

The Derivation of Orientation-Invariant Shape Representation in Visual Object Recognition.

E.C.Leek, S.J.Johnston,C.J.Atherton, N. A. Thacker, S.Danziger and A.Jackson.

Last updated
11 /4 / 2005



Imaging Science and Biomedical Engineering Division,
Medical School, University of Manchester,
Stopford Building, Oxford Road,
Manchester, M13 9PT.

The derivation of orientation-invariant shape representations
in visual object recognition

E.C.Leek¹, S.J.Johnston¹, C.J. Atherton⁴, N.A.Thacker², S. Danziger³ & A.Jackson²

¹Centre for Cognitive Neuroscience, School of Psychology, University of Wales, Bangor, UK.

² Imaging Science and Biomedical Engineering Department, University of Manchester, UK.

³Ben Gurion University, Israel.

⁴University of Central Lancashire, Preston, UK.

RUNNING HEAD: ORIENTATION-INVARIANT OBJECT RECOGNITION

Address for correspondence

Charles Leek PhD, Centre for Cognitive Neuroscience, School of Psychology, University of Wales,

Bangor, Gwynedd, UK, e-mail: e.c.leek@bangor.ac.uk. Tel (+44) 1248 382948. The work reported in

this paper was supported by project grant (5/S15963) from the Biotechnology and Biological Sciences

Research Council. UK (BBSRC).

ABSTRACT

While previous studies suggest that the recognition of misoriented objects may be orientation-dependent, or orientation-invariant, the functional independence of the mechanisms underlying these contrasting patterns of data remains unclear. In particular, it has been widely reported that orientation-invariant performance only emerges following spatial normalisation of misoriented objects on early blocks of trials, suggesting that the computation of orientation-invariant shape representations is not functionally independent of orientation-dependent processes. This issue is examined in two experiments contrasting recognition latencies, across blocks of trials, for symmetrical and asymmetrical 2D novel forms have previously been shown to elicit orientation-invariant performance in shape recognition tasks. The results show that orientation-invariance, with both stimulus types, does not depend on spatial normalisation of misoriented stimuli on early blocks of trials. In contrast to some recent claims, this finding shows that there are some kinds of orientation-invariant shape representations that can be computed independently of spatial normalisation mechanisms.

One fundamental issue in research on human visual object recognition concerns how we are able to identify objects across changes in stimulus orientation. The primary source of evidence about the cognitive mechanisms underlying this ability comes from studies of misoriented object recognition (e.g., Arguin & Leek, 1995; Arguin & Leek, in press; Biederman & Gerhardstein, 1993; Jolicoeur, 1985; Jolicoeur, 1990; Jolicoeur & Milliken, 1989; Jolicoeur & Humphrey, 1998; Lawson, 1999; Leek, 1998a; 1998b; Maki, 1986; McMullen & Farah, 1991; Murray, 1999; Rock, 1973; Takano, 1989; Tarr, Williams, Haywood & Gauthier, 1998; Tarr & Pinker, 1989; 1990; Tarr & Bülthoff, 1998). Numerous studies have shown that the time taken to recognise objects can be sensitive to stimulus orientation. Typically, response times (RTs), in shape recognition tasks, increase as a function of the angular disparity between stimulus orientation and an object's 'upright' or most familiar orientation/s (e.g., Jolicoeur, 1985; Jolicoeur & Milliken, 1989; Jolicoeur & Humphrey, 1998; Lawson, 1999; Leek, 1998b; Maki, 1986; McMullen & Farah, 1991; Murray, 1999; Rock, 1973; Tarr, 1995; Tarr & Pinker, 1989; 1990; Tarr & Bülthoff, 1998). This finding suggests that, at least under certain circumstances, the recognition of misoriented objects involves some form of spatial normalisation. Several hypotheses about the nature of this normalisation process have been advanced. These include accounts based on analogue 'mental rotation' (e.g., Jolicoeur, 1985; 1990; Rock, 1973; Shepherd & Metzler, 1971; Tarr & Pinker, 1989)¹, view interpolation (e.g., Edelman & Bülthoff, 1992), the response

¹ Some recent evidence has challenged this proposal – including data showing that normalisation rates are sometimes non-linear which is not clearly predicted by analogue mental rotation (e.g., Lawson & Jolicoeur, 1999; see also Lawson, 1999), and from case studies of neurological patients (e.g., Farah & Hammond, 1988; Turnbull & McCarthy, 1996), and functional brain imaging (e.g., Gauthier et al., 2002) reporting dissociations between mental rotation and misoriented object recognition. For this reason, in this paper we refer only to the more

properties (i.e., tuning functions) of neural population vectors (e.g., Perrett, Oram & Ashbridge, 1998), and the incremental spread of activation over networks of orientation (or viewpoint)-dependent shape descriptions (e.g., Edelman and Weinshall; 1991; 1998).

In contrast, there is also evidence that visual recognition can sometimes be orientation-invariant: that is, some objects can be identified equally quickly regardless of stimulus orientation (Biederman & Bar, 1999; Biederman & Gerhardstein, 1993; Cohen & Kubovy, 1993; Corballis, Zbrodoff, Shetzer & Butler, 1978; Corballis & Nagourney, 1978; DeCaro & Reeves, 2002; Eley, 1982; Farah, Rochlin & Kline, 1994; Leek, 1998a, McKone & Grenfell, 1999; McMullen & Farah, 1991; Murray, Jolicoeur, McMullen & Ingleton, 1993; Takano, 1989; Tarr & Pinker, 1990; Wiser, 1981). For example, in Tarr and Pinker (1990) participants memorised sets of two-dimensional novel shapes at a single image plane orientation. Recognition memory was later tested by presenting the memorised stimuli for identification, along with visually similar distracters, at both the practiced as well as several unfamiliar 'test' orientations. The results showed an orientation-dependent (i.e., monotonically increasing) pattern of RTs for only one of the four sets of stimuli. In order to account for these results, Tarr and Pinker (1990) argued that stimuli showing orientation-invariant performance could each be uniquely identified in terms of a one-dimensional object-centred shape description; that is, a representation that specifies the relative ordering of shape features along a single internal shape axis. In contrast, stimuli showing orientation-dependent performance could only be discriminated from each other through the specification of their feature configuration along two-dimensions (that is, in terms of x and y co-ordinates). Other studies have suggested orientation-invariance might also be achieved through the identification of local, so-called 'free floating' (Jolicoeur, 1990) or 'orientation-free' (Takano, 1989), features that serve to define object identity regardless of global shape orientation (e.g., Biederman & Bar, 1999; Jolicoeur, 1990; Jolicoeur & Milliken, 1989; Jolicoeur & Humphrey, 1998; Just & Carpenter, 1985; Murray et al., 1993; Takano, 1989). These features might include invariant properties of edges, vertices, and their local

configuration (e.g., Biederman & Bar, 1999; Takano, 1989; Thacker, Riocreux & Yates, 1994)².

However, the exact conditions under which recognition is orientation-dependent, or orientation-invariant, seem to depend on a variety of stimulus and task variables (e.g., Dickerson & Humphreys, 1999; Hamm & McMullen, 1998; Jolicoeur & Milliken, 1989; Leek et al., 1998a; Vanrie, Willems & Wagemans, 2001).

One unresolved issue is whether orientation-invariance can be mediated by processes that are functionally distinct from those involved in orientation-dependent spatial normalisation. It is clear that orientation-invariant performance, under some conditions, can be accounted for within the context of orientation-dependent theories of shape recognition (e.g., Tarr & Pinker, 1989; Tarr, 1990; Tarr & Bülthoff, 1998). For example, one well-documented finding is that orientation effects often diminish with practice, resulting in orientation-invariant performance on later blocks of trials (e.g., Jolicoeur, 1985; Jolicoeur & Humphrey, 1998; Lawson, 1999; Maki, 1986; Murray, 1999; Tarr & Pinker, 1989). This practice effect is consistent with the encoding of orientation (or viewpoint) specific representations across a range of stimulus orientations (e.g., Lawson & Jolicoeur, 1999; Leek, 1998a; 1998b; Tarr, 1995; Tarr & Pinker, 1989; Tarr & Bülthoff, 1998): if a sufficiently large number of different orientation-specific representations are encoded, recognition latencies may appear to be invariant to changes in stimulus orientation because the time taken to spatially normalise misoriented objects to the nearest stored shape description will be relatively constant or negligible. The hypothesis that the derivation of orientation-invariant shape representations is dependent on spatial normalisation mechanisms is also supported by findings from two other studies. Jolicoeur and Milliken (1989) examined the effects of naming upright, and rotated, objects on recognition latencies for the same and new objects in subsequent blocks. They found that effects of stimulus orientation diminished only when the same objects had been presented in non-upright orientations, or in the context of other misoriented stimuli, in earlier blocks of trials. Misoriented objects that had previously only been presented at upright orientations continued to undergo spatial normalisation on subsequent presentations. Additional supporting evidence has also been found by McMullen and

² It is also necessary to specify, on this account (although rarely done so) how the individual local features, or combinations of features, are identified independently of their orientation. One possibility is that they are

Farah (1993). They tested the generality of the findings of Tarr and Pinker (1990) in a task involving the recognition of line drawings of common objects. The results showed that while orientation effects for common objects with an axis of symmetry in the image plane do diminish more quickly than those for objects with no axis of symmetry, they still show orientation-dependent performance on early blocks of trials.

The data from these studies suggests that spatial normalisation, on initial presentations of objects, may be a prerequisite to the derivation of orientation-invariant shape descriptions. As such, they undermine the hypothesis that orientation-invariance may be mediated by processes that are functionally distinct from those involved in orientation-dependent spatial normalisation: If this were the case, it is not obvious why spatial normalisation should be required during the derivation of the orientation-invariant representations.

In this context, the generality of block, and context, effects (that is, orientation-invariant performance that is seemingly dependent either on spatial normalisation on early blocks of trials or on the prior presentation of objects in the context of misoriented forms) is of considerable theoretical interest. While the decline in orientation effects over blocks of trials has been reported in several studies (e.g., Jolicoeur, 1985; McMullen & Jolicoeur, 1992; Murray, 1999; Tarr & Pinker, 1989), the generality of block effects in the identification of other classes of object shapes that have been shown to elicit orientation-invariant performance has yet to be examined. In particular, both the studies of Jolicoeur and Milliken (1989), and of McMullen and Farah (1991), used only line drawings of common objects which, for the most part, are relatively rich in terms of their feature structure. Thus, the observation of block, and context, effects in those studies might reflect the operation of processes related to the identification of orientation-invariant shape features, rather than more generally to the encoding of other forms of orientation-invariant representation such as the axis-based global shape descriptions proposed by Tarr and Pinker (1989). If block effects do not generalise to other classes of stimuli, this would support the hypothesis that there some types of orientation-invariant mechanisms that are functionally distinct from those involved in orientation-dependent spatial normalisation.

We examined this issue in the current study. The goal of these experiments was to determine whether spatial normalisation also occurs on early blocks of trials with object shapes which, by hypothesis, can be encoded in orientation-invariant axis-based global shape representations. This was done using the same stimulus sets, and a modified version of the recognition-memory paradigm, reported in Tarr and Pinker (1990). During an initial learning phase, subjects were trained to recognise a target stimulus from a specific class of novel shapes. Objects were presented only at a single 'upright' orientation in the learning phase to avoid context effects influencing performance on subsequent blocks (Jolicoeur & Milliken, 1989). Recognition memory for the target shapes was later tested by presenting the target, and distracter, stimuli at the practiced orientation, as well as at several test orientations. Recognition latencies were then analysed, across blocks, to determine the time course involved in the derivation of object-centred representations. If spatial normalisation is indeed a general prerequisite for the derivation of orientation-invariant shape representations, then orientation-dependent performance should be found on early blocks of trials with both object sets. In contrast, if such representations can be computed using functionally independent mechanisms that do not involve spatial normalisation, then orientation-invariant performance should be found even on the initial block of test trials.

EXPERIMENT 1

METHOD

Participants

Participants were 14 undergraduate students from the University of Wales, Bangor. All had normal or corrected-to-normal eyesight. Informed consent was obtained from all subjects prior to testing³.

Materials & Apparatus

The stimuli were based on those used by Tarr and Pinker (1990). Set A consisted of seven asymmetrical novel shapes that previously elicited an orientation-dependent pattern of RTs in the Tarr and Pinker (1990) study (see Figure 1, Set A). Set B consisted of seven symmetrical objects that previously showed an orientation-invariant pattern (see Figure 1, Set B). Stimuli were scaled to fit within an 8x8 cm frame, which subtended 10.1° of visual angle from a viewing distance of 45 cm. Stimulus presentation was controlled by E-Prime software version 1.0 (Psychology Software Tools, Inc. Pittsburgh, USA.) running on a Pentium 4 PC with a 17" monitor operating at a resolution of 1024 x 768 pixels.

 Insert Figure 1 about here

Design & Procedure

A three-factor (Object set X Block X Stimulus orientation) repeated measures design was used. The experiment consisted of a separate training and test phase. In the training phase participants memorised one shape from each stimulus set. Each of the seven stimuli in each object set served as a target twice across subjects. They were asked to copy and redraw a stimulus from one of the novel object sets until this could be accomplished without error. During this phase the stimuli were always presented at the same single 'upright' orientation. At the 'upright' orientation the principal axis of each stimulus was aligned vertically with respect to the participants midline (in the drawing phase), or with the computer monitor (in subsequent phases of the experiment). The participants then

³ Pilot data showed that, at an individual subject level, not all participants show statistically reliable orientation-dependent effects with Set A stimuli. Since the aim of the current study is to examine performance on the first block of trials with orientation-invariant object sets, under conditions where reliable normalisation effects are shown with the contrasting object set, only subjects who showed normalisation rates of greater than 1 ms/deg

completed a short computerised training task. On each trial a single stimulus was presented at its upright orientation. The participants had to indicate whether or not the stimulus was the target object. There was a minimum of 12 target-present trials, and 6 distracter trials, maintaining a consistent 2:1 ratio in the practice and testing phases of the experiment. However, these trials continued until the subject reached a criterion level of accuracy of 80% correct. Once subjects had reached the criterion accuracy level they proceeded to the test phase of the experiment. Participants always completed both the learning and test phases with a target stimulus from the Set B block, prior to completing the learning and test phases with stimuli from Set A⁴.

In the test phase for both stimulus sets participants completed three 90 trial blocks. These blocks contained targets and distracters from only one of the two stimulus sets. In each block there were 60 target present trials, and 30 target absent trials (target-non target ratio 2:1). Target stimuli were presented 12 times at the upright zero degree orientation) and six times each at the following orientations: +30, +60, +90, +120, +240, +270, +300 and +330 degrees in the fronto-parallel plane. Distracters (n=6) were presented once each at the upright, and four times at test orientations. Each orientation was probed an equal number of times with distracters.

Participants were seated 45 cm in front of the monitor. Inclinations of the head during the experiment were prevented by a chinrest. Each trial began with 750 msec fixation (+) prompt appearing in the centre of the screen. The fixation was replaced by either a target or distracter stimulus shown at one of the tested stimulus orientations (ISI = 10 msec). The task was to indicate by a keypress (target, non-target) whether or not the stimulus was a previously memorised target object. The stimuli remained on the screen until the participant responded, or the reaction time (RT) exceeded 5000 msec. If the participant made an incorrect response, or did not respond before the

(e.g., Cohen & Kubovy, 1993; Tarr & Pinker, 1989 – see below) with Set A (blind to their performance with other object sets) were included in the current study.

⁴ A pilot study using the same stimuli showed an effect of task order. Participants who performed the task with the Set B stimuli first showed orientation-invariant performance with those objects, and orientation-dependent performance with Set A. However, participants who performed the task in the reverse order (Set A before Set B) showed orientation effects with both stimulus sets. This incidental finding is consistent with previous reports showing context dependent orientation effects (e.g., Jolicoeur & Milliken, 1989). For this reason, in the current study, only the context-independent task order was used: that is, Set B/C before Set A.

deadline, an error message was displayed. Trial order was randomised within blocks. Target present responses were always made with the dominant hand. The experiment lasted about 40 minutes.

RESULTS

All analyses were performed on RTs for correct responses in the test blocks of trials. RT and error data were collapsed across symmetrical orientations (i.e., +/-60, +/-120). RTs greater or less than three SDs from the mean in each condition were eliminated.

Analyses of block and orientation effects

Since target stimuli had been viewed more frequently at the zero degree upright orientation (learning plus test phases), analyses of orientation effects were conducted only across the range of stimulus orientations where frequency of presentation had been held constant (30-120 degrees). Figure 2 shows mean RTs as a function of Stimulus Orientation and Block for (a) Set A, and (b) Set B stimuli.

 Insert Figures 2(a) and 2(b) about here

A three-way ANOVA on RTs with Stimulus Set, Block and Orientation as within-subjects factors showed significant main effects of Block, $F(2, 26) = 11.55, p < .001$; and Orientation, $F(3, 39) = 14.74, p < .001$; and also significant two-way interactions between Stimulus Set and Orientation, $F(3, 39) = 5.64, p < .005$; and Block and Orientation, $F(6, 78) = 2.52, p < .05$. In addition, there was a significant three-way interaction between Stimulus Set, Block and Orientation, $F(6, 78) = 2.22, p < .05$.

These interactions were analysed on mean RTs separately for each stimulus set using two-way ANOVA with Block and Orientation as within-subjects factors. This analysis confirmed a significant Block x Orientation interaction for Set A, $F(6, 78) = 3.23, p < .01$, but not for Set B [$F(6, 78) = 0.5, n.s.$]. Crucially, for Set A, there was a significant effect of orientation in Block 1, $F(3, 39) = 10.42, p < .001$, but not for Set B [$F(3, 39) = 0.9, n.s.$]. This shows that orientation-invariant performance with Set B stimuli can be found without evidence of spatial normalisation on early blocks of trials. That is,

the distinction between orientation-dependent, and orientation-invariant processing, can be observed from the outset of testing.

These effects were further examined by calculating linear regression slopes of mean RTs against stimulus orientation across subjects for block and stimulus set (e.g., Tarr & Pinker, 1989). The slopes for Set A and Set B stimuli are shown in Figures 3(a) and 3(b).

 Insert Figures 3(a) and 3(b) about here

As Figures 3(a) and 3(b) show, there is a striking contrast in the patterns of regression slopes between Set A and Set B stimuli. Critically, in Block 1, Set A stimuli show a time cost associated with stimulus orientation equivalent to 2.37 ms/deg ($r^2 = 0.68$, $p < .01$), compared to 0.45 ms/deg ($r^2 = 0.48$, n.s.) for Set B. The mean slope for Set A is well within the range of slopes previously taken as evidence for spatial normalisation (e.g., 1.18 – 4.23 ms/deg range across subjects: Mean = 2.44 ms/deg; Exps 1 & 2; Cooper, 1975), while for Set B the mean slope is well outside of this range, and too shallow for spatial normalisation (e.g., Cohen & Kubovy, 1993; Cooper, 1975; Cooper & Shepard, 1973; Shepard & Metzler, 1988; Shepard & Cooper, 1983; Takano, 1989; Tarr & Pinker, 1989)⁵.

Regression slopes were analysed using a two-way ANOVA with Stimulus Set and Block as factors. There were significant main effects of Stimulus Set, $F(1, 13) = 17.99$, $p < .001$, and Block, $F(2, 26) = 6.36$, $p < .005$. The Stimulus Set X Block interaction was also significant, $F(2, 26) = 7.10$, $p < .005$. One-way ANOVA on the regression slope data showed a significant effect of Block for the Set A stimuli, $F(2, 26) = 15.94$, $p < .001$; but not for Set B [$F(2, 26) = 0.5$, n.s.]. Thus, these data suggest that while stimuli in Set A show patterns of RTs that increase with stimulus orientation (consistent with spatial normalisation), while the slopes decrease across blocks, Set B stimuli do not.

Error rates

⁵ Consistent with previous literature, we adopted a 1 ms/deg lower boundary of linear regression slopes for spatial normalization in the image plane (e.g. Cohen & Kubovy, 1993; Cooper, 1975; Cooper & Shepard, 1973; Shepard & Metzler, 1988; Shepard & Cooper, 1983; Takano, 1989; Tarr & Pinker, 1989). Thus, slopes of < 1 ms/deg are assumed to reflect orientation-invariant performance.

Error rates were very low (< 3% of trials in any condition). For Set A, mean error rates per block were: 2.6% (Block 1), 1.6% (Block 2) and 1.7% (Block 3). For Set B, mean error rates were: 2.8% (Block 1), 1.9% (Block 2) and 2.2% (Block 3). Three-way ANOVA showed no significant effects of Stimulus Set, Block or Orientation, and no significant interactions.

Discussion

The main findings of Experiment 1 were as follows: First, the data replicate, within-subjects, the dissociation between orientation-dependent (Set A) and orientation-invariant (Set B) recognition found by Tarr and Pinker (1990). Second, orientation-invariant performance with symmetrical objects (Set B) was not associated with spatial normalisation of early blocks of trials. Rather, recognition of these stimuli showed an orientation-invariant pattern from the outset of the test phase. This finding is inconsistent with the results of McMullen and Farah (1991) who found that orientation-invariance with line drawings of symmetrical common objects occurred, but only after spatial normalisation on an initial block of trials. Moreover, in contrast to Jolicoeur and Milliken (1989) this orientation-invariant pattern, in the test phase, was not dependent on the objects being presented in the context of misoriented forms on an earlier block of trials: in the training phase, all stimuli were only presented at a single upright orientation. In addition, the test blocks with Set B stimuli were completed prior to Set A stimuli. These results provide a boundary condition on the observation of block effects in the derivation of orientation-invariant shape representations, and provide some evidence for functionally distinct orientation-invariant object constancy processes in recognition.

The results of Experiment 1 motivate further investigation of the factors underlying the absence of block effects in the derivation of orientation-invariant representations. The absence of these effects might be restricted to objects that contain an unambiguous axis of symmetry, rather than more generally to forms that can be encoded in a one-dimensional object-centred coordinate system regardless of symmetry (Tarr & Pinker, 1990). This issue was examined in Experiment 2 using asymmetrical novel shapes.

EXPERIMENT 2

In Experiment 2 we contrasted recognition performance with the same orientation-dependent object set (Set A) used in Experiment 1, with a set of asymmetrical orientation-invariant objects (Set C) from the Tarr and Pinker (1990) object set. If the absence of block effects with orientation-invariant stimuli found in Experiment 1 is restricted to stimuli with an unambiguous axis of symmetry (McMullen & Farah, 1991), then spatial normalisation should be found on early blocks of trials during the derivation of orientation-invariant representations of the asymmetrical objects.

METHOD

Participants

Participants were 14 undergraduate students from the University of Wales, Bangor. All had normal or corrected-to-normal eyesight. Informed consent was obtained from all subjects prior to testing. None of the participants had taken part in Experiment 1.

Materials/Apparatus

The stimuli consisted of line drawings of two sets of novel objects from the original stimulus sets developed by Tarr and Pinker (1990). Set A consisted of seven asymmetrical novel shapes that previously elicited an orientation-dependent pattern of RTs in Experiment 1 and in Tarr and Pinker (1990) – see Figure 1, Set C. The second set consisted of seven asymmetrical objects that previously showed an orientation-invariant pattern in Tarr and Pinker (1990) (see Figure 1, Set C). Stimulus dimensions were identical to those in Experiment 1.

Design/Procedure

The design and Procedure were identical to Experiment 1.

RESULTS

All analyses were performed on RTs for correct responses in the test blocks of trials. RT and error data were collapsed across symmetrical orientations (i.e., +/-60, +/-120). RTs greater or less than three SDs from the mean in each condition were eliminated.

Analyses of block and orientation effects

Figure 4 shows mean RTs as a function of Stimulus Orientation (within the range of 30-120 degrees) and Block for (a) Set A, and (b) Set C stimuli.

 Insert Figures 4(a) and 4(b) about here

A three-way ANOVA on RTs with Stimulus Set, Block and Orientation as within-subjects factors showed significant main effects of Block, $F(2, 26) = 43.88, p < .001$; and Orientation, $F(3, 39) = 11.94, p < .001$; and also a significant two-way interaction between Stimulus Set and Orientation, $F(3, 39) = 2.42, p < .05$.

These effects were further examined on mean RTs separately for each stimulus set. As in Experiment 1, for Set A, there was a significant effect of orientation in Block 1, $F(3, 39) = 9.08, p < .001$. However, even at Block 1, there was significant orientation effect for Set C [$F(3, 39) = 1.28, n.s.$].

In order to confirm these findings orientation effects were also examined by calculating linear regression slopes of mean RTs against stimulus orientation across subjects for block and stimulus set (e.g., Tarr & Pinker, 1989). The slopes for Set A and Set C stimuli are shown in Figures 5(a) and 5(b).

 Insert Figures 5(a) and 5(b) about here

In Block 1, Set A stimuli show a time cost associated with stimulus orientation equivalent comparable to that found in Experiment 1: 2.33 ms/deg ($r^2 = 0.54, p < .05$), compared to 0.64 ms/deg ($r^2 = 0.38, n.s.$) for Set C. As was previously found in Block 1 for Set B, Set C showed a slope that would be considered too shallow for spatial normalisation (e.g., Cohen & Kubovy, 1993; Takano, 1989). Regression slopes were analysed using a two-way ANOVA with Stimulus Set and Block as factors. There was significant main effect of Stimulus Set, $F(1, 13) = 8.87, p < .01$, and a significant Stimulus Set X Block interaction, $F(2, 26) = 2.62, p < .05$. Further analysis, using One-way ANOVA,

confirmed that there was a significant effect of Block for the Set A stimuli, $F(2, 26) = 4.5$, $p < .05$; but not for Set C [$F(2, 26) = 0.1$, n.s.].

Error rates

Error rates were also low in Experiment 2 (< 3% of trials in any condition). For Set A, mean error rates per block were: 1.0% (Block 1), .11% (Block 2) and 1.3% (Block 3). For Set B, mean error rates were: 1.5% (Block 1), 0.4% (Block 2) and 1.0% (Block 3). Three-way ANOVA showed no significant effects of Stimulus Set, Block or Orientation, and no significant interactions.

Discussion

The results of Experiment 2 showed that Set A stimuli showed orientation-dependent performance that was strongest in Block 1, which diminished across subsequent blocks. In contrast, as was previously found for Set B stimuli in Experiment 1, Set C objects showed orientation invariant performance from Block 1. These findings extend those of Experiment 1, and show that orientation-invariant performance can also be observed with asymmetrical object shapes from the outset of the test phase.

GENERAL DISCUSSION

The aim of this study was to examine the generality of block effects in the encoding of orientation-invariant shape representations. In particular, we examined whether block effects would be found with classes of visual stimuli which, by hypothesis, can be encoded in one-dimensional object-centred global shape representations (Tarr & Pinker, 1990). The main findings of these experiments can be summarised as follows:

- In Experiments 1 and 2 we replicated, using a within-subjects design, the dissociation between orientation-dependent, and orientation-invariant, recognition for symmetrical and asymmetrical forms previously reported by Tarr and Pinker (1990).

- Analyses of RTs, and linear regression slopes, across blocks of trials showed that the orientation-dependent stimulus set (Set A) showed strong orientation effects in Block 1 that diminished across subsequent blocks. In contrast, both the symmetrical (Set B) and asymmetrical (Set C) object sets showed a pattern consistent with orientation-invariant performance from the outset of testing; that is, from Block 1.

To our knowledge, this is one of the first demonstrations of orientation-invariance in a recognition task where it has been shown that invariant performance, for image plane rotated stimuli, can be achieved from the outset of testing, rather than following spatial normalisation on early blocks of trials. This finding contrasts with those from some previous studies. McMullen and Farah (1991) reported rapid orientation-invariant performance for symmetrical, but not asymmetrical common objects, but only after spatial normalisation of misoriented forms on an initial block of trials. In addition, unlike in Jolicoeur and Milliken (1989), orientation-invariant performance in the present study was not dependent on the objects being presented in the context of misoriented forms on an earlier block of trials: in the training phase, all stimuli were presented only at a single orientation, and all test blocks with orientation-invariant stimuli were completed prior to blocks with orientation-dependent objects.

These data provide new constraints on hypotheses about the representations, and about the organisation of the functional processes, that mediate object constancy. They also have implications for our understanding of the origin of orientation effects in visual recognition tasks. We consider these issues in turn.

The derivation of orientation-invariant shape representations in object constancy

In the first place, the current data provide a new boundary condition on the block and context effects, reported in previous studies, which have been associated with the derivation of orientation-invariant shape representations. The absence of these effects, for both symmetrical and asymmetrical forms, provides support for the hypothesis that object constancy may be mediated by orientation-invariant mechanisms that are functionally distinct from those involved in spatial normalisation: That

is, the visual system is seemingly able to derive at least some forms of orientation-invariant representation without spatial normalisation. At the most general level, this finding adds to a growing body of research showing that object recognition may be either orientation-dependent or orientation-invariant depending on a variety of stimulus and task variables (e.g., Dickerson & Humphreys, 1999; Hamm & McMullen, 1998; Jolicoeur & Milliken, 1989; Leek et al., 1998a; Vanrie et al., 2001).

At a more specific level, the current findings challenge some hypotheses about the functional organisation of orientation-dependent, and orientation-invariant, processes in object constancy. Several researchers have proposed the existence of multiple routes to object constancy that include both orientation-dependent, and orientation-invariant, processes (e.g., Humphreys & Riddoch, 1984; Jolicoeur, 1990), but the relations between the proposed routes have received relatively little attention. Consider three possibilities. First, an orientation-invariant process may be the default recognition strategy, with spatial normalisation (or some other orientation-dependent mechanism) only employed when orientation-invariant processes have failed to result in unambiguous stimulus identification (e.g., Corballis, 1988; Turnbull, Carey & McCarthy, 1997). Although this hypothesis holds some computational appeal – on the assumption that invariant recognition is less time-consuming than spatial normalisation – it is challenged by the observation, noted already, of block effects in some studies (e.g., Jolicoeur & Milliken, 1989; McMullen & Farah, 1991): if orientation-invariance were the default strategy, it is not obvious why spatial normalisation should ever be found on early blocks of trials prior to the emergence of orientation-invariant performance. This hypothesis is also undermined by data showing that orientation dependent effects, in recognition tasks, are the rule, rather than the exception (for recent reviews, see Lawson & Humphreys, 1998; Lawson, 1999).

A second, contrasting, hypothesis is that orientation-dependent recognition is default (e.g., Tarr, 1995), and invariant mechanisms are only be employed if the default strategy fails to result in stimulus identification. While this hypothesis is consistent with evidence showing orientation effects in most visual recognition tasks, it is challenged by the current data. If orientation-dependence were the default strategy, it is not clear why orientation-invariant performance should ever be observed from the outset of stimulus presentation.

A third possibility is that both orientation-dependent, and at least some orientation-invariant, mechanisms, operate in parallel on any given input image (Jolicoeur, 1990). This possibility is consistent with the current data, and the findings of other studies. As noted earlier, there is evidence that orientation-invariance may result from the encoding of multiple orientation-specific representations (e.g., Tarr & Pinker, 1989), or from the identification, and representation, of image features that are invariant to changes in stimulus orientation (e.g., Murray et al., 1993). Both of these routes to invariance seem to depend on the repeated presentation of objects and on spatial normalization of misoriented forms on early blocks of trials (Tarr & Pinker, 1989; Murray et al., 1993). In contrast, as suggested by the present findings, some forms of orientation-invariant representation can be computed independently of spatial normalization processes. The functional mechanisms that support the computation of these orientation-invariant representations could operate in parallel with, but independently of, other processes that underlie the spatial normalization of misoriented forms. The dual operation of orientation-dependent, and orientation-invariant, mechanisms also has computational appeal. It avoids the problem of specifying how the visual system could select, a priori, a particular object constancy mechanism to process a given input image. Instead, all stimuli would be processed by both orientation-dependent, and orientation-invariant, mechanisms in parallel. Whether orientation-dependent, or orientation-invariant, performance results as a behavioural outcome would depend on several stimulus and task variables that differentially affect the operation of the underlying object constancy mechanisms. Some of these factors have already been described in the literature, including: (i) the geometric complexity of the stimuli (e.g., Bethel-Fox & Shepard, 1988; Cohen & Kubovy, 1993; Leek, 1998a; Takano, 1989; Tarr & Pinker, 1990; Wisner, 1981), (ii), stimulus familiarity (e.g., Bethel-Fox & Shepard, 1988; Leek, 1998a) and (iii) task context (e.g., Dickerson & Humphreys, 1999; Hamm & McMullen, 1998; Jolicoeur & Milliken, 1989; Schyns, 1998; Shepard & Metzler, 1988).

Axis-based global shape representations

What kinds of representations are these functionally independent orientation-invariant mechanisms likely to compute? Consistent with the hypothesis of Tarr and Pinker (1990), we might

suppose that they consist of one-dimensional global axis-based shape descriptions. On this account, an orientation-invariant shape description can be generated by encoding the ordering of shape features along a single object-internal global axis of elongation or symmetry. Arguably, such an account provides a potential explanation for why block, and context, effects have been reported in previous studies using other kinds of objects. The studies of Jolicoeur and Milliken (1989), and of McMullen and Farah (1991), both involved the image-plane rotation of featurally complex line drawings of 3D common objects. The complexity of many common objects could preclude their encoding in the kinds of one dimensional global axis-based representations described by Tarr and Pinker (1990) because their shape configuration must be specified in at least two dimensions. Instead, for these objects, orientation-invariance may be mediated by a different form of representation, for example, based on multiple orientation-specific representations (e.g., Tarr & Pinker, 1989), or on the identification of 'free floating' shape features, both of which can only be encoded following spatial normalisation on early blocks of trials (Jolicoeur & Milliken, 1989; McMullen & Farah, 1991).

However, there are also reasons to doubt that global axis-based shape representations underlie orientation-invariant recognition in the current study. Although rarely discussed in the object recognition literature, there are grounds for arguing that global axis-based object-centred representations are not orientation-invariant as is widely assumed. This counter-argument is based on a consideration of the way in which such representations encode information about the spatial configuration of objects features. In particular, the vectors assigned to locations along the principal axis (which specify the relative ordering of object features) are necessarily defined in relation to polar co-ordinates; that is, the locations of features relative to end points on the axis (Hinton & Parsons, 1981; Marr & Nishihara, 1978; Pinker, 1984). Indeed, as discussed by Marr (1982), such intrinsic polarity may usefully serve to define the canonical orientation of shape representations. However, for recognition to occur, the polarity of a stored shape description (that is, the polar co-ordinates) must be matched to similar co-coordinates in the representation computed from the input image. Thus, recognition mediated by global axis-based shape descriptions will only be invariant to changes in stimulus orientation if, among other things, the mechanisms involved in matching the polarity of the representations are themselves invariant. Since the polarity of the representation is necessarily

changed by rotation around the image plane (i.e., by a reversal of top-bottom relations) this is unlikely.

There are also other computationally motivated arguments against the use of global axis-based shape representations in recognition. One general, and frequently discussed, problem is the difficulty of reliably extracting global shape axes from input images of depth rotated forms because of foreshortening and occlusion (e.g., Humphreys & Riddoch, 1984; Humphrey & Jolicoeur, 1993; Marr, 1982; Pinker, 1984; Wechsler, 1990). For this reason other forms of 'orientation-invariant' representation should also be considered (e.g., Edelman & Weinshall, 1998; Lowe, 1987; Thacker, Riocreux & Yates, 1994; Wechsler, 1990). One candidate, recently discussed in the machine vision literature, is the pairwise geometric histogram, which encodes object shape in terms of the configuration of image feature (e.g., edges) pairs in relation to a locally defined reference axis (Thacker et al., 1994). Such representations are robust to changes in image plane stimulus orientation, as well as partial shape occlusion, because they are not based on the recovery of global shape properties from the input image such as principal axes of elongation or symmetry.⁶

The origin of orientation effects in visual recognition

The current findings are not only relevant to theories about the nature of orientation-invariance. They also provide constraints on hypotheses about the origins of orientation effects in visual recognition. While it has been most widely assumed that orientation effects reflect spatial normalisation processes that occur during perception (e.g., Arguin & Leek, in press; Jolicoeur, 1985; Leek, 1998a; 1998b; Tarr & Pinker, 1989; 1990), an alternative view is that these effects actually reflect processes occurring at a post-recognition stage; that is, following access to a long-term memory representation of object shapes (Corballis, 1988; DeCaro & Reeves, 2000; 2002). Most recently, DeCaro

⁶ A further relevant observation, in relation to the role of global shape axes in recognition, stems from the finding that orientation-invariant performance can be observed from the outset of testing. This suggests that orientation effects are not related to the identification of internal global shape axes per se: both of the orientation-invariant stimulus sets used in the current study possessed an unambiguous internal shape axis (and in one case, an axis of symmetry), but orientation effects were not observed during the initial blocks of trials. However, this does not mean that global shape axes play no role in shape representation. They may help define stimulus orientation (e.g., DeCaro & Reeves, 2002), or the canonical orientation of objects that undergo spatial normalisation during recognition (Marr, 1982). They may also facilitate the localisation of diagnostic shape features for use on subsequent presentations of the same objects (Jolicoeur, 1990).

& Reeves (2000; 2002) have proposed that orientation effects reflect a post-recognition process related to the verification of object orientation, and not object identity. The current data challenge this view. If the spatial normalisation process implied by orientation effects relates to the verification of object orientation, rather than shape recognition, then these effects should always be found – at least during the initial block of trials - in tasks involving judgements about misoriented forms. The current data show that this is not the case. Converging evidence also come from some other studies which challenge the post-recognition verification hypothesis. For example, Leek (1998a) contrasted orientation effects for misoriented mono- and polyoriented common objects (i.e., objects with a single predominant orientation in the environment and those that are frequently seen at a variety of different orientations). While mono-oriented forms show spatial normalisation effects consistent with those previously reported in the literature (e.g., Jolicoeur, 1985), polyoriented forms can be identified equally quickly regardless of stimulus orientation. If orientation effects are the consequence of a post-recognition verification of object orientation (independent of stimulus identity), then comparable orientation effects should have been found for mono- and polyoriented objects. Thus, the current findings, and those of Leek (1998a), suggest that, when orientation effects are observed, they result from the operation of spatial normalisation processes that are integral to the recognition process.

Orientation-invariance and the basic- versus subordinate-level category distinction

One final issue raised by the current data relates to the distinction between basic, and subordinate, level categorisation previously made in relation to orientation effects in visual recognition. It has been claimed, by some authors (e.g., Dickerson & Humphreys, 1999; Hamm & McMullen, 1998), that orientation effects arise from categorisation at a subordinate level of orientation-dependent shape representation where distinctions are made among specific exemplars of a basic (or entry) level category. In contrast, basic, or superordinate, levels of representation are orientation-invariant. Thus, tasks that require only recognition at these levels will not show orientation effects for misoriented stimuli. Some evidence for this claim comes from data reported by Hamm and McMullen (1998; see also Dickerson & Humphreys, 1999) showing strong orientation-dependent effects when subjects are required to match subordinate level names (e.g., collie dog) to

misoriented pictures of objects, irrespective of the level of categorisation involved in mismatch trials. But when the task involved matching at the basic (e.g., dog) or superordinate (e.g., animal) levels, orientation-invariant (or weaker orientation-dependent – Dickerson & Humphreys, 1999) performance was found. However, the current data challenge the view that subordinate level categorisation is necessarily orientation-dependent. The data from stimulus sets B and C (Experiments 1 and 2) show that orientation-invariant performance can also be found from the outset of a task even when subordinate level discrimination is required – that is, when subjects must make within-category discriminations among visually similar objects. This finding undermines an account of orientation-dependence, and orientation-invariance, in shape recognition solely in terms of the distinction between basic and subordinate levels of stimulus categorisation.

Summary

The results of two experiments show that orientation-invariant performance can be found with both symmetrical, and asymmetrical, 2D novel forms without spatial normalisation on early blocks of trials (e.g., Jolicoeur, 1985; Tarr & Pinker, 1989), or following the presentation of the stimuli in prior contexts of misoriented shape recognition (Jolicoeur & Milliken, 1989; McMullen & Farah, 1991). These findings suggest that the derivation of orientation-invariant shape representations need not rely on spatial normalisation (or the encoding of orientation-specific shape descriptions), but rather, may be mediated by cognitive processes that are independent of the spatial normalisation mechanisms presumed to underlie orientation-dependent recognition.

The current findings can be accounted for in terms of a model of object constancy which assumes the parallel operation of partially independent orientation-dependent, and orientation-invariant, object constancy mechanisms. However, it is unlikely that invariance is mediated by object-centred global shape representations of the kind proposed in some earlier studies (e.g., Tarr & Pinker, 1990) because such representations necessarily encode orientation-dependent information. Rather, orientation-invariance may be based on the encoding of shape representations that specify local, but not global, configural relations among shape features (e.g., Thacker et al., 1994).

Finally, the current data also challenge two recent hypotheses about the origin of orientation effects in visual recognition. First, the finding that orientation effects are not found on initial blocks of trials with certain stimulus types undermines the hypothesis that orientation effects reflect only post-recognition verification of object orientation, rather than the operation of spatial normalisation processes that occur during shape perception. If this were the case, then comparable effects of orientation should be found on the first block of trials with all classes of misoriented stimuli. The current data show that this is not the case. Second, orientation-invariant performance from the outset of testing, and under conditions where subjects had to discriminate targets from visually similar distracters, undermines the claim that orientation effects necessarily reflect access to subordinate level shape representations: although subordinate level discriminate was required by the current task, orientation effects were not found.

REFERENCES

Arguin, M. & Leek, E.C. (in press). Orientation-dependence in visual object recognition: Further evidence from repetition priming and picture naming. Perception & Psychophysics.

Arguin, M., & Leek, E.C. (1995) Mental Rotation for Visual Object Recognition: Identical Rates in Priming and Picture Naming. Investigative Ophthalmology & Visual Science, 36(4), 473.

- Bethel-Fox, C.E. & Shepard, R.N. (1988). Mental rotation: Effects of stimulus complexity and familiarity. Journal of Experimental Psychology: Human Perception and Performance, *14*, 12-23.
- Biederman, I. & Bar, M. (1999). One-shot viewpoint invariance in matching novel objects. Vision Research, *39*, 2885-2899.
- Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. Journal of Experimental Psychology: Human Perception and Performance, *19*, 6, 1162-1182.
- Chun, M.M. & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. Cognitive Psychology, *36*, 28-71.
- Cohen, D. & Kubovy, M. (1993). Mental rotation, mental representation, and flat slopes. Cognitive Psychology, *25*, 351-382.
- Cooper, L.A. (1975). Mental rotation of random two-dimensional shapes. Cognitive Psychology, *7*, 20-43.
- Cooper, L.A. & Shepard, R.N. (1973). Chronometric studies of the rotation of mental images. In W.G. Chase (Ed.). Visual information processing. New York: Academic Press.
- Corballis, M. C. (1988). Recognition of Disoriented Shapes. Psychological Review, *95*, 115-123.
- Corballis, M.C., Zbrodoff, N.J., Shetzer, L.I. & Butler, P.B. (1978). Decisions about identity and orientation of rotated letters and digits. Memory and Cognition, *6*, 98-107.
- Corballis, M.C. & Nagourney, B.A. (1978). Latency to categorize disoriented alphanumeric characters as letters or digits. Canadian Journal of Psychology, *32*, 186-188.
- DeCaro, S.A. & Reeves, A. (2002). The use of word-picture verification to study entry-level object recognition: Further support for view-invariant mechanisms. Memory and Cognition, *30*, 811-821.
- DeCaro, S.A. & Reeves, A. (2000). Rotating objects to determine orientation, not identity: Evidence from backward-masking/ dual-task procedure. Perception and Psychophysics, *62*, 1356-1366.
- Dickerson, J. & Humphreys, G.W. (1999). On the identification of misoriented objects: Effects of task and level of stimulus description. European Journal of Cognitive Psychology, *11*, 145-166.
- Edelman, S. & Weinshall, D. (1991). A self organising multiple-view representation of 3D objects. Biological Cybernetics, *64*, 209-219.
- Edelman, S. & Weinshall, D. (1998). Computational approaches to shape constancy. In V. Walsh & J. Kulikowski (Eds.), Perceptual constancy. Cambridge University Press. Cambridge.
- Edelman, S., & Bühlhoff, H. H. (1992). Orientation-dependence in the recognition of familiar and novel views of 3D objects. Vision Research, *32*, 2385-2400.
- Eley, M.G. (1982). Identifying rotated letter-like symbols. Memory and Cognition, *10*, 25-32.
- Farah, M.J. & Hammond, K.M. (1988). Mental rotation and orientation-invariant object recognition: Dissociable processes. Cognition, *29*, 29-46.

Gauthier, I., Haywood, W.G., Tarr, M.J., Anderson, A.W., Skudlarski, P. & Gore, J.C. (2002). BOLD activity during mental rotation and viewpoint-dependent object recognition. Neuron, 34, 161-171.

Hamm, J. P., & McMullen, P. A. (1998). Effects of orientation on the identification of rotated objects depend on the level of identity. Journal of Experimental Psychology: Human Perception and Performance, 24, 413-426.

Heil, M., Roesler, F., Link, M. & Bajric, J. (1998). What is improved if a mental rotation task is repeated – the efficiency of memory access, or the speed of a transformation routine? Psychological Research, 61, 99-106.

Hinton, G.E. & Parsons, L.M. (1981). Frames of reference and mental imagery. In A. Baddeley & J. Long (Eds.), Attention and Performance, Vol 9. Hillsdale, NJ: Erlbaum.

Humphrey, G.K. & Jolicoeur, P. (1993). Visual object identification: Some effects of image foreshortening, monocular depth cues and visual field on object identification. Quarterly Journal of Experimental Psychology, A46, 137-159.

Humphreys, G.W. & Riddoch, M.J. (1984). Routes to object constancy: Implications from neurological impairments of object constancy. Quarterly Journal of Experimental Psychology, A37, 493-495.

Jolicoeur, P. (1985). The time to name disoriented natural objects. Memory and Cognition, 13, 4, 289-303.

Jolicoeur, P. (1990). On the role of mental rotation and feature extraction in the identification of disoriented objects: A dual systems theory. Mind and Language, 5, 387-410.

Jolicoeur, P. & Humphrey, G.K. (1998). Perception of rotated two-dimensional and three-dimensional objects and visual shapes. In V. Walsh & J. Kulikowski (Eds.), Perceptual constancy. Cambridge University Press. Cambridge.

Jolicoeur, P. & Milliken, B. (1989). Identification of disoriented objects: Effects of context of prior presentation. Journal of Experimental Psychology: Learning, Memory and Cognition, 15, 200-210.

Just, M.A. & Carpenter, P.A. (1976). Eye fixations and cognitive processes. Cognitive Psychology, 8, 441-480.

Landy, M.S. (1998). Observer biases in the 3D interpretation of line drawings. Vision Research, 38, 2817-2832.

Lawson, R. (1999). Achieving visual object constancy across plane rotation and depth rotation. Acta Psychologica, 102, 221-245.

Lawson, R. & Humphreys, G.W. (1998). The neuropsychology of object constancy. In V. Walsh & J. Kulikowski (Eds.), Perceptual constancy. Cambridge University Press. Cambridge.

Lawson, R. & Jolicoeur, P. (1999). The effects of plane rotation effects on the recognition of brief masked pictures of familiar objects. Memory and Cognition, 26, 791-803.

Leek, E. C. (1998a). Effects of stimulus orientation on the identification of common polyoriented objects. Psychonomic Bulletin & Review, 5, 650-658.

- Leek, E. C., (1998b). The analysis of orientation-dependent time costs in visual recognition. Perception, 27, 803-816.
- Lowe, D. G. (1987) The Viewpoint Consistency Constraint. International Journal of Computer Vision, 1, 57-72.
- Maki, R.H. (1986). Naming and location the tops of rotated pictures. Canadian Journal of Psychology, 40, 368-387.
- Marr, D. & Nishihara, H. K. (1978). Representation and recognition of the spatial organisation of three-dimensional shapes. Proceedings of the Royal Society of London, B207, 187-217.
- McKone, E. & Grenfell, T. (1999). Orientation invariance in naming rotated objects: Individual differences and repetition priming. Perception and Psychophysics, 61, 1590-1603.
- McMullen, P.A. & Farah, M.J. (1991). Viewer-centred and object-centred representations in the recognition of naturalistic line drawings. Psychological Science, 2, 275-277.
- McMullen, P.A. & Jolicoeur, P. (1992). Reference frame and effects of orientation on finding the tops of rotated objects. Journal of Experimental Psychology: Human Perception and Performance, 18, 807-820.
- Moore, C. & Cavanagh, P. (1998). Recovery of 3D volume from 2-tone images of novel objects. In Tarr, M.J. and Bülthoff, H.H. (Eds.), Object recognition in man, monkey and machine. Cambridge, MA. MIT Press.
- Murray, J. E. (1995). The role of attention in the shift from orientation-dependent to orientation-invariant identification of disoriented objects. Memory & Cognition, 23, 49-58.
- Murray, J.E. (1999). Orientation-specific effects in picture matching and naming. Memory and Cognition, 27, 878-889.
- Murray, J. E., Jolicoeur, P., McMullen, P. A., & Ingleton, M. (1993). Orientation-invariant transfer of training in the identification of rotated natural objects. Memory & Cognition, 21, 604-610.
- Palmer, S.E. (1975). The effects of contextual scenes on the identification of objects. Memory and Cognition, 3, 519-526.
- 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 rotation. In M.J.Tarr & H.H.Bulthoff (Eds.). Object recognition in man, monkey and machine. MIT Elsevier. London.
- Pinker, S. (1984) Visual cognition. MIT. Cambridge. MA.
- Rock, I, (1973). Orientation and form. New York: Academic Press.
- Schyns, P. G. (1998). Diagnostic recognition: Task constraints, object information, and their interactions. Cognition, 67, 147-179.
- Shepard, R.N. & Metzler, D. (1988). Mental rotation: Effects of dimensionality of objects and type of task. Journal of Experimental Psychology: Human Perception and Performance, 14, 3-11.
- Shepard, R.N. & Metzler, D. (1971). Mental rotation of three-dimensional objects. Science, 171, 701-703.
- Shepard, R. N., & Cooper, L. A. (1982). Mental images and their transformations. Cambridge, MA: MIT Press.

Shepard, R.N. & Metzler, D. (1988). Mental rotation: Effects of dimensionality of objects and type of task. Journal of Experimental Psychology: Human Perception and Performance, 14 (1), 3-11.

Takano, Y. (1989). Perception of rotated forms: A theory of information types. Cognitive Psychology, 21, 1-59.

Tarr, M.J., Williams, P., Haywood, W.G. & Gauthier, I. (1998). Three-dimensional object recognition is viewpoint-dependent. Nature Neuroscience, 1 (4), 275-277.

Tarr, M. J. (1995). Rotating objects to recognize them: A case study on the role of viewpoint dependency in the recognition of three-dimensional objects. Psychonomic Bulletin & Review, 2 (1), 55-82.

Tarr, M. J. & Bülthoff, H. H. (1998). Image-based object recognition in man, monkey and machine. In M. J. Tarr and H. H. Bülthoff (Eds.) Object recognition in man, monkey and machine. London. MIT Press

Tarr, M. J. & Pinker, S. (1990). When does human object recognition use a viewer-centred reference frame? Psychological Science, 1 (4), 253-256.

Tarr, M. J. & Pinker, S. (1989). Mental rotation and orientation dependence in shape processing. Cognitive Psychology, 21, 233-282.

Thacker, N.A., Riocreux, P.A. & Yates, R.B. (1994). Assessing the completeness properties of pairwise geometric histograms. Image and Vision Computing, 13, 423-429.

Turnbull, O.H., Carey, D.P. & McCarthy, R.A. (1997). The neuropsychology of object constancy. Journal of the International Neuropsychological Society, 3, 288-298.

Turnbull, O.H. & McCarthy, R.A. (1996). Failure to discriminate between mirror-image objects: A case of viewpoint-independent object recognition. Neurocase, 2, 63-72.

Ullman, S. (1989). Aligning pictorial descriptions: An approach to object recognition. Cognition, 32, 193-254.

Wechsler, H. (1990). Computational Vision. Academic Press. London.

Wiser, M. (1981). The role of intrinsic axes in shape recognition. Proceedings of the 3rd Annual Conference on Cognitive Science, 184-186.

FIGURE CAPTIONS

Figure 1. Stimulus sets A (Experiments 1 and 2), B (Experiment 1) and C (Experiment 2) used in the studies (from Tarr & Pinker, 1990).

Figure 2. Mean RTs for (a) Set A and (b) Set B stimuli as a function of stimulus orientation (30-120 degrees) and block. Bars show standard error of the mean.

Figure 3. Linear regression slopes as a function of block for (a) Set A and (b) Set B stimuli. Dotted line indicates 1ms/deg cut-off for spatial normalisation. Bars show standard error of the mean.

Figure 4. Mean RTs for (a) Set A (b) Set C stimuli as a function of stimulus orientation (30-120 degrees) and block. Bars show standard error of the mean.

Figure 5. Linear regression slopes as a function of block for (a) Set A and (b) Set C stimuli. Dotted line indicates 1ms/deg cut-off for spatial normalisation. Bars show standard error of the mean.

FIGURE 1




FIGURE 2(a)




FIGURE 2(b)




FIGURE 3(a)




FIGURE 3(b)




FIGURE 4(a)




FIGURE 4(b)




FIGURE 5(a)




FIGURE 5(b)

