
Things are looking up: Differential decline in face recognition following pitch and yaw rotation

Simone K Favelle, Stephen Palmisano, Ryan T Maloney

Department of Psychology, University of Wollongong, Northfields Avenue, Wollongong 2522, NSW, Australia; e-mail: simone_favelle@uow.edu.au

Received 6 June 2006, in revised form 6 February 2007; published online 10 October 2007

Abstract. Previous research into the effects of viewpoint change on face recognition has typically dealt with rotations around the head's vertical axis (yaw). Another common, although less studied, source of viewpoint variation in faces is rotation around the head's horizontal pitch axis (pitch). In the current study we used both a sequential matching task and an old/new recognition task to examine the effect of viewpoint change following rotation about both pitch and yaw axes on human face recognition. The results of both tasks showed that recognition performance was better for faces rotated about yaw compared to pitch. Further, recognition performance for faces rotated upwards on the pitch axis was better than for faces rotated downwards. Thus, equivalent angular rotations about pitch and yaw do not produce equivalent viewpoint-dependent declines in recognition performance.

1 Introduction

The retinal projection of a face at one viewpoint is often drastically different from its projection at another viewpoint. In fact, the images of two different faces viewed front-on (full-face) will resemble each other more than the same face viewed front-on and in profile. Most empirical evidence suggests that faces are represented in a viewpoint-dependent manner. Behavioural studies have shown that recognition performance declines with increases in the angular deviation between the learned (exposed, stored, or encoded) view of a face and a novel (test) view (Hill et al 1997; O'Toole et al 1998; Schyns and Bühlhoff 1994; Troje and Bühlhoff 1996, 1998; Troje and Kersten 1999; Valentin et al 1997). Consistent with this observation, neurophysiological studies have found neurons in the superior temporal sulcus (STS) of the macaque that respond selectively to images of faces (monkey and human) at specific viewpoints, with the frequency of cell discharge decreasing monotonically with rotation away from these views (Desimone et al 1984; Perrett et al 1985, 1987, 1992).

The above-mentioned studies examined only the effects of rotation in depth, or around the vertical axis, on face recognition (which simulates a 'yaw' head movement relative to a stationary observer). Also common, though less studied, are orthogonal rotations of faces around the horizontal pitch axis (which simulate a 'pitch' head movement relative to a stationary observer; examples of both types of rotation are given in figure 1). Some authors have examined the effect of pitch rotations (also referred to as 'head tilt') on the perception of facial emotion (Mignault and Chaudhuri 2003; Otta et al 1994) including the Noh mask effect (Lyons et al 2000), gender categorisation (Campbell et al 1996, 1999), and facial attractiveness (Geldart et al 1999). Yet, despite our everyday experience with faces rotated around the pitch axis, very few studies have been primarily concerned with the effects of this type of rotation on face recognition.

To our knowledge, only two studies have specifically investigated the effect of pitch rotations on human face recognition. The first, by Wallraven et al (2002), aimed to explore the viewpoint-dependent nature of face recognition by attempting to replicate the results of an object recognition study with paperclip and amoeboid blobs (Bühlhoff and Edelman 1992). Wallraven et al (2002) used a recognition task in which the learned stimuli

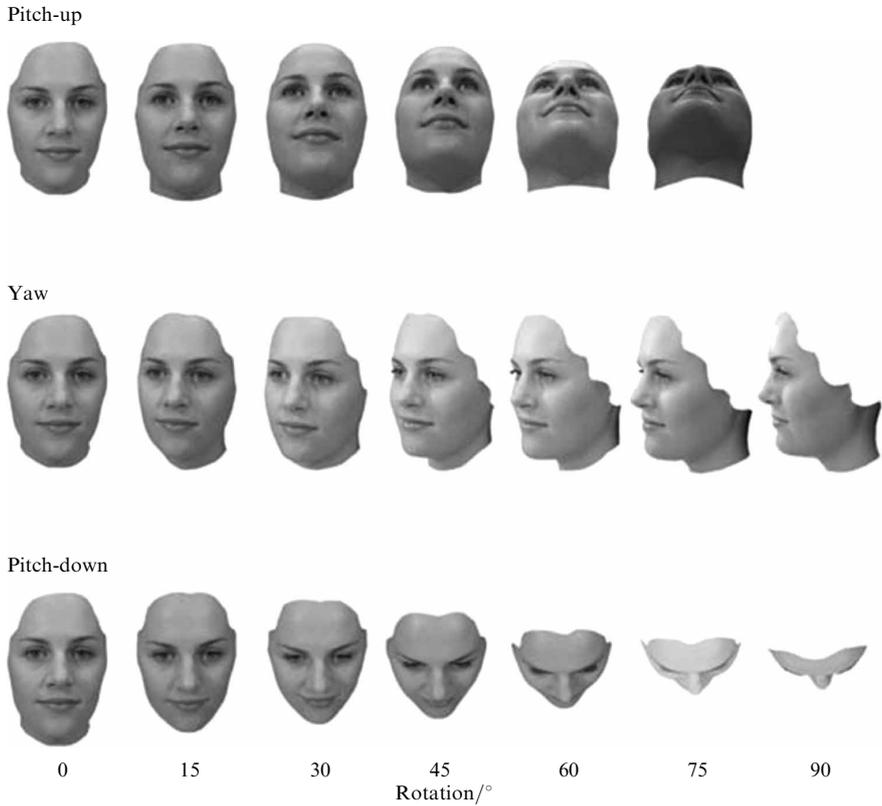


Figure 1. Examples of the stimuli generated. 0° is a full-face view and is the same for each axis. The top row shows images simulating pitch-up rotations (experiment 1a), the centre row shows yaw rotations, and pitch-down rotations (experiment 1b) are shown along the bottom row.

were ten laser-scanned faces rotated dynamically 5° to the left and 5° to the right of (i) a front-on view and (ii) a view 60° to the right of front-on (about yaw). During the test phase, static faces were shown at 15° increments around the orthogonal axis, in this case pitch (0° to 45° upwards or downwards) and yaw (0° to 60° leftwards or rightwards). While face recognition declined in a viewpoint-dependent manner following both yaw and pitch rotations, this rate of decline was faster for pitch rotations than for yaw. No difference in the rate of decline was found for pitch-up and pitch-down rotations. Wallraven et al used their results to distinguish between three different explanations of viewpoint-dependent recognition: (i) alignment of a 3-D model, (ii) linear combination of 2-D views, and (iii) view interpolation. Their finding of a differential effect of rotation direction does not fit with the alignment of a 3-D model explanation, and the linear combination of 2-D views approach cannot explain why performance was better for views between the 0° and 60° yaw training views than for views rotated about yaw in the opposite direction. Wallraven et al's results replicated Bülthoff and Edelman's (1992) earlier findings with unfamiliar objects in that recognition performance depended on the direction of rotation. Thus, they concluded that their results are best understood with a view-interpolation model of object recognition.

More recently, Liu et al (2005, experiment 3) investigated face recognition in a simultaneous-matching task using 3-D face models containing shape-from-shading cues. Participants matched a target face (which remained visible) to one of eight test faces. The test faces consisted of the target and seven distractor faces and were always presented at a different view to the target (pitch-up, pitch-down, and yaw rotations up to 40°).

Liu et al found that face recognition declined in a viewpoint-dependent fashion for both pitch and yaw rotations of the test face. However, in apparent conflict with the findings of Wallraven et al (2002), 40° pitch-down rotations were significantly more difficult to recognise than either 40° pitch-up or 40° yaw rotations (Liu et al 2005, experiment 3). In summary, Wallraven et al (2002) found differences in recognition performance between yaw and pitch (with pitch-up and pitch-down rotations producing equivalent recognition). Liu et al (2005), on the other hand, found similar recognition performance for all types of rotation except at an angular deviation of 40°. At 40° pitch-down rotation, performance was markedly poorer compared to pitch-up and yaw rotations (which did not differ).

In general, both the Wallraven et al (2002) and Liu et al (2005) studies are in line with a viewpoint-dependent account of face recognition; however, both studies found that face recognition or matching performance depended on the direction of rotation. In particular, the results seem to suggest that recognising faces rotated around the pitch axis can at times be more demanding than recognising faces rotated around the yaw axis. However, Wallraven et al were investigating recognition performance for faces rotated along orthogonal axes; it just happened to be the case that the orthogonal axis in their study was the pitch axis. In their experiment, faces were learned at 0° and an additional orientation along the yaw axis (60°). A possible explanation of Wallraven et al's finding that performance for pitch rotations was worse than for yaw rotations is that learning faces rotated about yaw enhanced performance around that axis.

Although neither Wallraven et al (2002) nor Liu et al (2005) were primarily concerned with the effects of rotating faces about the pitch axis, their results highlight the need for a focused investigation of the effect of pitch rotation on face recognition. In addition to the above-mentioned issue with the learning phase in Wallraven et al's (2002) study, the differences between the results of Wallraven et al and Liu et al might have arisen for a number of reasons. First, Wallraven et al (2002) used a recognition procedure that required a memory load, whereas Liu et al (2005) used a simultaneous-matching procedure in which memory load was minimal (or arguably non-existent) and all visual information required to complete the task was available concurrently. The possibility of direct visual comparison of the two faces in Liu et al's task might explain why performances along the different directions of rotation were more similar than those in Wallraven et al. Second, there were differences in the stimuli: Liu et al (2005, experiment 3) used shaded face models with homogeneous texture, whereas Wallraven et al (2002) used laser-scanned face models with realistic skin texture. However, different strategies are needed to recover structural information from shaded and textured faces (Troje and Bühlhoff 1996). For example, the edge and sharp-contrast information of localisable features provided by realistic skin texture is not available in shaded faces.

The aim of the current study was to directly investigate the effect of both pitch and yaw rotation on face recognition. In experiments 1a and 1b, we used a sequential matching task (since it could be considered the simplest form of recognition task—ie involving memory; see Ellis and Allport 1986). In experiment 2, we used an old/new recognition task similar to that used by Wallraven et al (2002), to tap into longer-term representations used for recognition. The emphasis in both tasks was on comparing faces rotated to varying degrees about the yaw and pitch axes to a full-face view for recognition.⁽¹⁾ The full-face view was used as the comparison point for recognition,

⁽¹⁾A fully factorial design in which each rotated face image is compared with every other face image was not employed because of the difficulties of equating angular deviations. For example, a 15° angular deviation between 0° and 15° is not necessarily equivalent to a 15° angular deviation between 60° and 75°. The current method ensures that a 15° angular deviation, for example, is always a comparison between 0° and 15° images.

since it is the only view common to both the pitch and yaw axes. In the sequential matching task, the aim in each trial was to match a rotated face to a full-face view. In the old/new recognition task, a full-face view was learned in the initial encoding phase and rotated faces were shown at test for recognition. The stimuli were high-quality digital images of real faces viewed at 15° increments around pitch (0° to 75° or 90°) and yaw (0° to 90°) axes. View-invariant properties such as skin blemishes and hair were removed from the images of the faces, since we were primarily interested in the matching of 3-D shape. All faces were lit primarily from above as in the Wallraven et al (2002) and Liu et al (2005, experiment 3) studies. A wide range of views in pitch was used since the consequences to face recognition of rotations of more than 45° along this axis are unclear. Static face images were used, since 'animating' the encoding sequences (as Wallraven et al did) has been shown to enhance recognition (Busey and Zaki 2004; Hill et al 1997; Lander and Chuang 2005). In addition, all views were displayed an equal number of times during the encoding phase. Because there is evidence that the visual system treats the left-hand and right-hand sides of faces more or less equivalently (Busey and Zaki 2004; Hill et al 1997; Troje and Bühlhoff 1996, 1998), only one direction of yaw rotation (the one simulating a face turning to the right) was examined.

In experiment 1a, face-recognition performance in a sequential matching task following rightwards yaw rotation was compared with performance following upwards pitch (pitch-up) rotation. In experiment 1b, the face recognition performance in a sequential matching task following rightwards yaw rotation was compared with performance following downwards pitch (pitch-down) rotation.⁽²⁾ In experiment 2, face recognition following rightwards yaw, pitch-up, and pitch-down rotations was compared in an old/new recognition task. On the basis of the findings of previous research, we predicted that, for both our short-term and longer-term memory tasks, face recognition would be viewpoint-dependent for faces rotated around both axes. Using an old/new memory task, Wallraven et al (2002) found worse recognition performance for 45° pitch-up rotations (relative to 45° yaw rotations), whereas Liu et al (2005) found no difference between pitch-up and yaw rotations using a simultaneous-matching task (involving minimal memory load). On the one hand, it might be expected that recognition performance for pitch-up rotation would be equivalent to yaw with the sequential matching task in experiment 1a (on the basis of Liu et al's findings and the similarities between sequential and simultaneous matching tasks), but worse than yaw with the old/new recognition task in experiment 2 (on the basis of Wallraven et al's findings with such a task). However, given that both tasks used in the current study require some level of recognition (memory rather than simply perception), it might be the case that both experiment 1a and experiment 2 will show worse performance for pitch-up than yaw rotations. Both Liu et al (2005) and Wallraven et al (2002) found worse recognition performance for 40° and 45° pitch-down rotations (relative to 40° and 45° yaw rotations), respectively. Thus, it was expected that both experiment 1b and experiment 2 will show worse performance for pitch-down than yaw rotations. Predictions regarding comparisons between pitch-up and pitch-down rotations were more tentative, since this comparison has not been made previously; however, given overall worse performance in pitch-down rotations it was possible that we might find a steeper decline in recognition performance for this axis in both experiments.

⁽²⁾ Sequential matching performance for pitch-up and pitch-down rotations was investigated in two separate experiments owing to the large number of trials that a fully within-subjects design (including yaw, pitch-up, and pitch-down rotations) would have generated. As the old/new matching task was quicker to run and required fewer trials than the sequential matching task, both pitch-up and pitch-down conditions could be examined within subjects.

2 Experiment 1a

2.1 Method

2.1.1 *Participants.* Thirty-seven naive Caucasian undergraduates from the University of Wollongong participated in exchange for course credit. All had normal or corrected-to-normal vision and none was familiar with the face stimuli/models.

2.1.2 *Stimuli.* The faces of ten female Caucasian ‘models’ were photographed from 12 different viewpoints (see figure 1 for an example of the stimuli constructed and the full range of rotations examined).⁽³⁾ The goal was to produce an ecological pattern of facial lighting (with more light originating from above than below the face). Faces were illuminated by four directional light sources located 1 m in front of the participant (just above and to the left, just above and to the right, just below and to the left, and just below and to the right of the face)⁽⁴⁾ and an ambient light source located directly above them. As the camera was moved with respect to the model’s stationary head, the relative pattern of light intensity across his/her face remained constant. Where visible, the average luminance of the participant’s forehead (49.36 cd m^{-2} ; SD 3.36 cd m^{-2}) was always greater than that of his/her chin region (26.91 cd m^{-2} ; SD 2.21 cd m^{-2}). The photographs were taken with a Fuji FinePix 4-megapixel camera. The camera was set at an ISO sensitivity of 100. Aperture, shutter speed, and focus were automatically controlled. Models removed all jewellery and glasses. Each model wore a black headband (placed at the model’s natural hairline to prevent hair protruding onto the forehead) and a black T-shirt. The camera was placed 59 cm from the tip of the models’ noses for each camera angle, by means of a semicircular frame. The models were asked to sit as still as possible and fixate directly ahead at a cross located at 0° on the semi-circular frame whilst maintaining a neutral expression.⁽⁵⁾ The camera was moved to take photos of the left-hand side of the models’ faces at 0° (full-face), 15° , 30° , 45° , 60° , 75° , and 90° (profile) away from 0° (ie about the yaw axis). Pitch-up images were taken at 15° , 30° , 45° , 60° , and 75° below 0° on the pitch axis. 75° was the lowest camera placement possible below the face models’ chins.

The face photographs were processed with Adobe Photoshop 7.0 as follows. The hair region, ears, and whole back of the head were removed completely. To remove cues to hairstyle, different faces were superimposed at corresponding angles of rotation during this process to ensure that the hairline followed a similar path for each face. As we were primarily interested in recognition based on 3-D shape, blemishes, moles, or other such distinctive features which may have unfairly aided performance, were also removed. Faces were set against a white background. All faces were matched for size. The 0° face images had a height on-screen of approximately 16 cm and a width of 10 cm (subtending a visual angle of $14.7 \text{ deg} \times 9.2 \text{ deg}$). While the height of the faces was held constant for yaw rotations, the width of the face images increased with their angular deviation from 0° . The widest face image was produced by 90° yaw rotation and was approximately 12 cm on-screen (subtending an area of $14.7 \text{ deg} \times 11 \text{ deg}$). During pitch-up rotations, face width was held constant while face height decreased with each increase in angular deviation from 0° . The smallest face image was produced by a 75° pitch-up rotation and was 12 cm high and 10 cm wide (subtending a visual area of

⁽³⁾ Female faces were used as stimuli in the current study since we were using photographic images of real people and facial hair cues on male faces, despite being clean-shaven, presented a potential confound.

⁽⁴⁾ The two lights located just above the face were oriented horizontally (ie at 90° to gravity), whereas the lights located just below the face were oriented 45° below the horizontal.

⁽⁵⁾ Slight movements by the models in-between takes are inevitable. We kept the photo session as short as possible (approximately 10–15 min) to minimise the impact of this issue. Face position and expression were checked via a grid superimposed on the camera viewfinder at each take by the photographer.

approximately $11 \text{ deg} \times 9.2 \text{ deg}$). The pattern mask used in this experiment subtended a visual area of approximately $18 \text{ deg} \times 22 \text{ deg}$ and consisted of elements from a variety of the face images.

2.1.3 Apparatus. The images were presented in full colour on a 48 cm flat-screen monitor with resolution of 1024×768 pixels. The trial sequence was controlled by RSVP experimental software (Version 4.0.5; www.tarlab.org) on a Macintosh G4 computer. A chin-rest maintained participants' head position and distance from the monitor (62 cm). Participants made their responses on a keyboard placed on the table in front of the chin-rest.

2.1.4 Design. This was a two-factor design with both rotation type (rightwards yaw or pitch-up) and angular deviation between encoding and test phases (0° to 90°) as within-subjects factors. There were 12 conditions in experiment 1a: 0° ; 15° , 30° , 45° , 60° , 75° , and 90° yaw rotations; and 15° , 30° , 45° , 60° , and 75° pitch-up rotations.

2.1.5 Procedure. Participants were tested individually according to a sequential matching paradigm. Participants viewed pairs of sequentially presented faces and were asked to judge whether the two images were of the same person, regardless of the orientation at which the faces were displayed. Participants were instructed to indicate whether the faces were the 'same' or 'different' by pressing corresponding keys on a keyboard as quickly and as accurately as possible. The experiment began with 20 practice trials, with three different faces that were not used in the experimental trials that followed.

Each trial began with a fixation cross that remained on-screen for 500 ms. This was followed by the first (encoding) face, which could be displayed at any one of 12 different orientations (0° , 15° to 90° yaw rotations, or 15° to 75° pitch-up rotations). Encoding views remained on-screen for 250 ms and were then immediately followed by the pattern mask that appeared for 500 ms. Next, the second (test) face was presented (which was always 0°), and remained visible for 250 ms. The pattern mask appeared immediately after the test face, remaining on-screen for 500 ms. The screen then remained blank for 2000 ms, or until the participant responded. The inter-trial interval was 1000 ms. All faces were displayed in the centre of the screen.

The experiment consisted of 480 randomly ordered trials. Participants were given self-timed rest periods every 60 trials, and recommenced the experiment by pressing the space bar. In half of the trials the test view was the same face as the encoding view ('same' trials). For the remaining trials the test face was a distractor face, different from the one shown at encoding ('different' trials). Distractor faces were drawn pseudo-randomly from the 0° views of the remaining nine faces. The randomisation was constrained such that no distractor face was paired with the same target face (at any angle of rotation) more than twice throughout the experiment. Computer beeps provided feedback on incorrect responses.

2.2 Results

Participants' 'same/different' (yes/no) responses were converted into hit (H) and false-alarm (F) rates, where a hit was a correct 'same' response to a face during the test phase, and a false alarm was an incorrect 'same' response to a 'different' face during the test phase. H and F rates were converted into z -scores and used to calculate d' (see Macmillan and Creelman 1991). Reaction-time data (RTs) were also collected.

2.2.1 Sensitivity analyses. The d' values were subjected to a 2 (rotation type: rightwards yaw and pitch-up) \times 5 (angular deviation: 15° to 75°) repeated-measures ANOVA. Greenhouse–Geisser corrections were applied to the df wherever the assumption of sphericity was violated. The ANOVA revealed a significant interaction between rotation type and angular deviation for sensitivity ($F_{4,144} = 10.9$, $p < 0.001$). As can be seen in

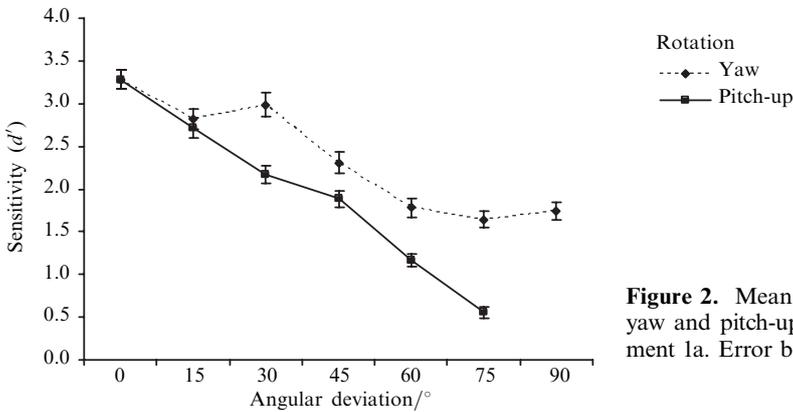


Figure 2. Mean sensitivity (d') for yaw and pitch-up rotations in experiment 1a. Error bars represent 1 SEM.

figure 2, this interaction indicates that the rate of decline in recognition performance, with increasing angular deviation between encoding and test views, was greater for pitch-up rotations than that for yaw rotations. Bonferroni adjusted a-posteriori pairwise comparisons between yaw and pitch-up at each angular deviation indicated that there was no difference between these two rotation types at 15° ($p = 0.30$), but that there was significantly greater sensitivity to yaw than pitch-up rotations at all other angular deviations (all $ps < 0.01$) with the largest difference occurring at 75°. Simple-main-effects analysis of the effect of rotation for yaw and pitch-up separately showed a significant effect of angular deviation for both yaw ($F_{4,144} = 52.0$, $p < 0.001$) and pitch-up rotations ($F_{4,144} = 105.2$, $p < 0.001$), with d' decreasing as the angular deviation between encoding and test views increased.

2.2.2 Reaction-time analyses. A 2 (rotation type: rightwards yaw and pitch-up) \times 5 (angular deviation: 15° to 75°) repeated-measures ANOVA was also conducted on the participants' mean RTs for correctly answered 'same' trials. The results reflected the patterns in the d' data. The interaction between angular deviation and rotation type reached significance ($F_{4,144} = 2.93$, $p < 0.05$). Looking at figure 3, we find a trend (though not as clear as for the d' data) suggesting a steeper rate of increase in RT for pitch-up rotations than for yaw rotations as the angular deviation between encoding and test views increased. While Bonferroni adjusted a-posteriori pairwise comparisons between yaw and pitch-up at each angular deviation indicated that there was no difference between these two rotation types at 60° ($p = 0.72$), the RTs for yaw rotation were significantly faster than for pitch-up at all other angular deviations (all $ps < 0.01$). Simple-main-effects analysis of the effect of rotation for yaw and pitch-up separately showed

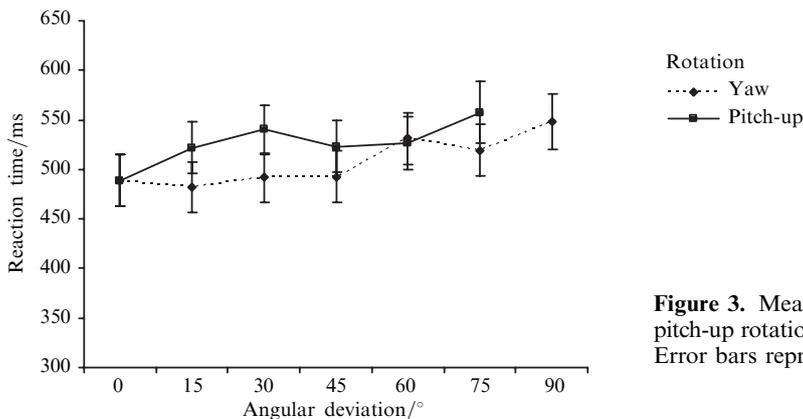


Figure 3. Mean RT for yaw and pitch-up rotations in experiment 1a. Error bars represent 1 SEM.

a significant effect of angular deviation for both yaw ($F_{4,144} = 5.9, p < 0.001$) and pitch-up rotations ($F_{4,144} = 3.1, p < 0.05$), with RTs decreasing as the angular deviation between encoding and test views increased.

2.3 Discussion

Recognition performance (as indexed by both sensitivity and reaction time) was shown to be viewpoint-dependent during both yaw and pitch-up rotations. Consistent with the findings of Wallraven et al (2002), but unlike those of Liu et al (2005), there was a stronger effect of the angular deviation between encoding and test views during pitch-up rotations than during yaw rotations. The superior performance of participants during yaw rotations compared to pitch-up rotations demonstrates that the viewpoint-dependence observed in face recognition is not solely determined by the degree of angular deviation between the encoding and test views. In experiment 1b we sought to clarify if similar effects would be found for pitch-down rotations.

3 Experiment 1b

3.1 Method

3.1.1 *Participants.* Thirty-seven naive Caucasian undergraduates from the University of Wollongong participated in exchange for course credit, and were not familiar with the face stimuli. All had normal or corrected-to-normal vision.

3.1.2 *Stimuli and apparatus.* Pitch-down images were taken at 15°, 30°, 45°, 60°, 75°, and 90°; with 90° being directly above the model's head (figure 1). Face models and stimuli were the same as those described in experiment 1a, with the following exception: whereas experiment 1a was concerned with pitch-up and rightwards yaw rotations, experiment 1b was concerned with pitch-down and rightwards yaw rotations. As with pitch-up rotations, face width was held constant for pitch-down rotations while face height decreased with increasing angular deviations away from 0°. The smallest face image was 90° pitch-down, at 4 cm high (subtending a visual area of approximately 3.7 deg × 9.22 deg). The pattern mask and experimental apparatus were identical to those of experiment 1a.

3.1.3 *Design.* Experiment 1b was also a two-factor design with both rotation type (rightwards yaw and pitch-down) and angular deviation (0° to 90°) as within-subjects factors. There were 13 conditions in experiment 1b: angular deviations of 0°, 15°, 30°, 45°, 60°, 75°, and 90° in both rightwards yaw rotation and pitch-down rotation.

3.1.4 *Procedure.* The stimulus presentation times and trial sequence were the same as in experiment 1a, except that the orientations at which faces were shown in the encoding phase were 0°, 15° to 90° rotated about yaw, or 15° to 90° rotated about pitch-down. 20 practice trials preceded a total of 520 randomly ordered experimental trials. Participants were given self-timed rest periods every 65 trials.

3.2 Results

3.2.1 *Sensitivity analyses.* As in experiment 1a hit and false-alarm rates were converted into z -scores and used to calculate d' . A 2 (rotation type: yaw and pitch-down) × 6 (angular deviation: 15° to 90°) repeated-measures ANOVA was conducted. Greenhouse–Geisser corrections were applied to the df wherever the assumption of sphericity was violated. This revealed a significant interaction between rotation type and angular deviation ($F_{5,180} = 15.35, p < 0.001$). As can be seen in figure 4, this interaction indicated that the rate of decline in recognition performance following pitch-down rotation was much steeper than that following yaw rotation as the angular deviation between encoding and test views increased. Note that although close, performance does not fall to chance levels. One-sample t -tests indicate that recognition sensitivity (d') following 75° and 90° pitch-down rotations was significantly greater than 0 (ie chance-level performance)

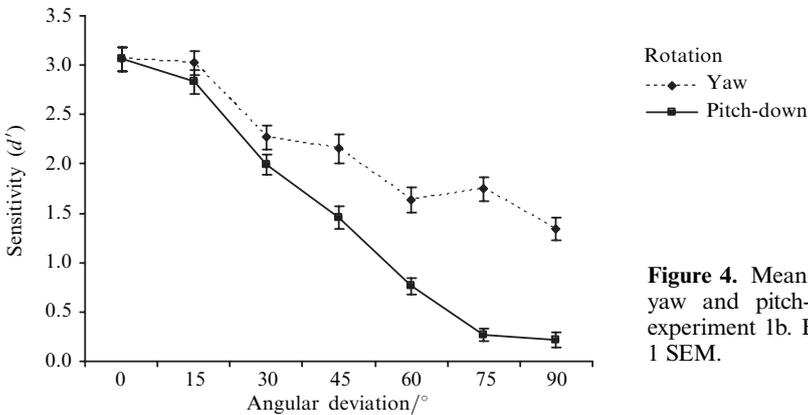


Figure 4. Mean sensitivity (d') for yaw and pitch-down rotations in experiment 1b. Error bars represent 1 SEM.

($t_{36} = 4.06$, $p < 0.001$, and $t_{36} = 2.90$, $p = 0.006$, respectively). Bonferroni adjusted a-posteriori pairwise comparisons between yaw and pitch-down at each angular deviation indicated that there was no difference between the two rotation types at 15° ($p = 0.20$), but that there was a significant difference between them at all other angular deviations (all $ps < 0.05$) with the greatest difference occurring at 75° and 90°. Simple-main-effects analysis of the effect of rotation for yaw and pitch-down separately showed a significant effect of angular deviation for both yaw ($F_{5,180} = 44.5$, $p < 0.001$) and pitch-down rotations ($F_{5,180} = 134.6$, $p < 0.001$), with d' decreasing as the angular deviation between encoding and test views increased.

3.2.2 RT analyses. Participants' mean RT data for correctly answered 'same' trials were submitted to the same analyses as the sensitivity data. A 2 (rotation type: yaw, and pitch-down) \times 6 (angular deviation: 15° to 90°) repeated-measures ANOVA revealed a significant interaction between the two factors ($F_{3,97,143.03} = 9.39$, $p < 0.001$), suggesting that RT increased more with the angular deviation in pitch-down rotation conditions than it did in yaw rotation conditions (figure 5 shows the mean RT for experiment 1b). Bonferroni adjusted a-posteriori pairwise comparisons between yaw and pitch-down at each angular deviation indicated that there was a significant difference between these two rotation types at all angular deviations (all $ps < 0.05$). Importantly, the mean difference in RT between yaw and pitch-down increased from 22 ms at 15° to 130 ms at 90°. Simple-main-effects analysis of the effect of rotation for yaw and pitch-down separately showed a significant effect of angular deviation for both yaw ($F_{5,180} = 4.6$, $p < 0.001$) and pitch-down rotations ($F_{5,180} = 47.8$, $p < 0.001$), with RTs increasing as the angular deviation between encoding and test views increased.

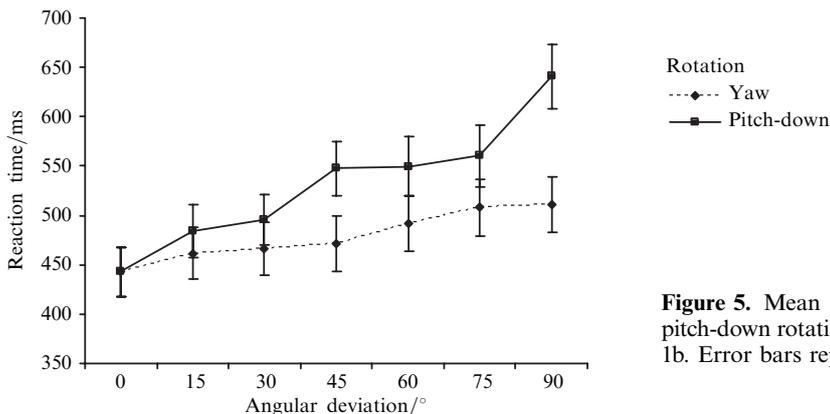


Figure 5. Mean RT for yaw and pitch-down rotations in experiment 1b. Error bars represent 1 SEM.

3.3 Discussion

Following the pattern of results found in the first experiment, experiment 1b showed that yaw and pitch-down rotations with equivalent angular deviations from 0° produced markedly different rates of decline in recognition performance (both sensitivity and RT). Recognition performance following rotation about yaw was more robust to increases in the angular deviation between views (compared to performance following pitch-down rotations). In the current experiment, recognition sensitivity started to diverge for pitch-down and yaw rotations of more than 30° (see figure 4). This pattern of performance appears to fall in-between that found by Wallraven et al (2002), where the divergence between pitch-down and yaw rotations started to emerge at 15° and Liu et al (2005), where this divergence occurred at 40°. Importantly, we found that this divergence in performance continued for rotations exceeding 45°.

3.4 Comparison of pitch-up and pitch-down rotations

To compare recognition performance for the two directions of pitch rotation (pitch-up and pitch-down), d' data from experiments 1a and 1b were subjected to a split-plot ANOVA with angular deviation (15° to 75°) as a within-subjects factor and rotation type (pitch-up/pitch-down) as a between-subjects factor.⁽⁶⁾ Greenhouse–Geisser corrections were applied to the df wherever the assumption of sphericity was violated. There was a significant interaction between rotation type and angular deviation ($F_{3,5,252.5} = 4.70, p < 0.01$). Recognition sensitivity diverged at 45°, such that performance was more robust for pitch-up rotations than for pitch-down rotations. However, for rotation exceeding 45°, the difference between pitch-up and pitch-down rotation remained relatively constant. Bonferroni adjusted a-posteriori pairwise comparisons between pitch-up and pitch-down at each angular deviation indicated that there was no difference between these two rotation types at 15° or 30° (both $ps = 0.56$), but that there was a significant difference between the rotation types at all other angular deviations (all $ps < 0.01$). This finding is partially in line with Wallraven et al (2002) and Liu et al (2005); however, we have clearly demonstrated that there is indeed a difference in recognition sensitivity following pitch-up and pitch-down rotations of faces and that this extends to rotations larger than 45°.

4 Experiment 2

Experiments 1a and 1b demonstrated a differential effect of rotation type on viewpoint-dependent face recognition in a sequential matching task. Face recognition showed greatest decline following pitch-down rotations, less decline following pitch-up rotations, and least decline following yaw rotations. Our results were partially consistent with those of Wallraven et al (2002) who used an old/new recognition task, tapping into longer-term memory, to study face recognition. In order to determine whether differences in processing demands between the two tasks may account for the inconsistent findings, we conducted a second experiment using the conditions included in experiments 1a and 1b in an old/new recognition task. In addition to bridging the gap between the current study and Wallraven et al, results from this experiment should allow us to establish whether rotation has a similar effect on short-term and longer-term representations used in face recognition.

In experiments 1a and 1b, the aim in each trial was to match an initially displayed rotated face to a full-face view. In experiment 2, a full-face view was encoded in the initial learning phase and rotated faces were shown at test for recognition. Thus, whereas in experiments 1a and 1b we investigated the recognition of yaw and pitch rotated faces

⁽⁶⁾A similar analysis was conducted for the yaw data across experiments to ensure that performance in both experiments could be reasonably matched. There was no overall difference between the two experiments in recognition sensitivity following yaw rotations ($F = 1.8$).

compared to a full-face view, in experiment 2 we investigated the recognition of yaw and pitch rotated faces from a full-face view.⁽⁷⁾

4.1 Method

4.1.1 *Participants.* Thirty-six naive Caucasian undergraduates from the University of Wollongong participated in exchange for course credit, and were not familiar with the face stimuli. All had normal or corrected-to-normal vision.

4.1.2 *Stimuli and apparatus.* The stimuli and experimental apparatus were identical to those of experiments 1a and 1b. That is, images of yaw, pitch-up, and pitch-down rotated faces were used.

4.1.3 *Design.* Experiment 2 was a two-factor design with both rotation type (rightwards yaw, pitch-up, and pitch-down) and angular deviation (0° to 90°) as within-subjects factors. There were 18 conditions in experiment 2: angular deviations of 0° , 15° , 30° , 45° , 60° , 75° , and 90° in both rightwards yaw rotation and pitch-down rotation and angular deviations of 15° , 30° , 45° , 60° , and 75° in pitch-up rotation.

4.1.4 *Procedure.* Participants were tested individually with an old/new recognition paradigm. In the learning phase, participants viewed five target faces displayed at 0° (20 s total for each face). In the subsequent testing phase, participants were shown the five target and five distractor faces for 1 s each and asked to judge whether the face had been shown in the learning phase (ie was a target/old) or not (ie was a distractor/new), regardless of the orientation at which the face was displayed.

Each trial in the learning phase began with a fixation cross that remained on-screen for 500 ms. This was followed by a target face displayed at 0° which remained on-screen for 10 s. No response was required, but participants were instructed that they would be asked to remember the target faces in a subsequent test. The inter-trial interval was 1000 ms. All faces were displayed in the centre of the screen. The five target faces were displayed in random order. This learning phase was then repeated, giving a total of 10 learning trials (ie 20 s viewing time for each target face).

Each trial in the testing phase began with a fixation cross that remained on-screen for 500 ms. This was followed by a face displayed on-screen for 1 s at one of 18 orientations (0° , rotated 15° , 30° , 45° , 60° , 75° , or 90° about yaw or pitch-down, or rotated 15° , 30° , 45° , 60° , or 75° about pitch-up). Participants were instructed to indicate whether the faces were 'old' (ie targets shown in the learning phase) or 'new' (ie distractor faces) regardless of the orientation at which the face was displayed. Responses were made by pressing corresponding keys on a keyboard as quickly and as accurately as possible. In half of the trials the test view was the same face as the encoding view seen in the learning phase ('old' trials). For the remaining trials the test face was one of five distractor faces, different from the five learned faces ('new' trials). The first block of the test phase of the experiment consisted of 180 randomly ordered trials. After this first test block, participants were given the learning phase again (ie 5 trials showing the target faces) before completing a second test phase block identical to the first (with a different randomisation of trials). Participants were given self-timed rest periods every 45 trials during the test blocks, and recommenced the experiment by pressing the space bar. No feedback was provided.⁽⁸⁾

⁽⁷⁾ There was a possibility that task differences in recognising rotated faces compared to a full-face view versus recognising rotated faces from a full-face view might have made the comparison of results between the two current experiments difficult. However, in control experiments with sequential matching tasks we had participants match either from or to a full-face view. We found no significant difference in recognition performance between these two conditions ($F = 0.006$).

⁽⁸⁾ A pilot study showed that performance in the recognition task was at ceiling with feedback; thus, feedback was not included in experiment 2.

4.2 Results

Data analysis was conducted on ‘old’ trials. Because of the unequal number of levels of rotation within each axis, three separate repeated-measures ANOVAs were conducted on the mean d' and RT data examining: (i) yaw versus pitch-up, (ii) yaw versus pitch-down, and (iii) pitch-up versus pitch-down. Greenhouse–Geisser corrections were applied to the df wherever the assumption of sphericity was violated.

4.2.1 Sensitivity analyses. Hit and false-alarm rates were converted into z -scores and used to calculate d' (see figure 6). A 2 (rotation type: yaw and pitch-up) \times 5 (angular deviation: 15° to 75°) repeated-measures ANOVA revealed a significant interaction between rotation type and angular deviation ($F_{4,140} = 2.7$, $p = 0.026$). Bonferroni adjusted a-posteriori pairwise comparisons indicated that there was no difference between the two rotation types at 15°, 30°, 45°, or 60° (all $ps > 0.05$), but that there was a significant difference between them at 75° ($p < 0.05$). Simple-main-effects analysis of the effect of rotation for yaw and pitch-up separately showed a significant effect of angular deviation for both yaw ($F_{4,140} = 6.6$, $p < 0.001$) and pitch-up rotations ($F_{4,140} = 26.7$, $p < 0.001$), with d' decreasing as the angular deviation between encoding and test views increased.

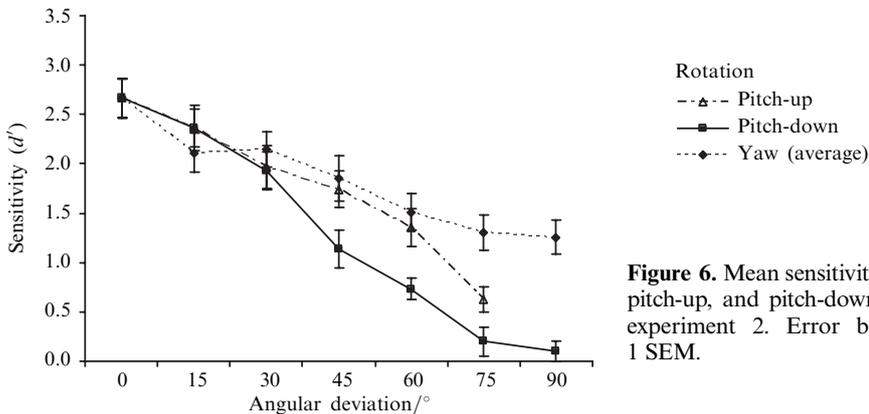


Figure 6. Mean sensitivity (d') for yaw, pitch-up, and pitch-down rotations in experiment 2. Error bars represent 1 SEM.

A 2 (rotation type: yaw and pitch-down) \times 6 (angular deviation: 15° to 90°) repeated-measures ANOVA revealed a significant interaction between rotation type and angular deviation ($F_{5,175} = 8.0$, $p < 0.001$). Bonferroni adjusted a-posteriori pairwise comparisons indicated that there was no difference between the two rotation types at 15° or 30° (both $ps > 0.05$), but that there was a significant difference between them at 45°, 60°, 75°, and 90° (all $ps < 0.01$). Simple-main-effects analysis of the effect of rotation for yaw and pitch-down separately showed a significant effect of angular deviation for both yaw ($F_{5,175} = 8.6$, $p < 0.001$) and pitch-down rotations ($F_{5,175} = 67.6$, $p < 0.001$), with d' decreasing as the angular deviation between encoding and test views increased.

A 2 (rotation type: pitch-up and pitch-down) \times 5 (angular deviation: 15° to 75°) repeated-measures ANOVA revealed a significant interaction between rotation type and angular deviation ($F_{4,140} = 3.2$, $p = 0.019$). Bonferroni adjusted a-posteriori pairwise comparisons indicated that there was no difference between the two rotation types at 15° or 30° (both $ps > 0.05$), but that there was a significant difference between them at 45°, 60°, and 75° (all $ps < 0.05$). Simple-main-effects analysis of the effect of rotation for yaw and pitch-up separately showed a significant effect of angular deviation for both pitch-up ($F_{4,140} = 39.3$, $p < 0.001$) and pitch-down rotations ($F_{4,140} = 54.0$, $p < 0.001$), with d' decreasing as the angular deviation between encoding and test views increased.

4.2.2 RT analyses. RT data analysis was conducted on correct 'old' trials (see figure 7). A 2 (rotation type: yaw and pitch-up) \times 5 (angular deviation: 15° to 75°) repeated-measures ANOVA revealed no interaction between rotation type and angular deviation ($F = 1.6$). There were significant main effects of rotation type ($F_{1,35} = 10.9$, $p = 0.002$), RT was slower for pitch-up rotation than for yaw rotations, and angular deviation ($F_{4,140} = 7.1$, $p < 0.001$). RT showed a viewpoint-dependent pattern of performance such that RT increased as angular deviation between encoding and test views increased.

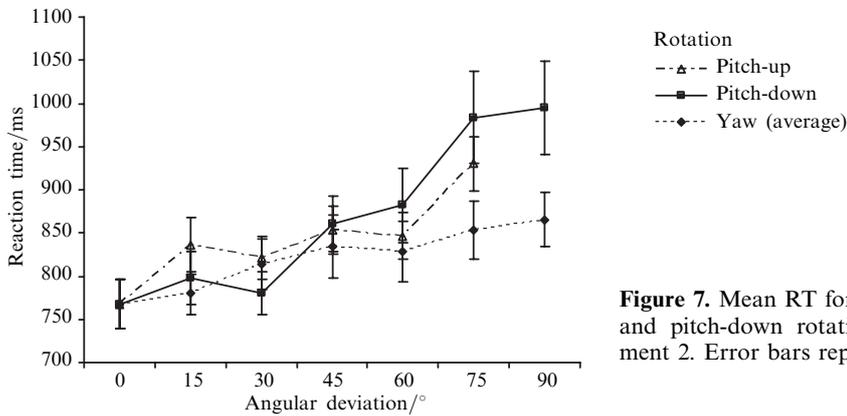


Figure 7. Mean RT for yaw, pitch-up, and pitch-down rotations in experiment 2. Error bars represent 1 SEM.

A 2 (rotation type: yaw and pitch-down) \times 6 (angular deviation: 15° to 90°) repeated-measures ANOVA revealed a significant interaction between rotation type and angular deviation ($F_{3,6,124.3} = 4.4$, $p = 0.003$). Bonferroni adjusted a-posteriori pairwise comparisons indicated no difference between rotation types at 15°, 30°, 45°, or 60° (all $ps > 0.05$), but RT was faster for yaw rotations than pitch-down at 75° and 90° (both $ps < 0.05$). Simple-main-effects analysis of the effect of rotation for yaw and pitch-down separately showed a significant effect of angular deviation for pitch-down rotations ($F_{2,7,124.3} = 22.6$, $p < 0.001$), with RT increasing as the angular deviation between encoding and test views increased, but not for yaw ($p > 0.05$).

A 2 (rotation type: pitch-up and pitch-down) \times 5 (angular deviation: 15° to 75°) repeated-measures ANOVA revealed no significant interaction between rotation type and angular deviation ($p > 0.05$). There was a significant main effect of angular deviation ($F_{2,3,79.4} = 17.0$, $p < 0.001$). RT showed a viewpoint-dependent pattern of performance, such that RT increased as angular deviation between encoding and test views increased. There was no effect of type of rotation ($p > 0.05$).

5 General discussion

The goal of the present study was to directly examine the effects of the angular deviation between encoding and test face views on face-recognition sensitivity for rotations along both pitch and yaw compared to a full-face view. Results from both the sequential matching and old/new recognition tasks showed that face-recognition sensitivity for pitch-up, pitch-down, and yaw rotation was affected in a viewpoint-dependent manner. The interactions between rotation type and angular deviation indicated that the rate of viewpoint-dependent decline in sensitivity following pitch-up or pitch-down rotation was greater than that found following equivalent yaw rotation. Sensitivity for pitch-down rotation diverged from yaw at a similar point for both tasks (sequential matching at 30°; old/new recognition at 45°). However, there was a substantial difference between the two tasks regarding the point at which sensitivity for pitch-up rotation diverged from yaw (sequential matching at 30°; old/new recognition at 75°).

Upon comparing the effects of pitch-up and pitch-down rotations directly, it was found that recognition sensitivity following pitch-up rotations was reliably better than following pitch-down rotations when the angular deviations were larger than 30° for the sequential matching task and 45° for the old/new recognition task. That is, sensitivity to pitch-down rotation diverged from pitch-up at a similar point in both tasks.

Our pattern of results shares only some similarities with those of Wallraven et al (2002) and Liu et al (2005), suggesting that task differences or memory load alone cannot account for the disparity in the findings of these studies. The current study employed both a sequential matching task and an old/new recognition task to tap into face recognition at two different levels. Sequential matching involves short-term memory use, whereas an old/new recognition task involves learning and the use of longer-term memory. Wallraven et al also used an old/new recognition task while Liu et al used an 8-alternative forced-choice simultaneous matching task, which could be argued to involve minimal use of memory. Interestingly, while the results of our sequential matching task are in line with those of Wallraven et al (2002), indicating that both pitch-down and pitch-up rotations were significantly more disruptive to recognition sensitivity than yaw rotations, our old/new recognition results are only partially supported in that sensitivity to pitch-down and not pitch-up rotation was poorer than to yaw up to 40°. We found that sensitivity to pitch-up rotation diverged from yaw only at 75° in the old/new recognition task. On the other hand, while this pattern of results for our old/new recognition task supports the findings of Liu et al (2005), the results for the sequential matching task (which is more similar to a simultaneous matching task) do not. Both tasks in the current study showed that sensitivity following pitch-up rotations was reliably better than following pitch-down rotations when the angular deviations were larger than 30° or 45°. This is contrary to the finding of Wallraven et al (2002) that the effects of pitch-up and pitch-down rotation on face-recognition sensitivity were more or less equivalent, but in line with Liu et al's (2005, experiment 3) finding that performance for pitch-down rotation diverged from pitch-up at approximately 40°.

Two types of recognition task were employed in this paper to address the potential issue of different effects of rotation on short-term and longer-term memory. After finding that sensitivity to pitch-down rotation of faces is poor compared to pitch-up and yaw rotations with a sequential matching task, we replicated this pattern of results using an old/new recognition task. This suggests that rotation has a similar and robust effect on both the short-term and longer-term memory processes involved in face recognition.

The current results are compatible with viewpoint-dependent theories of face recognition (eg Tarr and Bülthoff 1998; Troje and Bülthoff 1996). The viewpoint-dependent decline in performance persisted even with the inclusion of natural colour and texture information in the stimuli, which remained relatively stable (and hence less viewpoint-dependent) throughout changes in viewpoint in all directions. In addition to the sensitivity data, the RT data were also shown to be viewpoint-dependent—with RT increasing in a linear fashion with increasing angular deviation of the yaw, pitch-up, or pitch-down rotation. Transformation theories account for these linear increases in RT with increasing angular deviation by postulating a mechanism, such as mental rotation, that transforms the image of the novel view to match the stored representation. According to such theories, RT increases with increasing angular deviation from encoding as greater amounts of mental rotation are required. However, we also found a steeper rate of increase in RT following pitch rotations than following the equivalent yaw rotations. Transformation theorists could potentially explain this differential performance by proposing that humans are better at utilising cues to the amount and direction of mental rotation required for yaw rotation. Following a yaw rotation, these cues

could be provided by the amount of contour information available compared to, say, a profile view. Following a pitch rotation, the vertical position of features (eyes, nose, mouth) within a face could serve as cues to both the direction and amount of rotation from the 0° or full-face orientation (for example, full-face views with internal features shifted upwards in the face frame appear tilted backwards—Collishaw et al 2005). However, it is not clear how transformation theories could account for the differences we found between performance following pitch-up and pitch-down rotations. Further, there are more general problems with mental-rotation accounts of recognition, including difficulties in explaining task and practice effects, and observations that the reported rates of mental rotation appear to be too slow to account for the recognition of misoriented objects (Perrett et al 1998). Perrett et al (1998) also argue that the increasing RT patterns cited as evidence for mental rotation could be the result of any process that takes more time with increasing differences between novel and stored views.

The observed decreases in recognition sensitivity and increases in RT in the present study can be explained without invoking mental rotation. Critically, our results indicate that equivalent rotations around orthogonal facial axes (pitch and yaw) do not produce equivalent viewpoint-dependent declines in recognition performance. This finding is in line with interpolation models of object and face recognition (eg Bülthoff and Edelman 1992; Wallraven et al 2002). Such models do not rely on normalisation or transformation processes (such as mental rotation) to match novel views to stored ones. Rather, they utilise mechanisms that measure the perceptual similarity between different views or exemplars. Generalisation between views is done by establishing the location of the novel view within a view space and measuring the similarity of its features to the nearest known view—‘interpolating’ across the view space. In an attempt to explain their data in terms of interpolation theory, Wallraven et al (2002) argued that face recognition was worse for pitch than for yaw rotation because the change in image-feature information was greater for the same angular deviation in the former compared to the latter. That is, they suggested that the perceptual similarity of full faces to those rotated in pitch was less than the perceptual similarity of full faces to those rotated in yaw—the result being that it was more difficult to generalise across views following pitch rotation.

An interpolation account could be used to explain the observed differences in face-recognition sensitivity following the different types of rotation in the current study. Perceptual similarity could be determined on the basis of the amount of information available in the two images by using a basic measure such as the number of available pixels corresponding to the face. The pitch-down faces in the current study had less information (ie fewer pixels corresponding to the face) than pitch-up faces at extreme rotations (eg 60° – 90°) which may explain the superior performance for the pitch-up faces. Similarly, the greater amount of information available for yaw rotations at 60° – 90° may explain the superior performance compared to both pitch-up and pitch-down rotations. However, the current results show significant differences in performance for the yaw, pitch-up, and pitch-down axes at rotations of 30° . The minimal differences in information available for yaw, pitch-up, and pitch-down faces at 30° rotations would not allow us to predict this result. In addition, the magnitude of information differences between yaw, pitch-up, and pitch-down at extreme rotations (for example, compare stimuli and results for rotations of 75°) does not reflect the pattern of results at these viewpoints. Thus, rather than the amount of information, we propose that perceptual similarity is more likely to be based on differences in the availability of certain types of information (ie featural and configural information) following pitch-up and pitch-down face rotations.

Featural information in faces generally refers to nameable parts of the face (eg eyes, nose, mouth, chin, etc). All faces have a similar layout (or first-order configuration) of features: two laterally separated eyes above a centrally placed nose above a centrally placed mouth. Within this particular layout, individual faces differ in the spacing or distance between features. These distances between featural 'anchors' are second-order spatial relations and are typically referred to as configural information (see Maurer et al 2002 for a review of configural processing in faces). Our skill in perceiving and recognising faces has been attributed to a proficiency in processing configural information (Freire et al 2000; Leder and Bruce 2000; Leder et al 2001; Searcy and Bartlett 1996). However, both featural and configural information are important for face perception and recognition (Schwaninger et al 2002). Changes made to featural information can result in changes to configural information (Rhodes et al 1993) and vice versa (Rakover 2002). Looking at the stimuli and data for the current study (see figures 1, 6, and 7), following a 30° rotation in pitch-up or pitch-down from the full-face view, information about most internal features was still available and in both tasks performance was equivalent. However, following 45° rotation, performance diverged such that sensitivity to pitch-down rotations was always worse than that for pitch-up rotations. In pitch-down rotation, the eyes and lips become progressively more occluded with each increase in angular deviation compared to equivalent pitch-up rotation. However, in pitch-up rotation, there was less occlusion of features, with the eyes, nose, mouth, and even the eyebrows still visible at 75°. Thus, with increasing featural information loss and loss of 'anchors' for second-order configural information, performance in pitch-rotation conditions was impaired compared to full-face views and yaw rotation.

Perrett and colleagues (1998) offered a physiological explanation of viewpoint effects on face recognition that is in line with interpolation models. They proposed that people take longer to recognise familiar faces in unusual orientations because the neural representation of the face is less strongly activated. On the basis of neurophysiological data, they proposed a mechanism whereby repeated exposure to a particular face view increased the number of cells tuned to that view and decreased the time required for the population response to exceed the threshold level for recognition. Note that face-selective cells have been found that respond maximally to yaw rotations as well as cells (although fewer in number) that respond to pitch rotations (Perrett et al 1985). Importantly, these 'view-selective' face cells, formed by exposure to one particular view, were still able to aid recognition even after a dramatic change in view. Although these cells were selective to a particular view, their response remained above baseline even after a 135° rotation of the face about the yaw axis. Thus, their ability to generalise from the familiar view of the face to the novel test view of the face should depend on how well they tolerate differences in orientation, size, etc (relative to the familiar face view). If the test view falls within the tuning functions of face cells, then both the speed and the accuracy of recognition will depend on the magnitude of the difference between the two views (although accuracy should always be above chance in this situation). Thus, one possible interpretation of the current findings is that view-selective face cells were less tolerant to pitch rotations than they were to yaw rotations (eg the tuning functions were narrower for pitch than for yaw).⁽⁹⁾ Alternatively, since only a fraction of the face-selective cells identified by Perrett et al (1985) were maximally responsive to images of faces with the head tilted up (pitch-up) or lowered (pitch-down) compared to yaw, it may be the case that recognition of faces rotated in yaw is facilitated compared to that of faces rotated in pitch.

⁽⁹⁾ Given that performance remained above chance following the pitch and yaw rotations tested, the face views examined within this experiment appear to have been within the tuning function of these face cells.

Two factors that might account for faces rotated about pitch being less similar to a full-face view than faces rotated about yaw are: (i) the bilateral symmetry of a face, and (ii) the nature of the lighting. First, information lost by rotations in yaw could have been compensated for by the bilateral symmetry of the face. Troje and Bühlhoff (1996) have argued that an observer can use his/her knowledge of the bilateral symmetry of faces to generate a 'virtual view' symmetric to the one he/she is actually seeing.⁽¹⁰⁾ This virtual view could then be used as a second view to recover 3-D structure. It could also be used to recover featural information lost when the amount of yaw rotation exceeded 45° and features such as one eye and one corner of the mouth were occluded. Second, the lighting for the faces used as stimuli in the current study (as in the real world) was predominantly from above. Thus, while modest shading changes would have occurred in the conditions investigated here, these changes would have been less dramatic for rotations about yaw than pitch, since views from below the face (pitch-up rotations) would have partially occluded the overhead light source, resulting in some shadowing under the chin. This appeared to be the case for the current stimuli (see figure 1). In relation to this point, Hill and Bruce (1996) have demonstrated that face perception and recognition performance depended on the lighting direction as well as on the viewpoint (only yaw axis rotation was examined). They also found that the effects of viewpoint were dependent on lighting direction (light from above or below). It has yet to be shown whether lighting direction effects contribute to asymmetries in viewpoint-dependent performance (eg pitch-up versus pitch-down). However, this possibility is intriguing. In our experiment, the face lighting was predominantly from above. It is possible then that our pattern of results for pitch-up and pitch-down rotation might reverse if lighting was predominantly from below. In order to further investigate the role of lighting in face recognition, we plan to examine the effect of lighting direction on both pitch and yaw face rotation in future research.

The current study is the first aimed at specifically investigating human face recognition for faces rotated about both the yaw and pitch axes. As such it has raised several questions for future research. One of these is the further exploration of the amount and type of information lost during rotation as an explanation for the differential results between pitch and yaw. To address this issue we could compare performance following pitch and yaw rotations to conditions in which features are not occluded, ie rotation about the roll axis (Valentine and Bruce 1988).

In summary, the present study has filled an apparent gap in the face-processing literature by investigating the effect of a common source of viewpoint variation—rotation about pitch—on face recognition based on 3-D shape. We directly compared face-recognition performance following a wide range of pitch-up, pitch-down, and yaw rotations to full-face views and found similar results for two different recognition tasks tapping into short-term and longer-term visual memory. In both tasks, face recognition was found to be viewpoint-dependent for rotations about both pitch and yaw axes; however, performance depended on both the direction and axis of rotation. Recognition sensitivity for pitch-down rotation was poorer than that for pitch-up or yaw rotation. Interestingly, despite a substantial amount of occlusion occurring following many of the rotations, 75° and 90° pitch-down rotations in particular, recognition sensitivity never fell to chance levels. While sensitivity was relatively poor at these extreme views, the available information was still sufficient to allow better-than-chance recognition to occur. Interpolation theories account for these results by suggesting that the decline in recognition performance reflects the decline in perceptual similarity and loss of featural 'anchors' between full-face views and rotated views.

⁽¹⁰⁾Based on their finding that generalisation to the symmetric view is better than to other views.

References

- Bülthoff H H, Edelman S, 1992 "Psychophysical support for a two-dimensional view interpolation theory of object recognition" *Proceedings of the National Academy of Sciences of the USA* **89** 60–64
- Bülthoff H H, Edelman S, Tarr M J, 1995 "How are three-dimensional objects represented in the brain?" *Cerebral Cortex* **5** 247–260
- Busey T A, Zaki S R, 2004 "The contribution of symmetry and motion to the recognition of faces at novel orientations" *Memory & Cognition* **32** 916–931
- Campbell R, Benson P J, Wallace S B, Doesbergh S, Coleman M, 1999 "More about brows: How poses that change brow position affect perceptions of gender" *Perception* **28** 489–504
- Campbell R, Wallace S B, Benson P J, 1996 "Real men don't look down: Direction of gaze affects sex decisions on faces" *Visual Cognition* **3** 393–412
- Collishaw S M, Hole G J, Schwaninger A, 2005 "Configural processing and perceptions of head tilt" *Perception* **34** 163–168
- Desimone R, Albright T D, Gross C G, Bruce C, 1984 "Stimulus-selective properties of inferior temporal neurons in the macaque" *Journal of Neuroscience* **4** 2051–2062
- Ellis R, Allport D A, 1986 "Multiple levels of representation for visual objects: A behavioural study", in *Artificial Intelligence and Its Applications* Eds A G Cohn, J R Thomas (New York: John Wiley) pp 245–257
- Freire A, Lee K, Symons L A, 2000 "The face-inversion effect as a deficit in the encoding of configural information: Direct evidence" *Perception* **29** 159–170
- Geldart S, Maurer D, Henderson H, 1999 "Effects of the height of the internal features of faces on adults' aesthetic ratings and 5-month-olds' looking times" *Perception* **28** 839–850
- Hill H, Bruce V, 1996 "Effects of lighting on the perception of facial surfaces" *Journal of Experimental Psychology: Human Perception and Performance* **22** 986–1004
- Hill H, Schyns P G, Akamatsu S, 1997 "Information and viewpoint dependence in face recognition" *Cognition* **62** 201–222
- Lander K, Chuang L, 2005 "Why are moving faces easier to recognize?" *Visual Cognition* **12** 429–442
- Leder H, Bruce V, 2000 "When inverted faces are recognized: the role of configural information in face recognition" *Quarterly Journal of Experimental Psychology A* **53** 513–536
- Leder H, Candrian G, Huber O, Bruce V, 2001 "Configural features in the context of upright and inverted faces" *Perception* **30** 73–83
- Liu C H, Collin C A, Farivar R, Chaudhuri A, 2005 "Recognizing faces defined by texture gradients" *Perception & Psychophysics* **67** 158–167
- Lyons M J, Campbell R, Plante A, Coleman M, Kamachi M, Akamatsu S, 2000 "The Noh mask effect: Vertical viewpoint dependence of facial expression perception" *Proceedings of the Royal Society of London, B* **267** 2239–2245
- Macmillan N A, Creelman C D, 1991 *Detection Theory: A User's Guide* (New York: Cambridge University Press)
- Maurer D, Le Grand R, Mondloch C J, 2002 "The many faces of configural processing" *Trends in Cognitive Sciences* **6** 255–260
- Mignault A, Chaudhuri A, 2003 "The many faces of a neutral face: Head tilt and the perception of dominance and emotion" *Journal of Nonverbal Behavior* **27** 111–132
- O'Toole A J, Edelman S, Bülthoff H H, 1998 "Stimulus-specific effects in face recognition over changes in viewpoint" *Vision Research* **38** 2351–2363
- Otta E, Lira B B P, Delevati N M, Cesar O P, Pires C S G, 1994 "The effect of smiling and of head tilting on person perception" *Journal of Psychology* **128** 323–331
- Perrett D I, Hietanen J K, Oram M W, Benson P J, 1992 "Organisation and functions of cells responsive to faces in the temporal cortex" *Philosophical Transactions of the Royal Society of London, B* **335** 23–30
- Perrett D I, Mistlin A J, Chitty A J, 1987 "Visual neurones responsive to faces" *Trends in Neurosciences* **10** 358–364
- Perrett D I, Oram M W, Ashbridge E, 1998 "Evidence accumulation in cell populations responsive to faces: an account of generalisation of recognition without mental transformations" *Cognition* **67** 111–145
- Perrett D I, Smith P A J, Potter D D, Mistlin A J, Head A S, Milner A D, Jeeves M A, 1985 "Visual cells in the temporal cortex sensitive to face view and gaze direction" *Proceedings of the Royal Society of London, B* **223** 293–317
- Rakover S S, 2002 "Featural vs. configural information in faces: A conceptual and empirical analysis" *British Journal of Psychology* **93** 1–30

-
- Rhodes G, Brake S, Atkinson A P, 1993 "What's lost in inverted faces?" *Cognition* **47** 25–57
- Schwaninger A, Lobmaier J S, Collishaw S M, 2002 "Role of featural and configural information in familiar and unfamiliar face recognition" *Lecture Notes in Computer Science* **2525** 643–650
- Schyns P G, Bülthoff H H, 1994 "Viewpoint dependence and face recognition", in *Proceedings of the XVI Meeting of the Cognitive Science Society* Eds A Ram, K Eiselt (Hillsdale, NJ: Lawrence Erlbaum Associates) pp 789–793
- 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
- Tarr M J, Bülthoff H H, 1998 "Image-based object recognition in man, monkey and machine" *Cognition* **67** 1–20
- Troje N F, Bülthoff H H, 1996 "Face recognition under varying poses: The role of texture and shape" *Vision Research* **36** 1761–1771
- Troje N F, Bülthoff H H, 1998 "How is bilateral symmetry of human faces used for recognition of novel views?" *Vision Research* **38** 79–89
- Troje N F, Kersten D, 1999 "Viewpoint-dependent recognition of familiar faces" *Perception* **28** 483–487
- Valentin D, Abdi H, Edelman B, 1997 "What represents a face? A computational approach for the integration of physiological and psychological data" *Perception* **26** 1271–1288
- Valentine T, Bruce V, 1988 "Mental rotation of faces" *Memory & Cognition* **16** 556–566
- Wallraven C, Schwaninger A, Schuhmacher S, Bülthoff H H, 2002 "View-based recognition of faces in man and machine: re-visiting inter-extra-ortho" *Lecture Notes in Computer Science* **2525** 651–660

ISSN 0301-0066 (print)

ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 36 2007

www.perceptionweb.com

Conditions of use. This article may be downloaded from the Perception website for personal research by members of subscribing organisations. Authors are entitled to distribute their own article (in printed form or by e-mail) to up to 50 people. This PDF may not be placed on any website (or other online distribution system) without permission of the publisher.