ELSEVIER

Available online at www.sciencedirect.com

SCIENCE DIRECT®

Cognition 95 (2005) 73–93

COGNITION

www.elsevier.com/locate/COGNIT

Orientation-invariant object recognition: evidence from repetition blindness

Irina M. Harris*, Paul E. Dux

Macquarie Centre for Cognitive Science, Macquarie University, Sydney, NSW 2109, Australia

Received 5 September 2003; revised 10 December 2003; accepted 19 February 2004

Abstract

The question of whether object recognition is orientation-invariant or orientation-dependent was investigated using a repetition blindness (RB) paradigm. In RB, the second occurrence of a repeated stimulus is less likely to be reported, compared to the occurrence of a different stimulus, if it occurs within a short time of the first presentation. This failure is usually interpreted as a difficulty in assigning two separate episodic tokens to the same visual type. Thus, RB can provide useful information about which representations are treated as the same by the visual system. Two experiments tested whether RB occurs for repeated objects that were either in identical orientations, or differed by 30, 60, 90, or 180°. Significant RB was found for all orientation differences, consistent with the existence of orientation-invariant object representations. However, under some circumstances, RB was reduced or even eliminated when the repeated object was rotated by 180°, suggesting easier individuation of the repeated objects in this case. A third experiment confirmed that the upside-down orientation is processed more easily than other rotated orientations. The results indicate that, although object identity can be determined independently of orientation, orientation plays an important role in establishing distinct episodic representations of a repeated object, thus enabling one to report them as separate events.

Keywords: Object recognition; Object constancy; Implicit recognition; Repetition blindness

The human visual system is extremely efficient at recognising objects under a wide variety of conditions, including changes in illumination, size and viewpoint. This remarkable feature is termed *object constancy*. A number of theories of object recognition have been proposed to account for object constancy, and in particular for object recognition from

---

* Corresponding author. Tel.: +61-29850-6710; fax: +61-29850-6059.
E-mail address: iharris@maccs.mq.edu.au (I.M. Harris).

doi:10.1016/j.cognition.2004.02.006

different viewpoints. These theories can be grouped into two broad families: those that posit that object recognition is viewpoint-dependent and those proposing that object recognition is viewpoint-invariant.

Viewpoint-dependent recognition is postulated on the basis that there is often a monotonic decrease in performance (as measured by reaction time and/or accuracy) with increasing misorientation from a preferred view of an object (Jolicoeur, 1985; Tarr & Pinker, 1989). Viewpoint-dependent theories state that objects are represented in a viewer-centred frame of reference determined by the location of the object relative to the observer. With respect to the nature of the object representation, some theorists believe that representations of objects are stored in a single canonical orientation (e.g. Palmer, Rosch, & Chase, 1981), while others argue that multiple views are stored, corresponding to different instances encountered during the course of one's experience with the object (Bülthoff & Edelman, 1992; Tarr & Pinker, 1989). Each model has implications for the kinds of transformations involved: the former requires that novel views of objects be recognised by transforming the input image to the stored canonical representation (Jolicoeur, 1985; Ullman, 1989), the latter by transformation to the nearest stored view (Tarr, 1995; Tarr & Pinker, 1989). A further possibility that has been proposed is that recognition of novel views of objects might occur through interpolation between previously stored views (e.g. Bülthoff & Edelman, 1992; Edelman, 1999).

In contrast, viewpoint-invariant recognition is inferred if there is little cost to recognition when changing the orientation or viewpoint (Biederman & Gerhardstein, 1993; Corballis, Zbrodoff, Shetzer, & Butler, 1978). Viewpoint-invariant theories of object recognition postulate that objects are represented on the basis of distinctive features and their inter-relations, which remain constant across changes in viewpoint (Biederman, 1987; Corballis, 1988; Marr, 1980). Perhaps the most influential theory in this category is that of David Marr (Marr, 1980; Marr & Nishihara, 1978). According to this theory, one step towards object recognition is to construct a visual representation that provides information about edges and surfaces as defined from the viewer's perspective. This '$2\frac{1}{2}$-D sketch' is considered the richest purely bottom-up visual representation, but it is not sufficient to enable object recognition. Recognition can only be achieved once a 3D object representation is constructed, in which the object features are defined with respect to a reference frame centred on the object rather than on the viewer. This object-centred representation allows recognition from various views, because its structure is not affected by rotations and changes in viewing conditions. Several other authors have proposed that recognition could proceed largely independently of any reference frame, for example through the identification of salient features (Corballis, 1988; Deutsch, 1955; Humphreys & Riddoch, 1984; Warrington & James, 1986). Corballis (1988) has argued that, while this is a relatively crude recognition process, it is sufficient to activate a stored representation that contains more elaborate information about the object, including information about the object's internal reference frame (i.e. the location of its principal axes). This reference frame could then be applied to the visual input to refine the initial recognition process. Note that the principal difference between Corballis' and Marr's theories is that Marr's theory states that a viewpoint-dependent representation is a precursor to the formation of an object-centred representation, whereas Corballis argues that recognition is achieved before a reference frame is assigned to the object, with the obvious implication that the activation of an object representation in memory may bypass a viewer-centred representation.

Recently, there have been a number of reports of neurological patients who demonstrate intact, or near intact, recognition of objects presented in different orientations, while at the same time being unable to determine the orientation of the object (Harris, Harris, & Caine, 2001; Karnath, Ferber, & Bülthoff, 2000; Turnbull, Beschin, & Della Sala, 1997; Turnbull, Laws, & McCarthy, 1995). For example, such patients have no difficulty in recognising an upside-down dog, but are as likely as not to say that it is in its canonical (upright) orientation. It has been argued that these patients provide evidence for the existence of an orientation-independent route to object recognition (Turnbull et al., 1997). As such, they appear to provide support for viewpoint-invariant theories of object recognition, in particular that proposed by Corballis (1988).

However, the finding that recognition of rotated objects can be achieved in the absence of knowledge of object orientation need not necessarily imply that object recognition itself is orientation-invariant. For example, it is still possible that recognition is based on viewpoint-dependent representations which are not available to conscious awareness, while a separate mechanism is responsible for explicitly coding the orientation of the object, and it is the latter mechanism that is impaired in patients with agnosia for object orientation (Harris et al., 2001). This idea receives some support from the fact that a number of patients with agnosia for object orientation demonstrate some residual sensitivity to orientation, particularly the upright canonical orientation of objects (and sometimes also the upside-down orientation), even though they cannot use this information in a flexible and explicit way (Harris et al., 2001; Karnath et al., 2000). Such findings are somewhat difficult to reconcile with a completely orientation-independent object representation.

In the present study, we address the question of whether object representations unavailable to conscious awareness are orientation-dependent or orientation-invariant, using the paradigm of *repetition blindness* (RB) (Kanwisher, 1987). RB refers to a failure to detect the second occurrence of a repeated item under conditions of rapid serial visual presentation (RSVP), when it is presented within approximately 400 ms of the first member of the repeated pair. This failure to report a repeated item is not due to a limitation in short term memory capacity or forward masking, because a *different* item presented in the same temporal position is usually detected and remembered successfully. Thus, RB represents an example of implicit recognition of the repeated item, because some form of recognition must take place to give rise to the repetition effect, yet the repeated item is not consciously perceived. It is tempting to infer that RB reflects some kind of perceptual or neural refractoriness, such that, once a representation of an item has been activated, its recognition threshold is elevated for a short period of time, making it more difficult to recognise on a subsequent occasion. However, there is evidence that this is not actually the case. Previously it has been found that subjects are in fact *better* at reporting an item if they have seen it earlier in the list, provided they did not have to respond to that first occurrence (Kanwisher, 1987, Exp. 3). Given this finding, Kanwisher (1987) proposed that RB results from a failure to assign two separate tokens to the same visual type. *Types*, which are long-term representations stored in memory, are distinguished from *tokens*, which are episodic representations of particular visual types activated during a perceptual task (see also

Kahneman & Treisman, 1984; Kahneman, Treisman, & Gibbs, 1992). Thus, Kanwisher's 'token individuation hypothesis' states that RB arises because, once a type has been token individuated, it becomes unavailable for subsequent token individuation for a short period of time. This precludes the encoding of the repeated items as separate events.

In a subsequent elaboration of the type–token model, Chun (1997) proposed a two-stage process of tokenisation. In the first stage, identity information (visual type) and its position in the RSVP stream (spatio-temporal token) are extracted in parallel by two different processing modules. The spatio-temporal tokens confer episodic distinctiveness of target relevant events. However, before this information becomes accessible for overt report, it must undergo consolidation, with visual types and spatio-temporal tokens being bound together to create an *object token*, similar to the notion of object file proposed by Kahneman and Treisman (1984). Object tokenisation is an attention-demanding and capacity-limited process, creating a potential bottleneck that introduces a delay for subsequent items, during which the initial representation (either type or spatio-temporal token) of that item is lost. Chun's model explains RB as a failure to consolidate multiple spatio-temporal tokens into separate object tokens.

RB can be useful in addressing the question of which representations are treated as the same by the visual system, in the absence of spatio-temporal information provided by episodic tokens. For example, if object representations stored in memory are orientation-invariant, we would expect to see equivalent RB for identical (repeated) objects presented in different orientations, because they would converge on the same object representation. On the other hand, if object representations are orientation-dependent, then one would not necessarily predict the same amount of RB for differently oriented instances of the same object. One possible scenario is that objects are represented in memory as a collection of specific views (Tarr & Bülthoff, 1995; Tarr & Pinker, 1989), in which case one would expect no RB between repeated objects presented in different orientations, because they would activate distinct memory representations. Alternatively, if a single viewpoint-dependent representation (e.g. the canonical view) is stored in memory and a normalisation process is required to match a rotated exemplar to this stored representation, then one might expect the RB effect to vary as a function of the degree of rotation from the canonical view.

Kanwisher, Yin, and Wojciulik (1999) have reported a series of experiments using RB to explore various aspects of object recognition, amongst them the effect of orientation and viewpoint changes on the accuracy of perceiving a repeated object. With respect to orientation, they found that a 30° rotation in the picture-plane resulted in as much RB as presenting the two objects in the same orientation, and concluded that the object representations giving rise to RB are orientation-invariant. This conclusion was strengthened by the finding that RB also occurred for objects photographed from different viewpoints (i.e. depth rotations). However, one limitation of the Kanwisher et al. experiment is that it only tested for relatively small differences in picture-plane orientation. Some studies have shown that rotations of 30° in the picture-plane can have negligible effects on recognition, whereas larger rotations lead to significant decrease in recognition accuracy, suggesting an orientation-dependent process for these larger rotations (Lawson & Jolicoeur, 1998, 2003). Similarly, neurophysiological recordings in monkeys' superior temporal sulcus revealed that although neuronal responses to faces are primarily viewpoint-specific, these neurons also respond reasonably strongly to orientations close to their preferred orientation (Ashbridge, Perrett, Oram, & Jellema, 2000). Given this evidence, it seems a little premature to conclude that the object representations giving rise to RB are completely orientation-invariant.

Thus, the aim of the present study was to determine whether object representations generated outside conscious awareness are orientation-dependent or orientation-invariant. To this end, we used a RB paradigm based on that used by Kanwisher et al. (1999), but expanded to include a larger range of orientations. If significant RB is obtained across all these orientations, then we can make a strong case that these object representations are indeed orientation-invariant. Conversely, a failure to find RB for some orientations, or a systematic modulation of the RB effect as a function of orientation, would indicate that these representations are at least partially sensitive to orientation. If such orientation-dependence were observed, a further aim was to determine which orientations lead to less RB or, in other words, appear to be better discriminated by the visual system.

1. Experiment 1

Experiment 1 investigated whether RB occurs for two repeated objects that were either in identical orientations (both upright) or differed by 30, 60, 90 or 180°. The task required participants to report three pictures presented in rapid succession, for 100 ms each. The critical items were the first picture (C1, or critical item 1) and the third picture (C2, or critical item 2); these were either the same object (repeated condition) or different objects (non-repeated condition). The intervening item between C1 and C2 served as a distractor. Lower performance on repeated trials compared to non-repeated trials is indicative of RB.

1.1. Method

1.1.1. Subjects

Twenty-four first year psychology students (7 males), aged 19–46 years (mean = 25), participated in the experiment in exchange for course credit.

1.1.2. Apparatus

Stimuli were presented on a Dell Flat Trinitron monitor with 120 Hz vertical refresh rate controlled by a Dell PC computer. The experiment was constructed and run using DMDX software (Forster & Forster, 2003).

1.1.3. Stimuli and design

The stimuli were 64 pictures from the Snodgrass and Vanderwart (1980) corpus, chosen to have a well-defined canonical upright orientation. They were presented as black line drawings on a white background and subtended a visual angle of approximately 12° at the viewing distance of 45 cm. Three masks were created using random geometrical shapes of similar line thickness and density as the pictures.

Three independent variables were manipulated in a 2 × 2 × 5 within-subject design (see Fig. 1). The first variable was the relationship between the two critical items: repeated

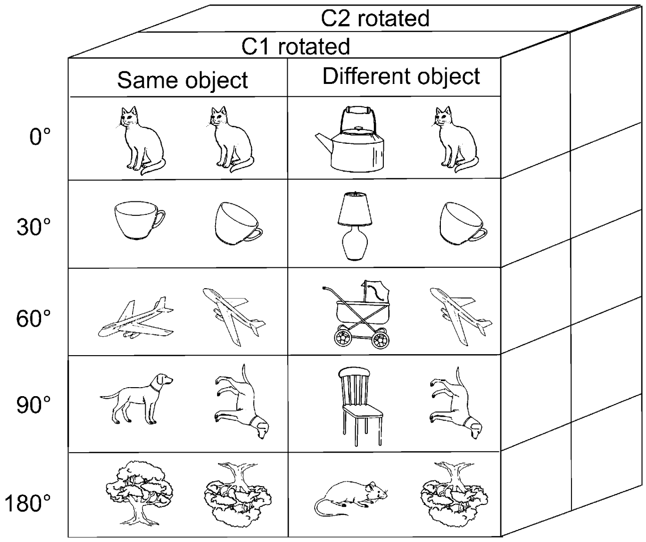

Fig. 1. Graphical representation of the experimental design. Three factors were manipulated: (1) the relationship between the critical items (same or different object); (2) the orientation of the rotated item; and (3) whether the first (C1) or second (C2) item of the pair was rotated. The critical items depicted here were separated by an intervening object (not shown).

(same object) or non-repeated (different object). The second variable was whether C1 or C2 was rotated (the alternative item always being upright). The third factor was the orientation of the rotated item (0, 30, 60, 90, 180°). Thirty-two of the pictures served as critical items. Sixteen of these always appeared in the C2 position. On repeated trials, these were paired with the identical picture appearing in the C1 position, while on non-repeated trials they were paired with the other 16 pictures at C1 (see Fig. 1). The two groups of 16 pictures had similar complexity and familiarity ratings (mean visual complexity: 3.04 and 3.31, respectively, $t(15) = 0.30$; mean familiarity: 3.72 and 3.56, respectively, $t(15) = 0.17$, Snodgrass & Vanderwart, 1980). The remaining 32 pictures served as intervening stimuli between the two critical items. In this experiment, all intervening items between C1 and C2 were displayed in the upright orientation. Thus, there were 16 trials in each condition, making a total of 320 experimental trials ($2 \times 2 \times 5 \times 16$).

1.1.4. Procedure

The experiment was conducted under normal illumination conditions and lasted about 1 h. The subjects were seated approximately 45 cm in front of the computer monitor and gave their responses verbally; these were coded by the experimenter. Before the experiment, subjects completed a familiarisation phase in which they saw all the pictures (in their upright orientation) and named them. Any naming errors were corrected at this stage.

The experiment was self-paced and the participants initiated each trial by pressing the space bar. Each trial consisted of three pattern masks, followed by the three pictures

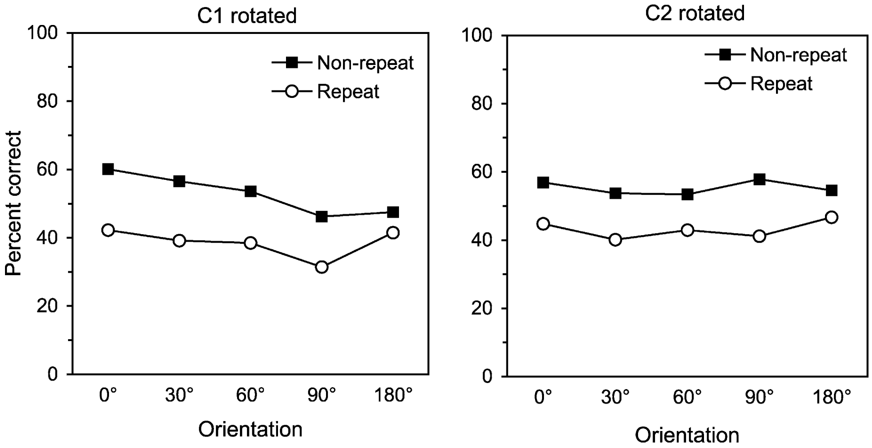

Fig. 2. Mean percent correct recall of both C1 and C2 in Experiment 1, as a function of orientation, depicted separately for trials in which C1 or C2 was rotated.

(C1, intervening item, C2), followed by the same three pattern masks in reverse order. Each of these frames was presented for 100 ms. At the end of the trial, a message appeared on the screen prompting the subject to recall all the pictures seen. They were told that sometimes a picture was repeated and, if that was the case, they should name it twice. In addition to the 320 experimental trials, 32 filler trials containing only two pictures (and an extra mask) were included, in order to discourage subjects from guessing the presence of an undetected repeated picture. The order of the trials was random. Before the start of the experiment, subjects received 32 practice trials at increasingly fast presentation rates ranging from 500 ms per picture (8 trials), to 150 ms per picture (8 trials), to 100 ms per picture (16 trials).

1.2. Results

Trials on which neither one of the critical items (C1 or C2) were recalled were discarded. This means that, for repeated objects, only trials for which we could be sure that the object representation of interest had been activated (i.e. the object was named at least once) were counted.[1] For symmetry, the same criterion was applied to non-repeated items. The data were then scored in terms of percent of trials in which both C1 and C2 were correctly recalled and are presented in Fig. 2, plotted separately for trials in which C1 or C2 was the rotated item.

A repeated measures ANOVA revealed a significant main effect of Repetition (same-object trials vs. different-object trials), $F(1,23) = 8.76$, $P = 0.007$, with lower overall accuracy for repeated pictures than for non-repeated pictures. There was also a significant difference between trials in which C1 was rotated compared to those in which C2 was rotated, $F(1,23) = 5.78$, $P = 0.025$, and a significant main effect of Orientation

---

[1] We opted for this method, rather than the more stringent method of counting only trials on which C1 was reported, because it is often impossible to distinguish which one of the repeated items is reported in these tasks.

Table 1
Size of RB effect in Experiments 1 and 2, as a function of whether C1 or C2 was rotated and the angle of rotation

	Experiment 1 C1 rotated	C2 rotated	Experiment 2 C1 rotated	C2 rotated
0°	18	12	12	13
30°	17	14	14	9
60°	15	10	15	11
90°	15	17	9	18
180°	$6^{\#}$	8	3 n.s.	23

The size of the RB effect is expressed as the percent difference between non-repeated and repeated items. Significant RB ($P < 0.01$) was present in all cases, except as indicated. $^{\#}P = 0.04$, uncorrected for multiple comparisons.

$F(4,92) = 4.73$, $P = 0.002$. Repetition did not interact with type of trial, $F(1,23) = 0.91$, suggesting similar patterns of RB regardless of whether C1 or C2 was rotated. However, there was a significant Repetition × Orientation interaction, $F(4,92) = 3.28$, $P = 0.015$, which was due to a reduction in the RB effect when pictures were rotated by 180° (see Table 1). The 3-way interaction was not significant, $F(4,92) = 0.96$, $P = 0.43$, illustrating the fact that the reduction in RB for pictures rotated by 180° was present both when C1 and C2 were rotated.

Separate ANOVAs were performed for C1-rotated and C2-rotated trials. When C1 was rotated, there was a significant main effect of Repetition, $F(1,23) = 10.08$, $P = 0.004$, a main effect of Orientation, $F(4,92) = 9.84$, $P < 0.001$ and a significant Repetition × Orientation interaction $F(4,92) = 2.77$, $P = 0.032$. Simple comparisons revealed robust differences between repeated and non-repeated items for 0, 30, 60 and 90° orientation differences ($F > 25$, $P < 0.001$), with the size of RB remaining fairly constant across these conditions (see Table 1). There was also a small difference between repeated and non-repeated items in the 180° rotation condition, but this difference was less reliable and did not survive correction for multiple comparisons ($F(1,92) = 4.34$, $P = 0.04$, uncorrected for multiple comparisons).

When C2 was rotated, there was a significant main effect of Repetition, $F(1,23) = 6.68$, $P = 0.017$, but no effect of Orientation, $F(4,92) = 1.20$, $P = 0.31$ and no interaction between Repetition and Orientation, $F(4,92) = 1.33$, $P = 0.27$, suggesting that RB did not change significantly as a function of C2 orientation.

1.3. Discussion

The main finding of Experiment 1 is that RB occurred in all orientation conditions. Sizeable RB was obtained for orientation differences up to 90°, irrespective of whether C1 or C2 was rotated. Significant RB was also obtained for orientation differences of 180°, although this was much reduced in magnitude and was only reliably present when C2 was rotated. These results support Kanwisher et al.'s (1999) conclusion that the object representations involved in the recognition of objects in RSVP streams are orientation-invariant, and extend those findings to a much wider range of orientations.

It could be argued that the finding of RB for rotated objects does not necessarily imply the existence of orientation-invariant object representations, but could equally well be explained by orientation-dependent representations and processes that compensate for orientation differences. In other words, it is possible that RB arises at a later stage of processing, after viewpoint-dependent representations have been normalised and object constancy has been achieved. However, there are at least two reasons to doubt this interpretation. First, if RB reflected object constancy following some normalisation process that compensates for orientation differences, one might expect the size of the RB effect to be systematically affected by the magnitude of the orientation difference between the items. Our results do not conform to this pattern, as the size of RB was fairly constant across orientations, with the sole exception of 180°. Second, an inspection of the accuracy with which subjects reported the item presented between C1 and C2 revealed a difference between repeated and different-object trials. The middle item was reported considerably more often when it occurred between repeated objects (85% correct) than when it occurred between two different objects (72% correct), a difference that was highly significant ($F(1,23) = 223.01$, $P < 0.0001$). This finding suggests that the repeated objects required less processing and, consequently, more attentional resources were available for processing the middle item. Such a finding cannot be easily accommodated by the notion that RB is due to processes that compensate for orientation differences between viewpoint-dependent object representations, because presumably such processes, and their respective attentional demands, would be the same in repeated and different-object trials. On the other hand, this result is precisely what one would expect to see if fewer types were activated during the repetition trials compared to the different-object trials (i.e. 2 instead of 3). On balance, then, the pattern of results seems to be more easily accommodated by the idea that the repeated objects activate a unique orientation-invariant representation.[2] We return to this issue in Section 4.1.

The conclusion that the object representations that mediate recognition and RB are orientation-invariant is somewhat at odds with the finding that overall accuracy, for both repeated and non-repeated objects, decreased as a function of orientation when C1 was the rotated item. This finding is reminiscent of many reports in the literature that reaction time increases in a near-linear fashion as an object is rotated further away from the canonical orientation (Jolicoeur, 1985; Murray, 1995; Tarr & Pinker, 1989), a result that has been consistently interpreted as evidence for viewpoint-dependent object recognition. Again, though, some features of the results speak against an interpretation of this effect in terms of viewpoint-dependent recognition. First of all, as outlined in the introduction, if recognition were mediated by viewpoint-dependent representations, one would expect to see either no RB, or at least a fairly systematic modulation of the RB effect across changes in orientation. This was clearly not the case, as RB was present and of a similar magnitude across orientation changes, with the sole exception of 180° rotations. Second, the decrease

---

[2] A reviewer questioned whether this effect might be due to the fact that it may be easier to remember the middle item when it occurred between repeated items because that situation only requires the subject to remember two identities, rather than three. An explanation in terms of memory failure seems unlikely in the present study, given that the subjects only ever had to remember a maximum of three items, which is well within the limits of short-term memory. In addition, numerous studies in the RB literature have shown that RB under these conditions occurs at a perceptual, rather than memory, level (Chun, 1997; Chun & Cavanagh, 1997).

in overall accuracy across orientations was only seen when C1 was rotated but not when C2 was rotated. It is difficult to explain this discrepancy in terms of processes involved in the recognition of individual items, as one would expect these to be the same in both situations (the individual items were identical in the two types of trial, the only difference being the relative order of the upright vs. rotated critical items). Therefore, it seems more likely that the discrepancy in the patterns of performance is due to differences at the level of trial structure, rather than individual item recognition. This raises the interesting possibility that the viewpoint-dependent effects seen when C1 is rotated might be related to a processing bottleneck that occurs at a later, post-recognition, stage. We will return to this issue in Section 4.3.

An intriguing finding of this experiment is that the size of RB was reduced for rotations of 180°, in particular when C1 was rotated. One way to interpret this reduction in RB is to assume that upright and upside-down objects are sufficiently different from each other to activate different representations. This would imply some degree of viewpoint-dependency, although it is difficult to see why this should only be the case for 180° rotations and not for other large orientation differences such as 90°, or even 60°. An alternative explanation for the reduction in RB in this case is that the two instances of a repeated picture activate the same (orientation-invariant) representation, but it is easier to individuate the two when one is upright and the other is upside-down. This could happen if it were easier to determine the upside-down orientation compared to other orientations, a pattern that is in fact exhibited by some patients with object orientation agnosia (Harris et al., 2001; Karnath et al., 2000). Thus, having established relatively quickly that an object was upside-down, this would enable subjects to encode the two repeated objects as different instances, one upright and one upside-down. In contrast, if it were more difficult to determine the other orientations, then under RSVP conditions subjects may not be able to extract sufficient orientation information from the rotated item to be confident that it is indeed different from another occurrence of that same object on that trial. Such a situation would lead to larger RB.

If this latter explanation is correct, we would expect that increasing the orientation processing demands during the task (for example, by showing most of the objects in non-upright orientations) should result in more RB in the case of objects rotated by 180°, because it would interfere with subjects' ability to individuate repeated objects on the basis of orientation. This prediction was tested in Experiment 2.

2. Experiment 2

This experiment investigated whether increasing the orientation processing demands of the task would make it more difficult for subjects to use orientation as a way of differentiating two repeated objects, resulting in robust RB even for objects rotated by 180°. The orientation processing demands of the task were increased by presenting the intervening item (that presented between C1 and C2) rotated away from the upright; the experiment was otherwise identical to Experiment 1. This means that, on most trials of this experiment, two of the three pictures were presented in non-upright orientations (the exception being the trials on which both critical items were upright). Moreover, the pictures always differed from each other by at least 60°, with the expectation that this would place considerable strain on subjects' ability to extract orientation information under the RSVP conditions of the experiments.

2.1. Method

2.1.1. Subjects

Twenty-four undergraduate students (7 males) aged 17–41 years (mean = 22) took part in this experiment and were paid $10 for participation. None had participated in the previous experiment.

2.1.2. Stimuli and design

The design of the experiment was identical in every respect to Experiment 1, except that here the intervening item was rotated away from the upright. This middle item could be in one of three orientations: 60, 90 or 180°, chosen such that the intervening item always differed from the critical items by at least 60°.

2.2. Results

As in Experiment 1, the data were scored in terms of percent of trials in which both C1 and C2 were correctly recalled, excluding any trials on which both critical items went unreported. The results are presented in Fig. 3, plotted separately for trials in which C1 or C2 was rotated.

A repeated-measures ANOVA revealed a significant main effect of Repetition, $F(1,23) = 8.68$, $P = 0.007$, a significant difference between trials in which C1 or C2 was rotated, $F(1,23) = 11.53$, $P = 0.003$, and a significant effect of Orientation, $F(1,23) = 10.33$, $P < 0.001$. In this experiment, there was a significant Repetition × type of trial interaction, $F(4,92) = 6.50$, $P = 0.018$, which indicates different patterns of RB when C1 was rotated compared to when C2 was rotated. There was also a significant

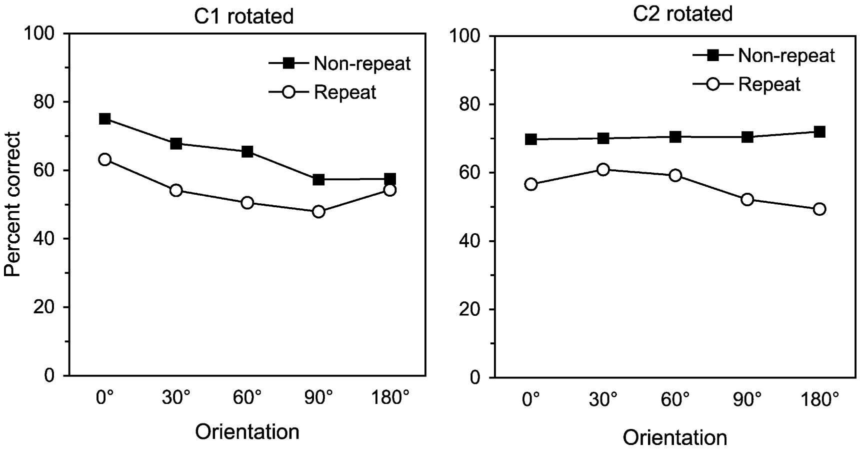

Fig. 3. Mean percent correct recall of both C1 and C2 in Experiment 2, as a function of orientation, depicted separately for trials in which C1 or C2 was rotated.

3-way interaction, $F(4,92)=5.93$, $P<0.001$, suggesting that RB was modulated differently by item orientation on trials when C1 was rotated compared to trials when C2 was rotated. This effect is clearly illustrated in Fig. 3, which shows a loss of RB when C1 was rotated by 180°, but an increase in RB when C2 was rotated by 180° (see also Table 1).

Separate ANOVAs were carried out for C1-rotated and C2-rotated trials. When C1 was rotated, there was a main effect of Repetition, $F(1,23)=6.26$, $P=0.02$, a significant effect of Orientation, $F(1,23)=17.89$, $P<0.001$, and a marginal Repetition × Orientation interaction, $F(4,92)=2.26$, $P=0.069$. Simple comparisons revealed significant differences between repeated and non-repeated items for orientation differences of 0, 30, 60 and 90° ($F>9.18$, $P<0.003$), but not 180°, $F(1,92)=1.06$, $P=0.30$.

When C2 was rotated, there was a significant main effect of Repetition, $F(1,23)=10.62$, $P=0.004$, but no overall effect of Orientation, $F(1,23)=1.64$, $P=0.174$. There was, however, a significant Repetition × Orientation interaction, $F(4,92)=4.33$, $P=0.003$. Simple comparisons revealed significant RB at all orientations ($F>12.85$, $P<0.001$), although the size of the RB effect did vary across orientation conditions, which accounts for the significant interaction. In particular, there was a dramatic increase in RB when C2 was rotated by 180° as well as a tendency for RB to increase when C2 was rotated by 90°, compared to smaller angles (see Table 1).

2.3. Discussion

This experiment replicated all the main findings of Experiment 1. As in Experiment 1, we found consistent RB for orientations up to 90°, regardless of whether C1 or C2 was rotated, but in addition, in the present experiment we obtained very large RB when C2 was rotated by 180°. This result makes the reduction in RB at 180° seen in Experiment 1 difficult to accommodate within a viewpoint-dependent framework. We elaborate this point further in Section 4.2.

The findings of this experiment support the conclusion that the reduction in RB between upright and upside-down objects seen in Experiment 1 is due to easier processing, and therefore better individuation, of these particular orientations. In the present experiment, preceding the upside-down C2 by a rotated distractor presumably made it more difficult to process its orientation and consolidate it into a stable object token, resulting in increased RB compared to Experiment 1. In contrast, when C1 was upside-down we found the same pattern of results as in Experiment 1, that is, a loss of RB. Presumably this is due to the fact that the first item in the sequence receives no interference from preceding items and, therefore, its orientation is more likely to be processed and identified as different from the upright C2.

3. Experiment 3

The results of the first two experiments suggest that the upright (0°) and upside-down (180°) orientations may be processed more easily than other orientations (30, 60, 90°). In this experiment, we investigated whether the accuracy with which people explicitly judge the orientation of a briefly presented object depends on the orientation of the object.

Since there was no evidence from Experiments 1 and 2 of any particular differences amongst the orientations of 30, 60 and 90°, in this experiment we used only objects that were upright, upside-down, or rotated by 90° (clockwise and anti-clockwise). This combination of orientations minimised working memory demands and provided a convenient and straightforward set that participants could respond to using the arrow keys on the keyboard.

3.1. Method

3.1.1. Subjects

Twenty-four first year psychology students (6 males), aged 18–38 (mean = 22), participated for course credit. None had participated in either of the previous experiments.

3.1.2. Apparatus

This was the same as in Experiments 1 and 2.

3.1.3. Stimuli and design

Ninety-six line drawings with an obvious canonical upright orientation were selected from the Snodgrass and Vanderwart (1980) corpus. Six masks similar to those used in Experiments 1 and 2 were also used. The 96 pictures were divided into four sets of 24 pictures, matched for familiarity and visual complexity based on norms from Snodgrass and Vanderwart. Each set was presented in one of four orientations: 0, 90, 180, or 270° (i.e. 90° anti-clockwise). Within each set of 24 items, one-third (8 pictures) were presented for each of the following durations: 42, 67, and 100 ms. Four versions of the experiment were constructed, with each set of 24 items appearing in a different orientation in each version, meaning that across the whole experiment each object appeared in all four orientations. The order of the trials was randomised, with the restriction that no more than three consecutive trials could have the same orientation or the same exposure duration.

3.1.4. Procedure

Following 15 practice trials, each subject completed one version of the experiment (96 trials in total) in a session lasting approximately 15 min. The experiment was self-paced and the participant advanced to the next trial by pressing the space bar. Each trial began with a fixation cross displayed for 1 s. This was followed by the picture (with variable exposure duration) and then a mask presented for 250 ms. The mask was replaced by a blank screen for 1 s, followed by the word ‘Ready’. The subjects had to indicate the orientation of the picture by pressing one of the four arrow keys on the keyboard, corresponding to ‘up’ (0°), ‘down’ (180°), ‘rotated right’ (90°) and ‘rotated left’ (270°).

3.2. Results

Percent correct responses for the different orientations and time exposures are shown in Fig. 4. The data were analysed using a 3 (Exposure duration: 42, 67, 100 ms) × 4 (Orientation: 0, 90, 180, 270°) within-subjects ANOVA. There was a significant effect of

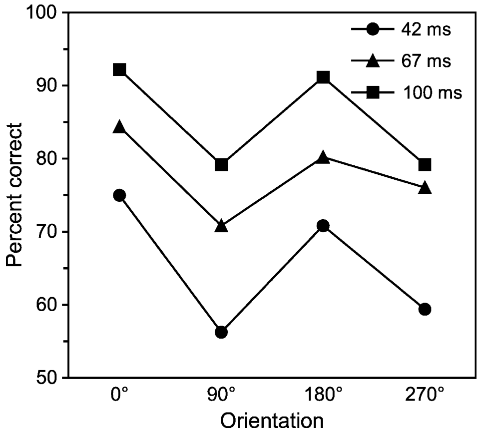

Fig. 4. Mean percent correct orientation judgements in Experiment 3, as a function of stimulus orientation and exposure duration.

Exposure duration, $F(2,23) = 47.15, P < 0.001$, with accuracy increasing systematically with exposure duration. There was also a significant effect of Orientation, $F(3,69) = 10.95$, $P < 0.001$, but no Exposure × Orientation interaction, $F(6,138) = 0.74, P = 0.62$. As can be seen in Fig. 4, the pattern of performance across orientations was remarkably consistent for all exposure durations. Across all exposures, participants were significantly more accurate at judging the orientation of upright and upside-down objects compared with objects rotated by 90°, as revealed by a contrast with weights $1 - 1\ 1 - 1$ for orientations 0, 90, 180, 270°, $F(1,69) = 31.01, P < 0.001$. Accuracy did not differ for upright compared to upside-down objects, $F(1,69) = 1.03$, $P = 0.31$, or between the two 90° orientations, $F(1,69) = 0.81, P = 0.37$.

3.3. Discussion

This experiment clearly demonstrates that it is easier to determine the orientation of an upright or upside-down object compared to objects rotated by 90°. This difference was remarkably consistent across the three exposure durations, ranging from very brief (42 ms) to reasonably lengthy (100 ms) exposure durations. Interestingly, contrary to what one might expect, it seems just as easy to extract the orientation of an upside-down object as it is to process the orientation of an upright object. In contrast, determining the orientation of an item rotated by 90° seems to require more time. We have argued previously that this is due to the fact that an additional step of locating the top-bottom axis of an object must be performed for rotations of 90°, a step which can be circumvented in the case of upright and upside-down objects (Harris et al., 2001).

It is important to note that even when the picture was displayed for 100 ms performance was not at ceiling. Extrapolating from these results, one would expect that extracting orientation information from a picture presented for 100 ms in a stream of other pictures, forward and backward masked, as in Experiments 1 and 2, should prove to be even more difficult. Thus, the present results provide an explanation as to why subjects were mostly unable to use orientation as an individuating feature in Experiments 1 and 2, resulting in RB for identical objects presented in different orientations. They also provide an explanation for why RB was sometimes reduced or lost when the critical items were upright and upside-down, because that situation allows for more rapid processing of orientation and, therefore, allows orientation to be used as an individuating feature.

4. General discussion

The present study used a RB paradigm to investigate whether the object representations generated outside conscious awareness are orientation-dependent or orientation-invariant. The principal finding was that RB occurred across changes in object orientation ranging from 30 to 180°. These results replicate and extend those reported by Kanwisher et al. (1999) and provide support for the suggestion that these object representations are orientation-invariant. Our results also revealed that orientation plays an important role in individuating two instances of a repeated object, leading to a reduction, and at times complete loss, of RB in cases where orientation is relatively easily extracted from brief visual displays. We will discuss each of these findings in turn, together with their implications for theories of object recognition.

4.1. Orientation-invariant object representations

Experiments 1 and 2 revealed sizeable RB when objects differed in orientation by up to 90°, regardless of whether C1 or C2 was rotated. Significant RB was also found for orientation differences of 180°, although this latter result was only obtained under certain conditions. Nevertheless, the fact that RB *could* be obtained for rotations of 180°, together with the lack of any systematic modulation of the RB effect as a function of orientation, is most consistent with the idea that the object representations that mediate RB, and implicitly, object recognition are orientation-invariant. Moreover, this conclusion is supported by the fact that the item intervening between C1 and C2 was recognised more accurately if it occurred between repetitions of the same object compared to when C1 and C2 were different objects. This finding implies that repetition trials require overall less processing, just as might be expected if they involved only two object representations (i.e. a single representation of the critical items and one distractor) as opposed to three (i.e. two separate critical items and one distractor).

A point of contention is whether these findings necessarily implicate orientation-invariant *visual* representations, as opposed to semantic representations that are, by definition, orientation-invariant (Arnell & Jolicoeur, 1997). In other words, it is possible that the RB obtained for our rotated objects might be based entirely on semantic representations, in which case the finding of orientation-invariance would be unsurprising. Kanwisher et al. (1999) have found some evidence of RB for different exemplars of the same basic-level object (e.g. grand piano and upright piano), as well as for semantically related objects (e.g. airplane, helicopter). Furthermore, Bavelier (1994) found RB between pictures and words depicting the same object, so an explanation in terms of semantic RB is plausible. However, we believe that there is enough evidence supporting the notion that the orientation-invariant RB seen here is based on visual representations. First, other studies have demonstrated RB for pseudo-objects that have no semantic representation (Arnell & Jolicoeur, 1997), which indicates that visual patterns of a structural complexity comparable to our objects can produce RB. Second, and more pertinent to the present study, Coltheart, Mondy, and Moore (2001) have found RB for non-objects rotated by 30° in depth and in the picture-plane, a finding which is consistent with orientation-invariant *visual* RB. Finally, in some of the conditions of the present experiments, we found a reduction in RB for objects rotated by 180°. Such a modulation of the RB effect would not be expected if RB was based purely on a semantic representation.

Thus, this study provides evidence for the rapid extraction of orientation-invariant object representations which are most likely visuo-perceptual in nature. Our findings suggest that these representations are formed early in the recognition process, they are elaborate enough to identify an object, but often fail to be perceived consciously. This is illustrated by the fact that subjects are sensitive to the repetition of an object, indicating that the repetition has been identified as such at some level, yet they fail to report the item, which indicates that it did not reach awareness. Within a type–token framework, these representations are conceptualised as *identity types* which, due to the spatio-temporal constraints of the RSVP procedure, are not successfully token-individuated and consequently are not perceived as separate events (Chun, 1997; Kanwisher, 1987).

Our results imply that orientation is not a defining feature of the object representations that mediate pre-conscious stages of object recognition. Together with the findings from patients with orientation agnosia, who are clearly able to recognise rotated objects despite not knowing their orientation, these results provide strong support for viewpoint-invariant theories of object recognition. Moreover, they indicate that viewpoint-invariance is apparent even in the early (pre-conscious) stages of recognition, a view that is not easily accommodated by Marr's (1980) theory, which states that a viewpoint-dependent representation precedes the formation of an object-centred representation. Thus, our results are more in line with theories that do not invoke any such viewpoint-dependent representations (e.g. Corballis, 1988; Deutsch, 1955).

Interestingly, our results of complete invariance with respect to orientation in the picture-plane go even beyond the predictions of some viewpoint-invariant models of object recognition, which argue that objects are represented as structural descriptions of component features (geons) and the relations between them (Biederman, 1987; Hummel & Biederman, 1992). Although these models predict viewpoint-invariance across rotations in depth (barring geon occlusion), they predict some costs associated with rotations in the picture-plane. These costs arise because the spatial relations between geons are perturbed by picture-plane rotations, such that an 'above' relation between two geons in an upright object becomes a 'beside' relation if the object is rotated by up to 180°, or a 'below' relation if the object is inverted (Hummel & Biederman, 1992). However, Biederman has suggested that such costs are possibly less likely to be due to orientation-dependent processing of individual geons, and more to do with perturbations of 'top-of' relations (Biederman, 1987, p. 140)—that is, a change in the global orientation of the object. In other words, recognition costs incurred by objects rotated in the picture-plane may be due to processes involved in determining the orientation of the object, rather than its identity, a position that has also been argued recently by a number of other authors (De Caro & Reeves, 2000; Harris & Dux, 2004, see below for further discussion).

Given this, one might surmise that our present findings of orientation-invariance in the picture-plane may in fact provide a stronger test for the existence of viewpoint-invariant object representations than invariance across depth rotations. However, it appears that these representations may not give rise to fully conscious recognition in the absence of information about the object's orientation, as we shall argue below.

4.2. Orientation as an individuating feature

A second important finding of this study, replicated in two experiments, is that RB was substantially reduced, and sometimes eliminated, when two identical objects differed by 180°. While this could be taken as evidence against our conclusion that the object representations are orientation-invariant, we argue that this reduction in RB is due to the fact that it is easier to individuate two objects when one is upright and the other is upside-down, rather than being due to the activation of different viewpoint-dependent representations. Several lines of evidence support this premise.

First, the results of Experiment 3 indicate that judging the orientation of upright and upside-down objects is significantly easier than judging other orientations (e.g. 90°). This result echoes earlier findings from a patient with agnosia for object orientation, who could discriminate between an upright and an upside-down object with much greater accuracy than between other orientations (Harris et al., 2001). Therefore, on trials in which the critical items are upright and rotated by 180° subjects would have a greater chance of processing not only the identity of the objects, but also their orientation. This additional information would help them encode the two repeated objects as different instances, based on the different orientation descriptions (upright vs. upside-down), and thus reduce susceptibility to RB. In contrast, for more difficult orientations subjects may not be able to extract sufficient orientation information from the rotated item to be confident that it is different from another presentation of that same object on that trial.

Second, the reduction in RB for objects rotated by 180° occurred in some circumstances but not others. Specifically, RB was reduced when the first object of the RSVP stream (C1) was rotated, but not when the third object (C2) was rotated. It is hard to see how viewpoint-dependent representations would explain this pattern of results, because it would imply that trials in which C1 is rotated and trials in which C2 is rotated employ different kinds of representations (viewpoint-dependent in the first case, viewpoint-invariant in the latter). On the other hand, this pattern of results can be accommodated by the idea that orientation information may be used to individuate repetitions of the same object. Specifically, we would expect that determining the orientation of the first item of the series is easier than determining the orientation of the last item, because the first item would gain privileged access to the tokenisation process, with no interference from preceding items. As such, when C1 is rotated by 180°, subjects would be more successful at consolidating its orientation and differentiating it from C2, resulting in a reduction or loss of RB. Conversely, when C2 is rotated by 180°, processing its orientation and creating an object token would be hampered to some extent by the preceding items in the stream. As a consequence, subjects would not have sufficiently reliable information to allow them to individuate the two instances of the object. Further support for this explanation comes from the fact that when C2 was rotated by 180°, the size of RB was modulated by the orientation of the preceding distractor. Smaller RB was seen in Experiment 1, where the distractor was upright, presumably because this distractor would not require much orientation processing. Much larger RB was seen in Experiment 2, where the distractor was rotated (by 60 or 90° in this specific case), because processing the orientation of the distractor would have taken considerable time and effort and this would have prevented subjects from processing the orientation of C2 in the short time available.

Thus, a second major conclusion that emerges from this study is that, although object identity can be determined independently of orientation, orientation nonetheless plays a crucial role in establishing distinct episodic representations of a repeated object, thus enabling one to report them as separate events. These results fit well with the model proposed by Chun (1997), in which identity information and spatio-temporal tokens (e.g. orientation, in this case) are extracted relatively independently from a visual stimulus, but need to be bound together into an object token before they are available for report. Chun suggested that RB results from a failure to consolidate different spatio-temporal tokens into separate object tokens, just as seems to be the case in our experiments, when the orientation of the repeated object is difficult to resolve in the time available.

4.3. Binding identity to orientation—the processing bottleneck?

Our conclusion that orientation-invariant object representations mediate recognition and RB may appear difficult to defend in the face of many findings reported in the literature, which suggest that object recognition is sensitive to viewpoint and orientation (e.g. Hayward, Tarr, & Corderoy, 1999; Jolicoeur, 1985; Murray, 1995; Tarr, 1995; Tarr & Pinker, 1989). Such studies have typically found that the time to name rotated objects increases systematically as the object is rotated further from the upright, a cost which is usually interpreted as arising from a normalisation of the visual input prior to recognition. However, other empirical results indicate that observers determine the identity of an object before they determine its orientation (Corballis et al., 1978; De Caro, 1998; De Caro & Reeves, 2000), a finding which seems to negate the need to compensate for orientation in order to determine identity. Furthermore, De Caro and Reeves (2000) have compared the effects of stimulus orientation on determining the identity vs. the orientation of a briefly presented, masked, object. They found that the time needed to verify the object's *orientation* increased systematically as a function of the object's rotation from the upright. In contrast, the time required to verify the object's *identity* was quite uniform across non-upright orientations, though longer overall than for upright objects. The authors concluded that the linear orientation effects seen in naming tasks are actually due to a process of determining the object's orientation, and that this process occurs after recognition.

This conclusion is supported by recent findings from a patient with agnosia for object orientation, who displayed a flat reaction time function when naming rotated objects (Turnbull, Della Sala, & Beschin, 2002). Note that this patient can recognise objects equally well at all orientations, but cannot interpret their orientations. Moreover, careful consideration of data from other patients with orientation agnosia also reveals that, although some of these patients demonstrate residual sensitivity to orientation (Harris et al., 2001; Karnath et al., 2000), these effects are only apparent on tasks that require orientation judgements, rather than recognition.

The findings discussed above are consistent with our present conclusion that recognition is mediated by orientation-invariant object representations. However, our present findings also indicate that establishing the orientation of an object is an important step in forming a distinct episodic representation of the object, thus enabling one to report it as a separate visual event. Thus, we agree with Chun (1997), that the type (i.e. object identity) needs to be bound to a spatio-temporal token (i.e. the object's orientation at a particular moment) before it is available for overt report (see Harris & Dux, 2004 for further elaboration of this idea). As outlined in the introduction, this binding process is attention-demanding and capacity-limited, which creates a potential bottleneck in processing (see also Treisman & Gelade, 1980). Thus, if the time taken to process either of these attributes exceeds the available time, this could result in a failure to consolidate the initial representations of these attributes into a durable object token. This scenario can explain, for example, the decrease in overall accuracy as C1 was rotated further from the upright that was found in Experiments 1 and 2. In those cases, the orientation of C1 would have been more difficult to determine, introducing a delay during which the representation of that object, and subsequent ones, would have decayed before being consolidated into a form accessible for overt report.[3] This explanation also provides a rationale for the naming costs incurred by rotated objects, even without the time constraints of RSVP, as the object's orientation would need to be determined, bound to the object's identity and consolidated, before the object is available for overt report.

To conclude, then, the present study has found evidence for the existence of orientation-invariant object representations that mediate the early stages of object recognition. At the same time, however, our findings indicate that determining the orientation of the object is an essential step in creating a conscious episodic representation of the object, which allows individuation and overt report.

Acknowledgements

This work was supported by an Early Career Researcher Grant from Macquarie University. Irina Harris was supported by an ARC Australian Post-doctoral Fellowship and Paul Dux was supported by an Australian Post-graduate Award. We thank Leila Petit and Stacey Kuan for assistance in running the experiments; Justin Harris, Mike Corballis and Steve Mondy for discussion of the research and comments on the manuscript; and two anonymous reviewers for their constructive criticisms.

---

[3] Further support for this notion comes from the observation that the accuracy with which the middle object (between C1 and C2) was reported also suffered when C1 was rotated, compared to when C2 was rotated.

References

Arnell, K. M., & Jolicoeur, P. (1997). Repetition blindness for pseudoobject pictures. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 999–1013.

Ashbridge, E., Perrett, D. I., Oram, M. W., & Jellema, T. (2000). Effect of image orientation and size on object recognition: Responses of single units in the macaque monkey temporal cortex. *Cognitive Neuropsychology*, *17*, 13–34.

Bavelier, D. (1994). Repetition blindness between visually different items: The case of pictures and words. *Cognition*, *51*, 199–236.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115–147.

Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional view-point invariance. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 1162–1182.

Bülthoff, H. H., & Edelman, S. (1992). Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proceedings of the National Academy of Science, USA*, *89*, 60–64.

Chun, M. M. (1997). Types and tokens in visual processing: A double dissociation between the attentional blink and repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 738–755.

Chun, M. M., & Cavanagh, P. (1997). Seeing two as one: Linking apparent motion and repetition blindness. *Psychological Science*, *8*, 74–78.

Coltheart, V., Mondy, S., Moore, J (2001). *Effects of repetition on the report of RSVP sequences of familiar and novel objects*. Paper presented at the Vision Sciences Society, Sarasota, USA.

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.

De Caro, S. A. (1998). On the perception of objects and their orientations. *Spatial Vision*, *11*, 385–399.

De Caro, S. A., & Reeves, A. (2000). Rotating objects to determine orientation, not identity: Evidence from a backward-masking/dual-task procedure. *Perception and Psychophysics*, *62*, 1356–1366.

Deutsch, J. A. (1955). A theory of shape recognition. *British Journal of Psychology*, *46*, 30–37.

Edelman, S. (1999). *Representation and recognition in vision*. Cambridge, MA: MIT Press.

Forster, K. I., & Forster, J. C. (2003). DMDX: A Windows display program with millisecond accuracy. *Behavior Research Methods, Instruments, and Computers*, *35*, 116–124.

Harris, I. M., Dux, P. E (2004). Putting orientation in the picture: A two-stage theory of object recognition. Submitted for publication.

Harris, I. M., Harris, J. A., & Caine, D. (2001). Object orientation agnosia: A failure to find the axis? *Journal of Cognitive Neuroscience*, *13*, 800–812.

Hayward, W. G., Tarr, M. J., & Corderoy, A. K. (1999). Recognizing silhouettes and shaded images across depth rotation. *Perception*, *28*, 1197–1215.

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review*, *99*, 480–517.

Humphreys, G. W., & Riddoch, M. J. (1984). Routes to object constancy: Implications from neurological impairments of object constancy. *Quarterly Journal of Experimental Psychology*, *36A*, 385–415.

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

Kahneman, D., & Treisman, A. M. (1984). Changing views of attention and automaticity. In R. Parasuranam, & D. R. Davies (Eds.), *Varieties of attention* (pp. 29–61). Orlando, FL: Academic Press.

Kahneman, D., Treisman, A. M., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, *24*, 175–219.

Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, *27*, 117–143.

Kanwisher, N. G., Yin, C., & Wojciulik, E. (1999). Repetition blindness for pictures: Evidence for the rapid computation of abstract visual descriptions. In V. Coltheart (Ed.), *Fleeting memories: Cognition of brief visual stimuli* (pp. 119–150). Cambridge, MA: MIT Press.

Karnath, H.-O., Ferber, S., & Bülthoff, H. H. (2000). Neuronal representation of object orientation. *Neuropsychologia*, *38*, 1235–1241.

Lawson, R., & Jolicoeur, P. (1998). The effects of plane rotation on the recognition of brief masked pictures of familiar objects. *Memory and Cognition*, *26*, 791–803.

Lawson, R., & Jolicoeur, P. (2003). Recognition thresholds for plane-rotated pictures of familiar objects. *Acta Psychologica*, *112*, 17–41.

Marr, D. (1980). Visual information: The structure and creation of visual representations. *Philosophical Transactions of the Royal Society of London—Series B: Biological Sciences*, *290*, 199–218.

Marr, D., & Nishihara, H. K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London Series B: Biological Sciences*, *200*, 269–294.

Murray, J. E. (1995). Imagining and naming rotated natural objects. *Psychonomic Bulletin and Review*, *2*, 239–243.

Palmer, S. E., Rosch, E., & Chase, P. (1981). Canonical perspective and the perception of objects. In J. Long, & A. Baddeley (Eds.), *Attention and performance IX*. Hillsdale, NJ: Erlbaum.

Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Perception and Performance*, *6*, 174–215.

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 and Review*, *2*, 55–82.

Tarr, M. J., & Bülthoff, H. H. (1995). Is human object recognition better described by geon structural descriptions or by multiple views? Comment on Biederman and Gerhardstein (1993). *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 1494–1505.

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

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*, 97–136.

Turnbull, O. H., Beschin, N., & Della Sala, S. (1997). Agnosia for object orientation: Implications for theories of object recognition. *Neuropsychologia*, *35*, 153–163.

Turnbull, O. H., Della Sala, S., & Beschin, N. (2002). Angosia for object orientation: Naming and mental rotation evidence. *Neurocase*, *8*, 296–305.

Turnbull, O. H., Laws, K. R., & McCarthy, R. A. (1995). Object recognition without knowledge of object orientation. *Cortex*, *31*, 387–395.

Ullman, S. (1989). Aligning pictorial descriptions: An approach to object recognition. *Cognition*, *18*, 193–254.

Warrington, E. K., & James, M. (1986). Visual object recognition in patients with right-hemisphere lesions: Axes or features? *Perception*, *15*, 355–366.