Perception, 2002, volume 31, pages 1221 – 1240

DOI:10.1068/p3252

Effects of geometric distortions on face-recognition performance

Graham J Hole, Patricia A George, Karen Eaves, Ayman Rasek
School of Cognitive and Computing Sciences, University of Sussex, Falmer, Brighton BN1 9QH, UK
Received 23 July 2001, in revised form 27 May 2002; published online 10 September 2002

Abstract. The importance of 'configural' processing for face recognition is now well established, but it remains unclear precisely what it entails. Through four experiments we attempted to clarify the nature of configural processing by investigating the effects of various affine transformations on the recognition of familiar faces. Experiment 1 showed that recognition was markedly impaired by inversion of faces, somewhat impaired by shearing or horizontally stretching them, but unaffected by vertical stretching of faces to twice their normal height. In experiment 2 we investigated vertical and horizontal stretching in more detail, and found no effects of either transformation.

Two further experiments were performed to determine whether participants were recognising stretched faces by using configural information. Experiment 3 showed that nonglobal vertical stretching of faces (stretching either the top or the bottom half while leaving the remainder undistorted) impaired recognition, implying that configural information from the stretched part of the face was influencing the process of recognition — ie that configural processing involves *global* facial properties. In experiment 4 we examined the effects of Gaussian blurring on recognition of undistorted and vertically stretched faces. Faces remained recognisable even when they were both stretched and blurred, implying that participants were basing their judgments on configural information from these stimuli, rather than resorting to some strategy based on local featural details. The tolerance of spatial distortions in human face recognition suggests that the configural information used as a basis for face recognition is unlikely to involve information about the absolute position of facial features relative to each other, at least not in any simple way

1 Introduction

We are able to recognise familiar faces reasonably well despite transient changes in their pose (orientation relative to the viewer), expression, and lighting conditions (eg Moses et al 1996; Troje and Bülthoff 1996). Any theory of the structural encoding phase of face recognition, whether it comes from the standpoint of computer science or cognitive psychology, needs to explain how we manage to compensate for these kinds of perturbations in making a match between seen and remembered faces (Bruce et al 1992; Bruce 1988, 1994). To understand specifically how humans recognise faces, one also has to explain something else: the *failures* in human face-recognition performance. Humans find it difficult to cope with certain kinds of alteration to facial images, such as presenting them in photographic negative (eg Galper 1970; Kemp et al 1990) or upside down (Yin 1969; review in Valentine 1991). The fact that we experience such difficulties provides some insight into the processes underlying the construction and use of structural codes in human face recognition. For example, presenting a face in negative appears to affect recognition primarily by disrupting processes that use shading information to produce a representation of the 3-D shape of a face (Johnston et al 1992; Kemp et al 1996; George and Hole 2000). The deleterious effects of inversion on face recognition have drawn attention to the importance of the configural[(1)] properties of upright faces for recognition (eg Young et al 1987; Rhodes et al 1993).

[(1)] See Searcy and Bartlett's (1996) summary of the various theoretical interpretations of 'configural', 'holistic', or 'relational' processing. The terminology in this area is somewhat confused at present: sometimes these terms have been used interchangeably but they have also been used to refer to quite different types of processing. We are using the term 'configural' processing to refer to a form of processing which is primarily based on information about the spatial relationships between local 'features' (eg eyes, nose, and mouth). In contrast, 'piecemeal' or 'featural' processing would be based primarily on details of the individual facial features themselves, rather than on their spatial relationships to each other.

Inversion is just one of a variety of possible geometrical transformations that could be performed on a face: what would be the effects on recognition of other types of transformation? Moderate changes of size and position of a face within the visual field might be expected to have little effect on face recognition (as has been demonstrated recently by Brooks et al 2002). However, the effects of other transformations are harder to predict. Exploring the effects of other types of alteration may throw further light on the processes underlying face recognition (Bruce et al 1992). In particular, it may provide further insight into the precise nature of configural processing. For example, stretching a face vertically or horizontally disrupts certain configural properties of the face: absolute distances such as the length of the nose are changed, as well as ratio distances such as the ratio between eye separation and nose length. If recognition is impaired under these circumstances, then this suggests that the values of some or all of these distances are important for recognition. Through the experiments reported in this paper, therefore, we examine the consequences for face recognition of different kinds of geometric distortion, principally affine linear transformations.

To our knowledge, there are few data on the effects on face recognition of geometric transformations of this kind. Some studies have examined our ability to cope with the distortions of virtual space in pictures that are produced by changes in viewpoint; however, these studies have a different emphasis to that of the present experiments, since they are concerned primarily with how we perceive the distortions themselves, rather than how we manage to continue to recognise what the picture portrays (see Busey et al 1990).

A number of studies (eg Haig 1984; Hosie et al 1988; Kemp et al 1990; Bruce et al 1991; Bruce 1994) have shown that we can detect very small changes in the positions of facial features, implying that we possess a very high level of sensitivity to the spatial arrangement of features within faces. However, these studies may give a misleading impression of how we actually use this information in recognising faces. Just because we are capable of making fine-grain discriminations of feature positions within a face, it does not mean that this is necessarily something that we actually do, routinely, in the normal course of face recognition. The ability to notice minor changes to the configural properties of a face is not necessarily incompatible with being able to tolerate a fair amount of 'noise' in the precise configural properties of the face for recognition purposes. In fact, the latter must be true, because if recognition were too heavily dependent on the minutiae of the configuration of a face, recognition would fail with every minor change of target-face orientation and lighting conditions with respect to the viewer.

Somewhat different conclusions about configural processing come from studies on the effects of caricature on face recognition (eg Rhodes et al 1987; Benson and Perrett 1991). These have examined participants' ability to recognise faces which have been subjected to a systematic transformation whose nature depends on the relationship between the individual face being transformed and some 'average' or 'prototypical' face. A caricature exaggerates the differences between the spatial characteristics of an individual face and those of the prototype. 'Anticaricatures' minimise these differences, and 'lateral' caricatures make systematic changes to a face that are unrelated to the differences between the individual face and the average (see Rhodes and Tremewan 1994; Lewis and Johnston 1998, 1999). The results of these experiments tend to differ somewhat depending on whether real or line-drawn faces are used (Rhodes et al 1987; Benson and Perrett 1991, 1994), but they generally demonstrate that face recognition is unaffected or even enhanced by some types of systematic distortion of the face (ie caricatures). For the present discussion, it is worth noting that anticaricatures and lateral caricatures remain recognisable until taken to extreme levels. These studies thus suggest that recognition can survive (or even be facilitated by) fairly large perturbations in the configural information of a face. Furthermore, they suggest that what is important for recognition is not the precise spatial interrelationship of the facial features, but the differences between faces, in terms of the differences either between individual faces and the average face (as in 'norm-based' models; eg Rhodes et al 1987; Benson and Perrett 1991, 1994) or between individual faces without reference to any norm ('exemplar-based' models; eg Lewis and Johnston 1999).

Facial distortions are not confined to experiments on caricature: they also occur in everyday encounters with faces because of changes in expression or pose, or when viewing a photograph of a face at an angle, because of perspective. More needs to be known about how recognition is achieved under these circumstances, and what kinds of distortion can and cannot be tolerated by the processes underlying recognition. In the experiments to be described in this paper, therefore, we examine the effects on face recognition of transformations of different types. In experiment 1 we investigated the effects of three global linear transformations: shearing, horizontal stretching, and vertical stretching. (Mathematically, these are all examples of affine transformations, which have in common the property that they preserve parallelism between lines in the original images.) These transformations are global in the sense that they apply to the entire image, and linear in the sense that they affect all parts of the image to the same extent. The effects on face recognition were assessed both in comparison with participants' performance with unmanipulated images, and in comparison with inverted faces (also a global linear transformation, but one which leaves the facial features undistorted, in contrast to the previous transformations mentioned). If configural processing relies on simple facial measurements (eg interpupillary distance, length of nose, etc) or their relationships to each other, then shearing and stretching would be expected to have marked effects on face recognition, since they distort these measurements considerably.

To anticipate the subsequent discussion, shearing and vertical and horizontal stretching were found to have little effect on face recognition. In experiment 2, we replicated this finding, using a different procedure to look at the effects of vertical and horizontal stretching in more detail. In experiments 3 and 4 we tried to identify what information was being used to recognise the stretched faces: were participants still able to use configural processing with them, or were they resorting to some other strategy, such as feature-based processing? In experiment 3 we compared performance with partially and globally stretched faces: if participants were using featural information to recognise distorted faces, their performance should be better with the partially stretched faces, since these retain a greater amount of undistorted information about the face than do globally stretched faces. In experiment 4 we used Gaussian blurring as a means of reducing featural information from normal and stretched faces, and thus encouraging participants to use configural processing (Collishaw and Hole 2000). If stretched faces are recognised by means of featural cues, then blurring should impair recognition severely; however, if stretched faces are recognised by means of their configural properties, blurring should have no more effect on their recognition than on the recognition of normal faces.

2 Experiment 1: Effects of linear global distortions on face recognition

2.1 *Method*

2.1.1 *Participants.* There were one hundred participants (twenty per group), mostly undergraduates from the University of Sussex. They were aged between 18 and 65 years. Twenty-three were male.

2.1.2 *Design.* An independent-measures design was used, with each participant allocated to one of five different conditions:

(i) Normal: the faces seen by participants in this group were monochrome 256 grey-level images. Each face was scaled to fit within a grey rectangle measuring approximately

5 cm high and 4 cm wide (4.76 deg by 3.81 deg at the viewing distance of 60 cm used in this experiment).
(ii) Inverted: the faces seen by participants in this group were identical to those shown in the normal condition, except that they were presented upside down.
(iii) Vertical stretch: the faces seen by this group were produced by vertically stretching the normal faces to twice their original height while preserving their original width. This was done by stretching the rectangle within which the normal versions were framed. These faces were therefore approximately 10 cm high and 4 cm wide (9.52 deg by 3.81 deg).
(iv) Horizontal stretch: the faces seen by this group were produced by horizontally stretching the normal faces to twice their original width, while preserving their original height. These faces therefore subtended 7.62 deg horizontally and 4.76 deg vertically.
(v) Shear: the faces seen by this group were produced by applying a shear transformation to the normal faces; this transformed the original rectangular image into a trapezoid whose acute corners were at angles of approximately 45°.
Examples of the stimuli for this experiment are shown in figure 1.

2.1.3 *Stimuli.* Face images were produced in Adobe Photoshop on an Apple Macintosh 6100/60 PowerMac computer, attached to an Agfa Studioscan 2 scanner. The photographs of celebrities were taken from magazines. Photographs of nonfamous people came from our own library of faces of ex-students, plus models in shopping catalogues, etc. Care was taken to ensure that the famous and nonfamous faces did not differ consistently in terms of attributes such as initial image resolution, pose, expression, and age. All faces used were full-face views of clean-shaven adult males.

2.1.4 *Procedure.* Each participant saw a different random sequence of 30 faces and was asked to decide whether each face was famous, as quickly and accurately as possible. Fifteen of the faces were famous, and the remainder were not (see the appendix for a list of the celebrities used). The faces were presented on the screen of an Apple Macintosh G3, by a program written in Supercard. No practice trials or feedback were given.

Each trial took the following form. The computer screen was blank except for the message "press space-bar for the next stimulus". Depressing the space bar caused the screen to clear. After a pause of approximately 1 s, the next face in the test sequence was displayed until the participant made a response. On each trial, the computer measured the time between the onset of presentation of the face, and the participant pressing either of two keys on the keyboard to register their decision ("famous" or "not famous").

2.2 *Results*

Table 1 shows the mean reaction time taken by participants in each experimental condition to make 'correct positive' decisions (ie to recognise correctly a celebrity as being famous). A one-way independent-measures analysis of variance (ANOVA) on these data showed that latencies differed significantly between conditions ($F_{4,95} = 8.61$, $p < 0.0001$). Latencies in the inverted group were significantly longer than those of all other groups (Newman–Keuls test, 0.05 significance level). However, the mean reaction times for the vertical-stretch, horizontal-stretch, and shear groups were not significantly slower than that of the normal group (independent-measures t-tests: $t_{38} = -1.96$, 0.27, and 0.46, respectively, all ns). Different affine transformations of faces (horizontal and vertical stretching, shearing, and inversion) thus produced variable effects on the time taken to recognise familiar faces: vertical stretching, horizontal stretching, and shearing produced no significant impairments in recognition time, whereas inversion markedly slowed recognition, both relative to the normal condition and compared with other affine transformations.

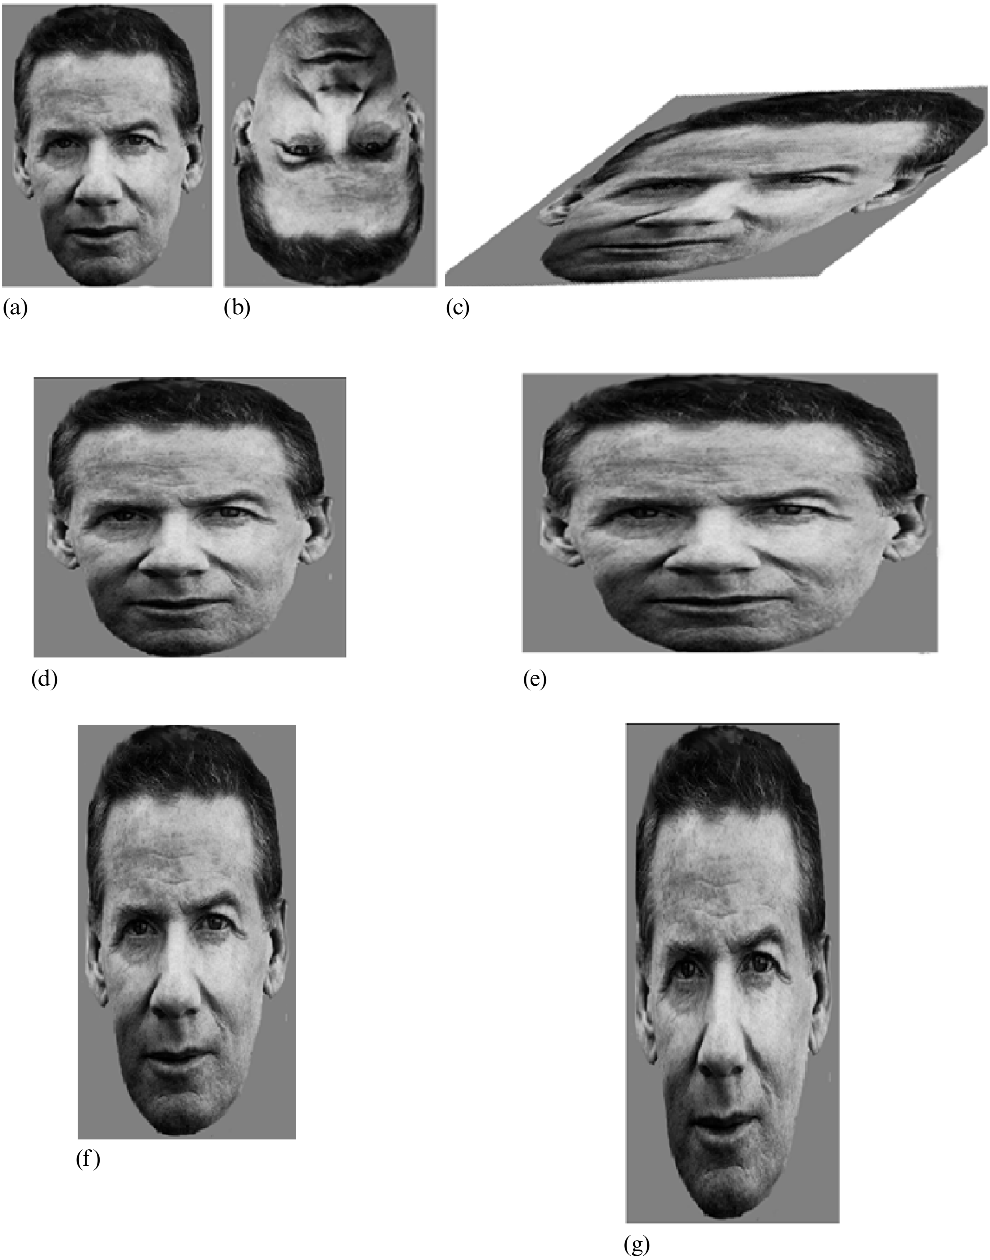

Figure 1. Examples of the manipulations used in experiments 1 and 2: (a) normal, (b) inverted, (c) shear, (d) horizontal stretch (150%), (e) horizontal stretch (200%), (f) vertical stretch (150%), (g) vertical stretch (200%). Manipulations (a), (b), (c), (e), and (g) were used in experiment 1; manipulations (a), (d), (e), (f), and (g) were used in experiment 2. For reproduction purposes, these are shown at relative sizes which are different to those used during the experiments (for full details, see text).

Table 1 also shows the mean number of correct positive decisions made by participants in each condition. A one-way independent-measures ANOVA on these data was highly significant ($F_{4,95} = 13.28$, $p < 0.0001$). The Newman–Keuls test (with a 0.05 significance level) was used to identify which groups were significantly different from each other. Participants in the normal and vertical-stretch conditions were most accurate,

Table 1. Summary of results of experiment 1.

Performance measure	Experimental condition				
	normal	inverted	vertical stretch	horizontal stretch	shear
Mean reaction time (ms) for 'correct positive' responses	1338	1921	1137	1365	1392
SD	(266)	(665)	(371)	(373)	(448)
Mean number correct (total = 15)	11.9	8.1	12.2	10.2	9.7
SD	(2.0)	(2.2)	(1.4)	(2.4)	(2.2)
Number correct as mean % of total famous faces presented	79	54	82	68	64
Mean A'	0.84	0.67	0.84	0.80	0.76
SD	(0.05)	(0.13)	(0.07)	(0.08)	(0.08)

and were not significantly different from each other. Correct recognition rates for these two groups were 79% and 82%, respectively. Performance in the normal condition was significantly more accurate than in the horizontal-stretch and shear groups (68% and 64%, respectively). The inverted condition was significantly poorer than all other groups; participants in the inverted condition recognised only 54% of the total number of celebrities to which they were exposed.

For each participant, data from 'hits' (correctly deciding that a face was famous) and 'false alarms' (incorrectly deciding that a face was famous) were used to compute A', a nonparametric estimate of sensitivity (Rae 1976). An A' value of 0.5 would indicate that performance was at chance levels, while a value of 1 would demonstrate perfect discrimination. As can be seen from table 1, which shows the mean A' value for each group, performance in all groups was well above chance; however, a one-way independent-measures ANOVA showed that the five groups differed significantly in sensitivity ($F_{4,95} = 12.89$, $p < 0.0001$). Independent-means t-tests were used to make pairwise comparisons of each group with the normal group. These confirmed that the mean A' scores for the inverted and shear groups were significantly lower than the mean for the normal group ($t_{38} = 5.33$, $p < 0.001$ and $t_{38} = 3.83$, $p < 0.0001$, respectively). The vertical-stretch and horizontal-stretch groups were not significantly different from the normal group ($t_{38} = 0.09$ and $t_{38} = 1.86$, respectively).

2.3 *Discussion*

Taken together, the latency and accuracy data suggest that the processes involved in face recognition vary in the extent to which they can cope with major global linear distortions of the input face image, depending on the nature of the transformation applied. Vertical stretching has no apparent effect on speed or accuracy, whereas inversion has marked effects. Shearing and horizontal stretching produced no effects on latencies, and only relatively small reductions in accuracy. That vertical stretching should produce no obvious effects on performance is quite surprising, given the marked alterations to facial spatial relationships that it produces. The tolerance of the visual system to vertical stretching, and to some extent to horizontal stretching and shearing too, is in marked contrast to the impact of inversion on recognition performance. Compared with the other conditions, inversion had pronounced effects on latency and accuracy of recognition—yet it was the only one of the experimental manipulations to leave facial spatial relationships completely untouched.

3 Experiment 2: A parametric study of vertical and horizontal stretching

Experiment 2 was carried out in order to investigate the effects of stretching in more detail. In order to increase the sensitivity of the experiment to the detection of any effects of stretching on recognition performance, a wholly within-subjects design was used instead of the between-subjects design of experiment 1. Each participant viewed a series of faces in five different versions—two different levels of horizontal and vertical stretching, plus the normal aspect ratios. Whereas participants in the previous experiment saw only a single block of face images which had received the same manipulation, in this experiment the various manipulations were randomly intermixed—so that on any given trial, a participant was unable to predict which type of manipulation would be experienced. This might be expected to provide a more taxing test of participants' abilities to cope with stretched faces.

3.1 *Method*

3.1.1 *Participants.* There were twenty-one participants, from a variety of occupational backgrounds. They were aged between 16 and 68 years. Ten were female.

3.1.2 *Design.* A wholly within-subjects design was used. The same set of faces was shown to each participant in five different versions which varied in aspect ratio.

(i) Normal: these versions of the faces were similar to those in the previous experiments, except that each face was scaled to measure 11 cm high and approximately 9.5 cm wide (scaling all faces to be the same height meant that they varied slightly in width, owing to individual variations in head shape). This corresponds to 6.28 deg by 5.43 deg, at the 100 cm viewing distance used in the experiment.

(ii) Vertical stretch 150%: these versions of the faces were vertically stretched to one and a half times their original height, while preserving their original width.

(iii) Vertical stretch 200%: these were stretched to twice their original height, while preserving their original width.

(iv) Horizontal stretch 150%: these were stretched horizontally to one and a half times their original width, while preserving their original height.

(v) Horizontal stretch 200%: these were stretched horizontally to twice their original width, while preserving their original height.

3.1.3 *Stimuli.* We used 38 pictures of celebrities and 38 pictures of nonfamous people to generate 380 monochrome face images, with the same criteria (clean-shaven males, roughly similar in age, expression, etc) as in the previous experiment (see the appendix for a list of the celebrities used).

3.1.4 *Procedure.* Each participant saw 380 face images (38 famous and 38 nonfamous faces, each in five different versions). The faces were presented in a different random order for each participant, by a Superlab Pro program on a PC. Subjects' responses were made via an additional mouse, attached to the serial port of the computer. The mouse buttons were labelled as 'F' for famous and 'NF' for nonfamous.

Participants were asked to decide, as quickly but as accurately as possible, whether or not each face was famous. Each face remained on the computer screen until the participant made a response by pressing the appropriate mouse button.

Subjects were given a short practice phase, with 20 of the 380 images (representing two celebrities and two nonfamous people, in all five versions of aspect ratio). Subjects then took 15–20 min to make their responses to the remaining 360 faces in the experiment proper. No feedback was given to subjects throughout the experiment.

3.2 *Results*

3.2.1 *Response times.* Table 2 shows the mean times taken by participants to make 'correct positive' decisions (ie to classify correctly a celebrity as 'famous'), for each of the five experimental conditions. A one-way repeated-measures ANOVA on these data revealed no significant differences between any of the experimental conditions ($F_{4,\,80} = 1.10$, ns): none of the levels of horizontal or vertical stretching appears to have affected participants' reaction times to decide whether or not a face was famous. Two-tailed paired-sample t-tests were used to make direct comparisons between reaction times in the normal condition and those in the two most extreme stretch conditions. No significant differences were found. (Normal versus 200% vertical stretch: $t_{20} = 0.52$; normal versus 200% horizontal stretch: $t_{20} = 0.83$; both ns.)

3.2.2 *Accuracy.* Table 2 shows the mean number of correct positive decisions made in each condition. A one-way repeated-measures ANOVA on these data revealed no significant differences between any of the conditions ($F_{4,\,80} = 1.29$, ns): none of the levels of horizontal or vertical stretching affected participants' accuracy at identifying famous faces.

Table 2 also shows the mean A' values for each condition, calculated as in experiment 1. There were no significant differences between the five conditions, suggesting that vertical and horizontal stretching had no discernible effects on participants' performance ($F_{4,\,80} = 1.13$, ns).

Table 2. Summary of results of experiment 2.

Performance measure	Experimental condition				
	normal	150% vertical stretch	200% vertical stretch	150% horizontal stretch	200% horizontal stretch
Mean reaction time (ms) for 'correct positive' responses	1006	984	988	1023	1038
SD	(302)	(269)	(272)	(285)	(296)
Mean number correct (total = 36)	26.3	27.2	26.6	26.7	26.5
SD	(8.4)	(8.9)	(8.3)	(8.4)	(8.2)
Number correct as mean % of total famous faces presented	73	76	74	74	74
Mean A'	0.77	0.79	0.77	0.78	0.79
SD	(0.27)	(0.27)	(0.27)	(0.26)	(0.25)

3.3 *Discussion*

With a more sensitive design than in the previous experiment, we found further evidence in experiment 2 that our ability to recognise faces is substantially unaffected by global linear stretching. Mean reaction times were much faster than in experiment 1: it is unclear whether this difference arises from differences in experimental design between the two experiments, or from the particular set of faces used. Reaction times may have been faster in experiment 2 because of repetition-priming effects: since each participant saw each face in all of its five different versions, it is possible that initial encounters with a face in one aspect ratio may have primed recognition for the other versions of the same face when they were encountered later within the run of trials. However, the results of this study are in agreement with the principal finding of the previous experiment—that stretching has no appreciable effect on face recognition.

4 Experiment 3: Effects of linear but nonglobal distortions

Experiments 1 and 2 showed that familiar faces can be recognised despite vertical stretching, an affine transformation which produces marked global alterations to the spatial properties of a face. How is recognition achieved with these faces? To recognise a face, there are at least two strategies that could be used: one could process the configural relationship between the features, or one could engage in piecemeal processing of information from the features themselves (Sergent 1984). There is evidence that either of these two processing modes can be used independently of the other as a basis for recognition, and that participants in face-recognition experiments may be quite flexible in which strategy they use (eg Hole 1994; Collishaw and Hole 2000). For example, face recognition survives levels of blurring that are sufficient to remove featural detail (eg Harmon 1973; Sergent 1986; Bachmann 1991; Costen et al 1994), implying that configural information alone may suffice for recognition purposes. Conversely, participants remain able to recognise faces from isolated features, especially the eyes (eg McKelvie 1976; Haig 1985; Lacroce et al 1993), which suggests that configural information is not essential for recognition. Recognition performance might be optimal when both types of processing are used, but either method may suffice to produce a reasonable level of performance (Collishaw and Hole 2000).

In experiments 1 and 2, vertical stretching would appear to pose problems for both configural and piecemeal processing strategies, unless one argues that the featural details of a face are somehow less disrupted by vertical stretching than are the configural properties. This seems a fairly implausible possibility, but it warrants consideration. Experiment 3 contained conditions in which participants saw faces that had been only partially stretched: either the top or the bottom half of the face was stretched vertically, leaving the other half of the face untouched. Compared with globally vertically stretched faces, these faces contain better-quality information about isolated facial features and their local configuration (in the undistorted part of the face). However, they share with the globally vertically stretched faces a significant distortion of the overall configuration of the face.

Performance with these partially stretched faces should provide further insight into the nature of configural processing. If participants are able to ignore the overall configural information from the partially stretched faces and make use of either local featural or local configural information, their performance with the partially stretched faces should be similar or even better than with globally vertically stretched faces. Specifically, one would expect bottom-stretched faces to be recognised almost as well as undistorted faces, because the top parts of these faces are undistorted and numerous studies (eg McKelvie 1976; Davies et al 1977; Haig 1985) have found the upper facial region to be most useful for recognition. The top-stretched faces contain undistorted information about the lower part of the face, which has been shown to be less effective for recognition.

The alternative possibility is that nonglobal linear distortions might affect recognition; if so, this would imply that local featural or configural information was not sufficient for recognition of stretched faces, and that the global configural properties of these faces continued to have a powerful influence on face processing.

4.1 *Method*

4.1.1 *Participants.* There were sixty participants (fifteen per group), mostly students from the University of Sussex. None had taken part in the previous experiments. They were aged between 19 and 46 years. Twenty-four were male.

4.1.2 *Design.* An independent-measures design was used, with each participant allocated to one of four different conditions.
(i) Normal: this condition was identical to the normal condition in experiment 1.
(ii) Global vertical search: the faces seen by this group were produced by vertically stretching the normal faces to twice their original height. This condition was identical to the vertical-stretch condition in experiment 1.
(iii) Top stretch: the faces seen by this group were produced by horizontally dividing each normal face into two halves mid-way down the nose, and then vertically stretching only the top half to twice its original height, leaving the width unaltered.
(iv) Bottom stretch: the faces seen by this group were produced by horizontally dividing each normal face into two halves mid-way down the nose, and then vertically stretching only the bottom half to twice its original height, leaving the width unaltered.
Examples of the stimuli used are shown in figure 2.

4.1.3 *Stimuli and procedure.* The participants' task, and the methods of stimulus preparation and presentation, were the same as those in experiment 1.

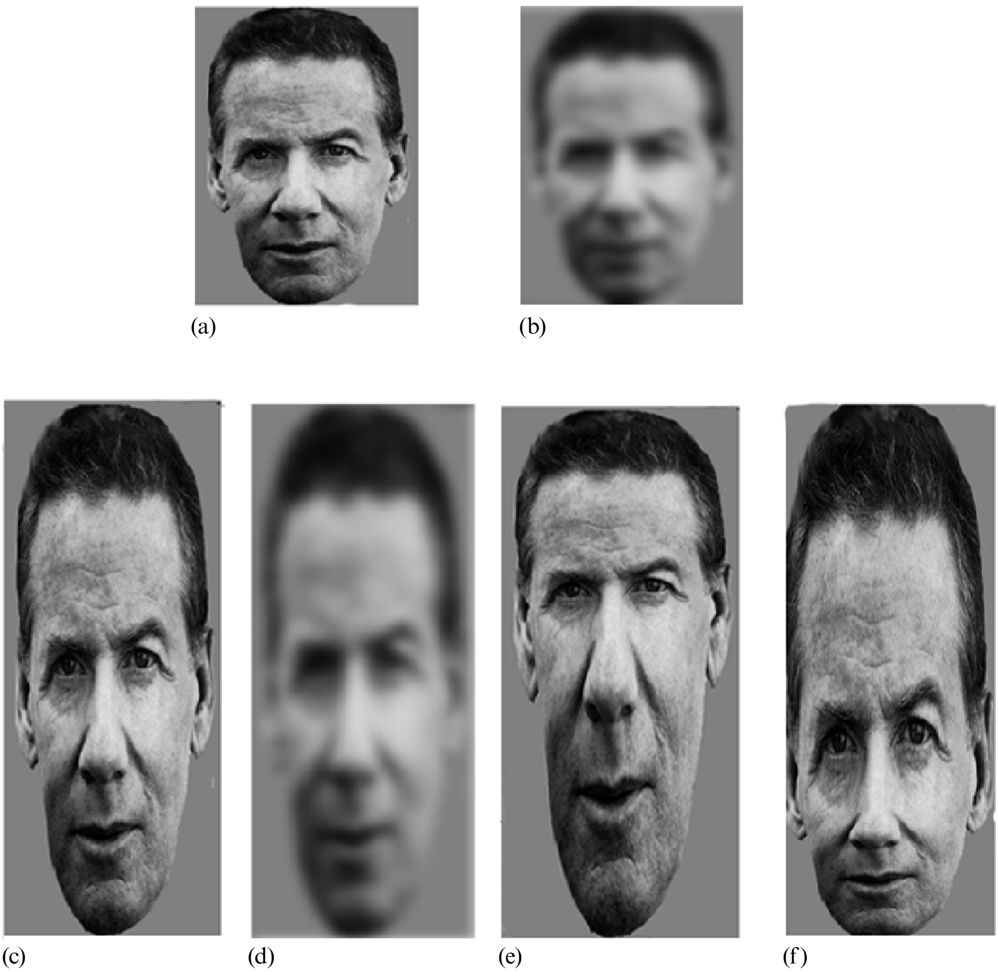

Figure 2. Examples of the manipulations used in experiments 3 and 4: (a) normal, (b) blurred normal, (c) vertical stretch, (d) blurred vertical stretch, (e) bottom stretch, (f) top stretch. Manipulations (a), (c), (e), and (f) were used in experiment 3; manipulations (a) to (d) were used in experiment 4. For reproduction purposes, these are shown at relative sizes which are different to those used during the experiments (for full details, see text).

4.2 *Results*

4.2.1 *Response times.* Table 3 shows the mean reaction time for participants in each experimental condition to make 'correct positive' decisions (ie to recognise correctly a celebrity as being famous). These data suggest that participants in the normal and global-vertical-stretch groups displayed similar mean reaction times, but that participants in the top-stretch group produced longer reaction times and participants in the bottom-stretch group took longer still. A one-way independent-measures ANOVA on the reaction-time data was significant ($F_{3,56} = 7.12$, $p < 0.001$). Newman–Keuls a posteriori tests (with a 0.05 significance level) revealed that the bottom-stretch group was significantly worse than all other groups, and that the remaining three groups were not significantly different from each other. However, an independent-means t-test which specifically compared the normal and top-stretch groups revealed that the mean reaction times for these groups were significantly different ($t_{28} = 2.38$, $p < 0.05$, two-tailed test).

Table 3. Summary of results of experiment 3.

Performance measure	Experimental condition			
	normal	global stretch	top stretch	bottom stretch
Mean reaction time (ms) for 'correct positive' responses	1069	1170	1413	1682
SD	(283)	(412)	(483)	(385)
Mean number correct (total = 15)	13.9	14.1	11.9	11.8
SD	(1.4)	(1.0)	(2.8)	(2.5)
Number correct as mean % of total famous faces presented	93	94	79	79
Mean A'	0.88	0.89	0.81	0.84
SD	(0.04)	(0.04)	(0.13)	(0.08)

4.2.2 *Accuracy.* Table 3 shows the mean number of correct positive decisions made by participants in each condition. A one-way independent-measures ANOVA on the number of correct positive responses was significant ($F_{3,56} = 5.22$, $p < 0.005$). Newman–Keuls tests revealed that participants in the normal and global-vertical-stretch conditions were most accurate, and were not significantly different from each other. Both conditions were significantly more accurate than the top-stretch and bottom-stretch groups. The latter groups were not significantly different from each other. This confirms the impression from the latency data that partially stretched faces are harder to recognise than undistorted or wholly stretched faces.

Table 3 also show the mean A' value for each condition. In all cases performance was well above chance levels. Although performance in all groups appears similar, a one-way independent-measures ANOVA revealed that the four groups differed significantly in sensitivity ($F_{3,56} = 2.82$, $p < 0.05$). Independent-means t-tests were used to identify the source of this result. It appears to arise primarily because the stretch group showed higher mean sensitivity than the top-stretch group ($t_{28} = 2.13$, $p < 0.05$). There was also a marginally significant difference between the normal and top-stretch groups ($t_{28} = 1.95$, $p = 0.06$), although not between the normal and bottom-stretch groups ($t_{28} = 1.65$, ns). The mean A' scores for the normal and stretch groups were not significantly different ($t_{28} = 0.37$, ns).

4.3 *Discussion*

This experiment replicates one of the results of experiments 1 and 2: faces to which global vertical stretching was applied were recognised as well (in terms both of speed and of accuracy), as undistorted faces. Nonglobal distortions (stretching confined to either the top or bottom half of the face) appear to have been coped with less well than similar distortions applied to the entire image. Which manipulation had the greatest effect on recognition (ie top or bottom stretching) varied depending on which measure is considered: stretching the bottom half of the face slowed recognition the most, but stretching the top half resulted in the greatest impairment on the accuracy measures. None of these manipulations disrupted recognition severely, however; all four groups displayed performance which was well above chance.

These results have a number of implications. First, they suggest that participants used configural information from the (globally) vertically stretched faces. If participants were resorting to identifying celebrities by processing of local featural or local configural information from the undistorted half of the face, one might expect performance in the top-stretch and bottom-stretch groups to be superior to that in the global-vertical-stretch condition. This is because partially stretched faces provide the viewer with some degree of uncorrupted information about some of the facial features (the ones in the undistorted half of the face) and their spatial interrelationships. Furthermore, bottom-stretch performance should have been superior to top-stretch performance because the former preserves undistorted information about the upper face region, which previous studies have found is most critical for recognition.

In fact, despite containing a greater amount of directly useful information about the identity of a face, partially distorted faces were harder to recognise than wholly distorted ones. In some respects, these results are similar to Young et al's (1987) findings with chimeric faces comprised of the top half of one face and the bottom half of another: when these halves were juxtaposed, an overall impression of a 'new' face was obtained, and participants found it very difficult perceptually to segregate the top half in order to identify it. The top-stretch and bottom-stretch faces in the present experiment could be regarded as chimeras consisting of an undistorted half-face paired with a half-face which is so distorted that it has effectively become an unfamiliar face.

There are at least two possible explanations for why performance is preserved with the globally stretched faces: it may be because participants can somehow perform the inverse transformation to the image, or it may be because configural processing involves some relationship between the features (some kind of ratio, for example) which is unaffected by vertical stretching. The results of this experiment do not enable us to distinguish between these two possibilities. However, they do tell us that, whichever of these processes is occurring, it is affected by nonglobal distortion. Thus, if participants are achieving recognition by performing the inverse transformation to stretching, then the present results suggest that they cannot apply such a transformation to only part of an image and then match this part with the remaining undistorted part. Alternatively, if participants are recognising the faces by exploiting some configural property (or properties) of a face that are preserved despite global vertical stretching, the present results imply that this property does not survive partial stretching of the faces concerned. Similar results have been reported by Cooper and Wojan (2000), who found that moving both eyes upwards in the head (ie a local translation which disrupts global configural information) disrupted face recognition.

5 Experiment 4: Effects of linear distortion plus blurring

In order to further explore the limits of the configural processing used with the distorted faces in the previous three experiments, experiment 4 was performed. As well as being distorted by global vertical stretching, faces were heavily blurred to reduce the amount of detailed information present. In previous research, the emphasis has been on using blurring as a means to reduce detailed information about the facial features, the implication being that this manipulation leaves the configuration of the face relatively unaffected. Obviously, the effects of blurring will depend on the precise level of blurring used. However, note that blurring not only reduces detailed information about the facial features themselves, but also increases uncertainty about their precise locations within the face—ie it reduces the precision of configural information.

If the precise location of facial features and/or their details are important for face recognition, then this combination of manipulations should impair recognition performance markedly: the distortion should make it difficult to use configural information, while the blurring should make it hard to use detailed featural and textural information, and also detailed information about feature location. If, on the other hand, such information is not critical for recognition to occur, then participants should remain able to recognise these faces.

5.1 *Method*

5.1.1 *Participants.* Sixty-eight participants (seventeen per group) took part, mostly students from the University of Sussex. They were aged between 18 and 47 years. Thirty were male. None had participated in the previous experiments.

5.1.2 *Design.* An independent-measures design was used, with each participant allocated to one of four different conditions.
(i) Normal: this condition was identical to the normal condition in the previous experiments.
(ii) Vertical stretch: this condition was identical to the global-vertical-stretch conditions of the previous experiments.
(iii) Blurred normal: this was identical to the normal condition, except that the faces were heavily blurred with the 'Gaussian blurring tool' in Adobe Photoshop.[2]
(iv) Blurred stretch: this was identical to the vertical-stretch condition, except that after stretching, the faces were heavily blurred (to the same degree as those in the blurred-normal condition) with the 'Gaussian blurring tool' in Adobe Photoshop.
See figure 2 for examples of the stimuli used.

5.1.3 *Stimuli and procedure.* The participants' task, and the methods of stimulus preparation and presentation, were the same as those in the previous experiments.

5.2 *Results*

5.2.1 *Response times.* Table 4 shows the mean times taken by participants in each experimental condition to make 'correct positive' decisions (ie to recognise correctly a celebrity as being famous). A two-way independent-measures ANOVA on the reaction-times data (two levels of blurring, blurred and unblurred, and two levels of stretching,

[2] The 'Gaussian blurring tool' in Adobe Photoshop (version 3.0) was used to limit the range of spatial frequencies available in each image. The algorithm used by this tool applies a weighted average to the pixels within a selected area. In this case the blurring tool was set to a value of 10 pixels. As can be seen from figures 2b and 2d, this procedure was effective in reducing fine-detail information in the faces to which it was applied. As a further check on what Photoshop actually did to the image when this routine was used, a Fourier analysis was performed on a 512×512 pixel image, after the Gaussian blurring routine had been applied to it. Although the resultant image did not conform perfectly to what one would expect of a Gaussian-filtered image, the filtering routine in Photoshop did produce a marked and reasonably progressive attenuation of higher spatial frequencies (above 225 cycles per image-width or so).

Table 4. Summary of results of experiment 4.

Performance measure	Experimental condition			
	normal	vertical stretch	blurred normal	blurred stretch
Mean reaction time (ms) for 'correct positive' responses	1123	1321	1435	1501
SD	(235)	(291)	(319)	(238)
Mean number correct (total = 15)	10.9	10.7	7.6	8.3
SD	(1.1)	(2.5)	(2.3)	(2.4)
Number correct as mean % of total famous faces presented	73	71	51	55
Mean A'	0.85	0.79	0.69	0.69
SD	(0.05)	(0.11)	(0.11)	(0.11)

stretched and unstretched) revealed a significant main effect of blurring ($F_{1,64} = 13.80$, $p < 0.001$), and a marginally significant effect of stretching ($F_{1,64} = 3.98$, $p = 0.05$). As can be seen from table 4, reaction times were longer for participants in the blurred-normal and blurred-stretch conditions. The interaction between blurring and stretching was not significant ($F_{1,64} = 0.99$, ns). Independent-means t-tests revealed significant differences between the normal and blurred-normal conditions ($t_{32} = 3.25$, $p < 0.005$), but the difference between the vertical-stretch and blurred-stretch groups was only marginally significant ($t_{32} = 1.98$, $p < 0.06$). Participants in the vertical-stretch condition were significantly slower than in the normal condition ($t_{32} = 2.18$, $p < 0.05$), but the blurred-stretch and blurred-normal conditions did not significantly differ ($t_{32} = 0.69$, ns). Blurring appears to have produced similar effects on reaction times in the case of both normal and stretched faces.

5.2.2 *Accuracy.* Table 4 shows the mean number of correct positive decisions made by participants in each condition. A two-way independent-measures ANOVA on the number of correct responses (two levels of blurring, blurred and unblurred, and two levels of stretching, stretched and unstretched) revealed a significant main effect of blurring ($F_{1,64} = 28.46$, $p < 0.0001$), no significant effect of stretching ($F_{1,64} = 0.26$, ns), and no significant interaction between blurring and stretching ($F_{1,64} = 0.91$, ns). It is clear that accuracy was similar for both stretched and normal faces, and that blurring decreased accuracy in a similar way for both types of stimulus. Note that accuracy in the blurred-stretch condition was low (55% of celebrities identified), but similar to that in the blurred-normal condition (51%).

Table 4 also shows the mean A' value for each permutation of blurring and stretching. A two-way ANOVA on these data confirmed the impression obtained from the 'correct positive' data: there was a significant effect on sensitivity of blurring ($F_{1,64} = 26.69$, $p < 0.0001$), but no effect of stretching, either as a main effect ($F_{1,64} = 1.59$, ns) or in interaction with blurring ($F_{1,64} = 1.21$, ns).

5.3 *Discussion*

Once again, vertical stretching had little effect on recognition performance, although in this experiment (unlike the previous three) reaction times for participants viewing vertically stretched faces were significantly slower than those for participants viewing unmanipulated faces. The addition of blurring significantly impaired recognition of vertically stretched faces, but no more than it did for normal faces. Since recognition of blurred-stretch faces was still possible, this suggests that participants were able to use whatever information remained available in these images—namely the approximate global configuration of the facial features, albeit in a heavily distorted form. If participants had been using detailed information in order to recognise stretched faces (ie texture, feature detail, or precise configural detail), then one might have expected the celebrity faces in the blurred-stretch condition to be unrecognisable, or at least to be significantly harder to recognise than in their unblurred vertical-stretch versions.

6 General discussion

Different kinds of global linear (affine) transformation (horizontal and vertical stretching, shearing, and inversion) produced markedly different effects on face recognition. As in previous studies (eg Yin 1969, 1970; Valentine 1988) inversion severely impaired recognition. Shearing produced moderate impairments. Horizontal stretching had little effect on recognition in experiment 1, and none at all in experiment 2. Vertical stretching had no significant effect on recognition in four separate experiments. Global linear transformations seem to affect recognition remarkably little, considering the extent to which they distort the spatial relationships within a face, which are presumed to be critical for recognition. Recognition was impaired, however, when vertical stretching was applied nonglobally (ie to only half of the face).

There are a number of reasons why recognition might be relatively unaffected by global stretching. First, it might be that the types of configural information which were distorted in the vertical-stretch, horizontal-stretch, and shear conditions of experiment 1 are not essential for recognition; participants might be using featural information instead, for example. It is likely that we can employ different strategies in order to recognise faces, and that we can switch between them in response to task demands (Hole 1994). Collishaw and Hole (2000) have suggested that there are at least two routes to recognition—via either configural or featural information—and that providing one of these routes remains available, recognition performance remains above chance levels.

The idea that our participants switched to a strategy of recognising faces by means of featural cues thus needs to be considered. However, there are a number of reasons why this is unlikely to be the explanation for our results. The most obvious is that one would have to account for why featural information should be more resistant than configural information to the effects of distortion. We know of no data on this point, but it seems implausible that this should be so. In any case, the data from the present experiments offer no support for such an explanation. First, if participants were relying on featural information, the performance of the blurred-stretch group in experiment 4 should have been much worse than it actually was: in addition to the disruption of configural information in the faces by stretching, detailed featural information was also unavailable to these participants owing to the Gaussian blurring. Collishaw and Hole (2000) have shown that disruption of configural *and* featural information (by a combination of inversion and blurring) reduces face-recognition performance to chance levels, which was clearly not the case in experiment 4. Second, if participants were using featural information, performance with the part-distorted faces of experiment 3 should have been superior to performance with the globally vertically stretched faces, since the former contained a greater amount of accurate information about the facial features than did the latter. In fact, the relatively poor performance with the part-distorted faces suggests that participants were trying to process their global configural properties: the effects of the part-distorted faces were reminiscent of those found with the chimeric faces investigated by Young et al (1987), Hole (1994), and Hole et al (1999): the 'erroneous' configuration produced by the part-distorted faces disrupted recognition, nullifying any benefit that might be obtained from an analysis of individual features in the undistorted part of the face. Data consistent with this interpretation were reported by Cooper and Wojan (2000). They presented participants with famous faces which had been distorted by shifting either one or both eyes upwards. Although these faces had identical featural information, participants were better at recognising faces when only one eye had been shifted than when both eyes were moved, even though the latter stimuli appeared more 'face-like'.

Overall, our results suggest that participants were recognising the distorted faces by using whatever configural information was retained. Given that previous studies have demonstrated that we are highly sensitive to the precise location of facial features (eg Haig 1984; Hosie et al 1988; Kemp et al 1990), why is configural processing unaffected by gross distortions of facial spatial properties? We can think of at least three possible explanations.

One is that we might be able to apply the inverse transformation to the distorted facial image, in order to recover an undistorted image for comparison with our internal representations of familiar faces. If so, one might expect such a transformation to take some time: however, in experiment 2 (in which participants saw faces which had been subjected to different amounts of horizontal and vertical stretching) it took no longer to match a distorted image to memory than an undistorted one. Note that, in this experiment, the different versions of face were randomly intermixed, so that a participant would have no idea of the type and extent of distortion that would appear from one trial to the next. If any time costs of performing an inverse transformation were likely to manifest themselves, it would have been in this experiment, yet none was found.

In itself, this is not a strong argument, because the transformation might occur too rapidly for us to measure. However, performing a compensatory transformation in order to recognise a face also poses a computation problem, since it presupposes that the visual system 'knows' what proportions the face should be transformed to. The compensatory system would have to possess rules for knowing when to stop compensating, This is not an insuperable difficulty, however. Firstly, the compensation process could be guided by knowledge of the proportions of the average face. This might be one advantage of having a 'prototypical' or 'norm' face against which to measure any particular face that is encountered, and it would be one reason for engaging in an initial 'holistic' preprocessing which identifies that the input is indeed a face rather than some other kind of object (see Hole et al 1999). Moses et al (1996) have suggested that in order to generalise from known views of a familiar face to novel views of it (eg in order to compensate for changes in viewpoint and illumination) the visual system could exploit 'class-based' knowledge about the general characteristics that all faces share—eg bilateral symmetry, overall configuration of features, 3-D shape, etc. Our distortion manipulations can be viewed as a more extreme example of the same problem, with possibly the same solution.

An alternative solution to the compensation problem is to transform stored faces to match the input stimulus, rather than vice versa. Biederman and his associates (eg Fiser et al 1996; Biederman and Kalocsai 1997; Biederman et al 1999) have examined the performance of the von der Malsburg face-recognition system (Lades et al 1993; Wiscott et al 1997). This makes use of relatively simple rules and yet has been able to simulate many of the features of human face recognition (Kalocsai et al 1994).

In this system an input face is analysed by an array of columns of Gabor spatial filters that vary in spatial frequency and orientation. Each column deals with a fairly circumscribed region of the input image. The values from these filters for the input image are compared with the corresponding values stored in the system for each face 'known' to it. 'Recognition' is based on a match between the values of the Gabor columns for a stored image (ie a previously seen face) and the values for the corresponding columns in an input image. To cope with changes in orientation, position, expression, and lighting, matching is done after deforming the stored image: the position of each filter column in the stored image is allowed to change until it best matches a column in the input image. To help eliminate false matches, a constraint on this process is that the topological relationship between the columns must be preserved (hence the term 'deformable template'). The stored face which requires the least deformation of its array of columns in order to achieve a match is taken to be the system's 'decision' about the identity of the input image.

A more recent version of the model (Wiscott et al 1997) centres each Gabor column on a facial landmark or 'fiducial point', such as the corner of an eye. Both versions of the system have been shown to be highly accurate at recognising faces despite moderate changes in position, orientation, and expression. A preliminary test of the model with stretched faces suggests that its peformance is relatively unaffected (Biederman, personal communication).

In discussing compensatory transformations, it should be acknowledged that an ability to cope with distorted faces might reflect properties of picture perception, rather than face recognition per se; the present results might stem from a generalised ability to compensate for image distortions. As Busey et al (1990) point out, we rarely experience veridical 2-D pictorial representations, since a 2-D image is undistorted only if viewed from a single viewpoint, orthogonal to the picture plane and at a specific distance from the picture. The present experiments do not enable us to decide between face-specific and more generalised compensatory abilities. A potentially fruitful direction for future research might be to compare the effects of affine distortions on recognition of faces, non-face objects which share a common basic organisation, and non-face objects which do not share a basic configuration.

A third means by which recognition of the distorted faces in our experiments might be accomplished is if the configural information used is of a form that is unaffected by our manipulations—in other words, if the facial characteristics that are being used remain invariant under transformation. For example, the ratio between head width and interpupillary distance remains invariant under vertical or horizontal stretching. We are not claiming that face recognition is indeed based on such a simple ratio. If it were, one might expect that the partially stretched faces of experiment 3 (in which this ratio was preserved) would be recognised as well as normal and globally vertically stretched faces. In practice, there were some indications that partial stretching impaired performance. (The results of experiment 3 are not fatal to this argument, however. The differences in performance between the partial-stretch conditions and the global-stretch and normal conditions were relatively small. Also it could be argued that the ratio of interpupillary distance to head width is used under normal conditions, but that under our experimental conditions, the 'erroneous' configural information produced by partial stretching in our experiment interfered with the use of this measurement.)

Computationally, there are a number of reasons why it would make sense to use some spatial property involving eye separation as a basis for face recognition. First, the eyes are clearly defined (especially in terms of the contrast between sclera and iris), which makes them easy to localise within the facial configuration (Bachmann 1991). Second, some facial features are too mobile to provide a reliable basis for recognition: for example, the width and size of the mouth are likely to change from moment to moment as a consequence of talking and expression changes. In marked contrast to this, eye separation is fixed: the only changes which occur are as a consequence of perspective changes, and it may be relatively easy to compensate for these, at least within the limits imposed by change from a full-face to a three-quarters view. Reliance on configural information from the upper face region would be consistent with behavioural data suggesting that people are able to recognise faces solely on the basis of information from around the region of the eyes (McKelvie 1976; Lacroce et al 1993).

These three explanations for why stretching has so little effect on face recognition are not mutually exclusive and, at present, one can only speculate on their plausibility as explanations for the current findings. One way to distinguish between the two categories of 'normalisation' explanation (ie between theories suggesting that an input face is transformed before being matched to stored faces, and theories suggesting that the opposite occurs) might be to look at the effects of familiarity on the recognition of distorted faces. Since unfamiliar faces will lack stored representations, there should be no template that can be deformed to match them (other than perhaps some 'norm' or prototypical face). A deformable-template theorist might therefore predict that unfamiliar-face recognition will be more affected by distortion than familiar-face recognition. If the input face is being normalised to match stored (veridical) faces, then there is no reason to suppose any interaction between familiarity and distortion. To evaluate the plausibility of the 'ratio primitives' notion, first one would need to demonstrate that sufficient ratios exist and can act as a basis for accurate face recognition; and, second, one would have to show that disrupting these ratios significantly impaired recognition.

At the very least, the present experiments strongly suggest that face recognition is unlikely to be based on a collection of simple measurements between various facial features. These results therefore have implications for 'multidimensional face-space' models of face recognition (Valentine 1988) which are ultimately based on representing faces in such a way; in order to cope with stretching and other distortions, these models will need to be modified, either by including a preliminary compensation or normalisation process, or by being restricted to measurements (such as interpupillary ratio and head-width ratio) which are inherently resistant to facial distortions.

Acknowledgments. We are grateful to Irving Biederman and our anonymous reviewers for their comments on earlier drafts of this paper.

References

Bachmann T, 1991 "Identification of spatially quantised tachistoscopic images of faces: how many pixels does it take to carry identity?" *European Journal of Cognitive Psychology* **3** 87 – 103

Benson P J, Perrett D I, 1991 "Perception and recognition of photographic quality facial caricatures: implications for the recognition of natural images" *European Journal of Cognitive Psychology* **3** 105 – 135

Benson P J, Perrett D I, 1994 "Visual processing of facial distinctiveness" *Perception* **23** 75 – 93

Biederman I, Kalocsai P, 1997 "Neurocomputational bases of object and face recognition" *Philosophical Transactions of the Royal Society of London, Series B* **352** 1203 – 1219

Biederman I, Subramaniam S, Bar M, Kalocsai P, Fiser J, 1999 "Subordinate-level object classification re-examined" *Psychological Research* **62** 131 – 153

Brooks B E, Rosielle L J, Cooper E E, 2002 "The priming of face recognition after metric transformations" *Perception* **31** 297 – 313

Bruce V, 1994 "Stability from variation: the case of face recognition" *Quarterly Journal of Experimental Psychology* **47A** 5 – 28

Bruce V, Burton A M, Craw I, 1992 "Modelling face recognition" *Philosophical Transactions of the Royal Society of London, Series B* **335** 121 – 128

Bruce V, Doyle T, Dench N, Burton M, 1991 "Remembering facial configurations" *Cognition* **38** 109 – 144

Busey T A, Brady N P, Cutting J E, 1990 "Compensation is unnecessary for the perception of faces in slanted pictures" *Perception and Psychophysics* **48** 1 – 11

Collishaw S M, Hole G J, 2000 "Featural and configural processes in the recognition of faces of different familiarity" *Perception* **29** 893 – 909

Cooper E E, Wojan T J, 2000 "Differences in the coding of spatial relations in face identification and basic-level object recognition" *Journal of Experimental Psychology: Learning, Memory, and Cognition* **26** 470 – 488

Costen N P, Parker D M, Craw I, 1994 "Spatial content and spatial quantisation effects in face recognition" *Perception* **23** 129 – 146

Davies G, Ellis H, Shepherd J, 1977 "Cue saliency in faces as assessed by the 'Photofit' technique" *Perception* **6** 263 – 269

Fiser J, Biederman I, Cooper E E, 1996 "To what extent can matching algorithms based on direct outputs of spatial filters account for human object recognition?" *Spatial Vision* **10** 237 – 271
Galper R E, 1970 "Recognition of faces in photographic negative" *Psychonomic Science* **19** 207 – 208
George P A, Hole G J, 2000 "The role of spatial and surface cues in the age-processing of unfamiliar faces" *Visual Cognition* **7** 485 – 510
Haig N D, 1984 "The effect of feature displacement on face recognition" *Perception* **13** 505 – 512
Haig N D, 1985 "How faces differ—a new comparative technique" *Perception* **14** 601 – 615
Harmon L D, 1973 "The recognition of faces" *Scientific American* **227** 71 – 82
Hole G J, 1994 "Configurational factors in the perception of unfamiliar faces" *Perception* **23** 65 – 74
Hole G J, George P A, Dunsmore V, 1999 "Evidence for holistic processing of faces viewed as photographic negatives" *Perception* **28** 341 – 359
Hosie J A, Ellis H D, Haig N D, 1988 "The effect of feature displacement on the perception of well-known faces" *Perception* **17** 461 – 474
Johnston A, Hill H, Carman N, 1992 "Recognising faces: effects of lighting direction, inversion, and brightness reversal" *Perception* **21** 365 – 375
Kemp R, McManus C, Pigott T, 1990 "Sensitivity to the displacement of facial features in negative and inverted images" *Perception* **19** 531 – 543
Kemp R, Pike G, White P, Musselman A, 1996 "Perception and recognition of normal and negative faces: the role of shape from shading and pigmentation cues" *Perception* **25** 37 – 52
Lacroce M S, Brosgole L, Stanford R G, 1993 "Discriminability of different parts of faces" *Bulletin of the Psychonomic Society* **31** 329 – 331
Lades M, Vortbrüggen J C, Buhmann J, Lange J, Malsburg C von der, Würtz R P, Konen W, 1993 "Distortion invariant object recognition in the dynamic link architecture" *IEEE Transactions on Computers* **42** 300 – 311
Lewis M B, Johnston R A, 1998 "Understanding caricatures of faces" *Quarterly Journal of Experimental Psychology* **50A** 321 – 346
Lewis M B, Johnston R A, 1999 "A unified account of the effects of caricaturing faces" *Visual Cognition* **6** 1 – 41
McKelvie S J, 1976 "The role of the eyes and mouth in the memory of a face" *American Journal of Psychology* **89** 311 – 323
Moses Y, Ullman S, Edelman S, 1996 "Generalization to novel images in upright and inverted faces" *Perception* **25** 443 – 481
Rae G, 1976 "A non-parametric measure of recognition performance" *Perceptual & Motor Skills* **42** 98
Rhodes G, Brake K, Atkinson A, 1993 "What's lost in inverted faces?" *Cognition* **47** 25 – 57
Rhodes G, Brennan S, Carey S, 1987 "Identification and rating of caricatures: implications for mental representations of faces" *Cognitive Psychology* **19** 473 – 497
Rhodes G, Tremewan T, 1994 "Understanding face recognition: caricature effects, inversion and the homogeneity problem" *Visual Cognition* **1** 275 – 311
Searcy J H, Bartlett J C, 1996 "Inversion and processing of component and spatial-relational information in faces" *Journal of Experimental Psychology: Human Perception and Performance* **22** 904 – 915
Sergent J, 1984 "An investigation into component and configural processes underlying face recognition" *British Journal of Psychology* **75** 221 – 242
Sergent J, 1986 "Microgenesis of face perception", in *Aspects of Face Processing* Eds H D Ellis, M A Jeeves, F Newcombe, A Young (Dordrecht: Martinus Nijhoff) pp 17 – 33
Troje N F, Bülthoff H H, 1996 "Face recognition under varying poses: the role of texture and shape" *Vision Research* **36** 1761 – 1771
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, 1991 "A unified account of the effects of distinctiveness, inversion and race in face recognition" *Quarterly Journal of Experimental Psychology* **43A** 161 – 204
Wiscott L, Fellous J-M, Krüger N, Malsburg C von der, 1997 "Face recognition by elastic graph matching" *IEEE Pattern Recognition and Machine Intelligence* **19** 775 – 779
Yin R K, 1969 "Looking at upside-down faces" *Journal of Experimental Psychology* **81** 141 – 145
Young A W, Hellawell D, Hay D C, 1987 "Configurational information in face perception" *Perception* **16** 747 – 759

APPENDIX

List of celebrities whose faces were used in these experiments

Keanu Reeves
Jason Donovan
Barry Norman
Jimmy Nail
Kenneth Roache ("Ken Barlow" in *Coronation Street*)
Des O'Connor
Jeremy Clarkson
Jim Davidson
David Gower
Harrison Ford
Paul Newman
Sylvester Stallone
Michael Barrymore
Michael French ("David Wicks" in *EastEnders*)
Vince Earl ("Ron Dixon" in *Brookside*)