faces, suggesting that absolute coding is not specific to faces, or to homogeneous objects with which we have expertise. Rather, absolute coding may provide a general model for the coding of any set of objects that share a configuration.

These experiments also have implications for understanding caricature effects. First, support for absolute coding suggests that caricatures derive their power from their distinctiveness (i.e. their location in low density regions of face space), which appears to offset the detrimental effect of being distorted from the veridical target (at least within the limits of distortion used in this and similar studies). Second, the fact that a similar caricature effect (caricatures recognized as well as veridicals and better than anticaricatures) was found for quads as for faces indicates that caricature effects are not restricted to face recognition or to homogeneous classes with which we are experts. We should note, however, that caricatures were not superportraits in these experiments, i.e. they were not recognized better than veridicals, and that a superportrait effect has only ever been observed when subjects have expertise (or at least extensive training) with the stimulus class [9]. Therefore, the possibility remains open that expertise is required for caricatures to be superportraits.

In conclusion, Valentine [20] has suggested that the multidimensional space framework is a useful heuristic for thinking about the representation of faces in memory. Our results show that of the two models that can be supported by this framework, absolute and norm-based coding, absolute coding has stronger empirical support. Furthermore, absolute coding is not specific to faces, or to homogeneous objects with which we have expertise, but may be used whenever we must discriminate objects that share a configuration.

Acknowledgements

This work was supported by grants from the Foundation for Research, Science and Technology (UOC320, UOC414) and the University of Western Australia, to the first author, and a grant from the National Institutes of Health (R01HD22166) to the second author. This work has benefited greatly from discussions with Rhea Diamond and Kamal Nigam. We thank John Breukelaar, Alison Clark, Catherine Hickford and Craig Webster for running subjects in Experiment 3, Catherine Hickford and Craig Webster for assistance with data analysis, and John Breukelaar for assistance with stimulus construction. We also thank two anonymous reviewers for their comments on an earlier draft of the paper.

References

[1] Marr D. Vision. San Francisco: Freeman, 1982.
[2] Biederman I. Recognition-by-components: a theory of human image understanding. Psychol Rev 1987;94:115–47.
[3] Diamond R, Carey S. Why faces are and are not special: an effect of expertise. J Exp Psychol: Gen 1986;115:107–17.
[4] Brennan SE. Caricature Generator. Unpublished Master's thesis. Cambridge, MA: MIT, 1982.
[5] Brennan SE. The caricature generator. Leonardo 1985;18:170–8.
[6] Rhodes G, Brennan S, Carey S. Identification and ratings of caricatures: implications for mental representations of faces. Cogn Psychol 1987;19:473–97.
[7] Benson PJ, Perrett DI. Perception and recognition of photographic quality facial caricatures: implications for the recognition of natural images. Eur J Cogn Psychol 1991;3:105–35.
[8] Rhodes G, Tremewan T. Understanding face recognition: caricature effects, inversion and the homogeneity problem. Vis Cogn 1994;1:275–311.
[9] Rhodes G. Superportraits: Caricatures and Recognition, Hove and London: The Psychology Press, 1996.
[10] Byatt G, Rhodes G. Recognition of own race and other race caricatures: implications for models of face recognition. Vis Res Special Issue on Models of Object Recognition. 1997 (this issue).
[11] Hancock PJB, Burton AM, Bruce V. Face processing: human perception and principal components analysis. Mem Cogn 1996;24:26–40.
[12] Burt DM, Perrett DI. Perception of age in adult Caucasian male faces: computer graphic manipulation of shape and color information. Proc R Soc Lond Ser B 1995;259:137–43.
[13] Hill H, Bruce V, Akamatsu S. Perceiving the sex and race of faces: the role of shape and color. Proc R Soc Lond Ser B 1995;261:367–73.
[14] Hochberg JE. Perception, 2nd ed. New Jersey: Englewood Cliffs, Prentice-Hall, 1978.
[15] Goldstein AG, Chance JE. Memory for faces and schema theory. J Psychol 1980;105:47–59.
[16] Ellis HD. Theoretical aspects of face recognition. In: Davies GH, Ellis HD, Shepherd J, editors. Perceiving and Remembering Faces. London: Academic Press, 1981.
[17] Valentine T, Bruce V. The effects of distinctiveness in recognising and classifying faces. Perception 1986;15:525–35.
[18] Bruce V. Recognising Faces. Hove and London: Erlbaum, 1988.
[19] Rhodes G, McLean IG. Distinctiveness and expertise effects with homogeneous stimuli: towards a model of configural coding. Perception 1990;19:773–94.
[20] Valentine T. A unified account of the effects of distinctiveness, inversion, and race in face recognition. Quart J Exp Psychol 1991;43A:161–204.
[21] Moscovitch M, Winocur G, Behremann M. What is special about face recognition? Nineteen experiments on a person with visual object agnosia and dyslexia but normal face recognition. J Cogn Neurosci 1997;9:555–604.
[22] Bartlett JC, Hurry S, Thorley W. Typicality and familiarity of faces. Mem Cogn 1984;12:219–28.
[23] Langlois JH, Roggman LA. Attractive faces are only average. Psychol Sci 1990;1:115–21.
[24] Langlois JH, Roggman LA, Musselman L. What is average and what is not average about attractive faces? Psychol Sci 1994;5:214–20.
[25] Rhodes G, Tremewan T. Averageness, exaggeration and facial attractiveness. Psychol Sci 1996;7:105–10.
[26] Bruce V, Coombes A, Richards R. Describing the shapes of faces using surface primitives. Image Vis Comput 1993;11:353–63.
[27] Bruce V, Burton AM, Dench N. What's distinctive about a distinctive face? Q J Exp Psychol 1994;47A:119–41.
[28] Johnston RA, Ellis HD. The development of face recognition. In: Valentine T, editor. Cognitive and Computational Aspects of Face Processing: Explorations in Face Space. London: Routledge, 1995.
[29] Going M, Read JD. Effects of uniqueness, sex of subject, and sex of photograph on facial recognition. Percept Mot Skills 1974;39:109–10.
[30] Cohen ME, Carr WJ. Facial recognition and the Von Restorff effect. Bull Psychon Soc 1975;6:383–4.
[31] Light LL, Kayra-Stuart F, Hollander S. Recognition memory for typical and unusual faces. J Exp Psychol Hum Learn Mem 1979;5:212–9.
[32] Winograd E. Elaboration and distinctiveness in memory for faces. J Exp Psychol Hum Learn Mem 1981;7:181–90.
[33] Johnston RA, Ellis HD. Age effects in the processing of typical and distinctive faces. Q J Exp Psychol 1995;48A:447–65.
[34] Benson PJ, Perrett DI. Visual processing of facial distinctiveness. Perception 1994;23:75–93.
[35] Mauro R, Kubovy M. Caricature and face recognition. Mem Cogn 1992;20:433–40.
[36] Stevenage SV. Can caricatures really produce distinctiveness effects? Br J Psychol 1995;86:127–46.
[37] Calder AJ, Young AW, Benson PJ, Perrett DI. Self priming from distinctive and caricatured faces. Br J Psychol 1996;87:141–62.
[38] Rhodes G, Byatt G, Tremewan T, Kennedy A. Facial distinctiveness and the power of caricatures. Perception 1997;24:207–23.
[39] Carey S. Becoming a face expert. Philos Trans R Soc Lond Series B 1992;335:95–103.
[40] Nigam K. Mechanisms for the recognition of quadrilaterals and faces. Massachusetts: Massachusetts Institute of Technology, 1994.
[41] Attneave F. Transfer of experience with a class-schema to identification-learning of patterns and shapes. J Exp Psychol 1957;54:81–8.
[42] Farah MJ. Is face recognition 'special'? Evidence from neuropsychology. Behav Brain Res 1996;76:181–9.

[43] Bharick HP, Bharick OO, Wittlinger RP. Fifty years of memory for names and faces: a cross-sectional approach. J Exp Psychol Gen 1975;104:54–75.
[44] Rosch EH. On the internal structure of perceptual and semantic categories. In: Moore TE, editor. Cognitive Development and The Acquisition of Language. New York: Academic Press, 1973.
[45] Tversky A. Features of similarity. Psychol Rev 1977;84:327–52.