Perception, 2003, volume 32, pages 1465 – 1498

DOI:10.1068/p5031

Interactions between view changes and shape changes in picture – picture matching

Rebecca Lawson, Heinrich H Bülthoff¶, Sarah Dumbell
Department of Psychology, University of Liverpool, Eleanor Rathbone Building,
Bedford Street South, Liverpool L69 7ZA, UK; e-mail: rlawson@liverpool.ac.uk; ¶ Max-Planck-Institut
für biologische Kybernetik, Spemannstrasse 38, D 72076 Tübingen, Germany
Received 9 December 2002, in revised form 19 May 2003; published online 24 December 2003

Abstract. Four experiments are reported in which pictures of different morphs of novel, complex, 3-D objects, similar to objects which we must identify in the real world, were presented. We investigated how changes of viewpoint influence our ability to discriminate between morphs. View changes had a powerful effect on performance in picture – picture matching tasks when similarly shaped morphs had to be discriminated. Shape changes were detected faster and more accurately when morphs were depicted from the same rather than different views. In contrast, view change had no effect when dissimilarly shaped morphs had to be discriminated. This interaction between the effects of view change and shape change was found both for simultaneous stimulus presentation and for sequential presentation with interstimulus intervals up to 3600 ms. The interaction was found after repeated presentations of the stimuli before the matching task and after practice at the matching task as well as after no such pre-exposure to the stimuli or to the task. The results demonstrate the difficulty in activating abstract, view-insensitive representations to help to achieve object constancy, even when matching over long interstimulus intervals or after stimuli have already been seen many times.

1 Introduction

In this paper we examine interactions between the effects of view changes and shape changes on object recognition. The human visual object-recognition system must achieve object constancy over view changes in order to reliably identify objects as members of a class of familiar stimuli, such as dogs or chairs or apples. To do this, it must typically ignore considerable variation in the visual input due to semantically irrelevant changes in the viewing conditions, such as alterations in the lighting conditions and in the position of the observer relative to the object. Achieving visual object constancy allows us to generalise over variation in the visual input, which should be ignored because such variation does not alter the identity of an object.

In contrast to these semantically irrelevant view changes, a shape change often does affect important semantic-category information about an object. Quite small shape changes can alter a drawing pin to a nail, a mug to a jug, a stool to a table, and so on. Not all shape changes are important—many minor shape changes are semantically irrelevant and so should be ignored, such as the variability in the shape of wine glasses, trees, and knives. There are also some major shape changes which should be generalised over. For example, animals change their shape radically as they move to different positions, yet we need to identify a running dog and a sitting dog as belonging to the same semantic category. Overall, though, unlike view changes, shape changes are often semantically meaningful.

There will often be conflict in the attempt to optimise the achievement of these two goals of the visual system, namely achieving object constancy by generalising over semantically irrelevant view changes whilst being sensitive to shape changes in order to discriminate between semantically distinct categories of objects. If the differences in the visual input due to view changes were distinct in nature from the differences caused by shape changes, or much smaller, then this trade-off could be avoided.

Unfortunately, though, view changes and shape changes often result in similar effects on the visual input stimulus. Rotations in depth can alter the global outline shape of an object and the visibility, size, and apparent spatial relations between parts of the object, all of which can also result from shape changes. Any attempt to improve the process of achieving object constancy will therefore risk reducing the accuracy of object-shape discrimination and vice versa. The visual system must therefore try to reach an appropriate compromise in achieving these two aims. In this paper, we examine empirical evidence for such a trade-off.

As noted by Cutzu and Edelman (1998), most studies of object recognition and categorisation to date have either manipulated view changes whilst maintaining shape constant, or have manipulated shape changes whilst maintaining view constant. In addition, whilst studies in the former area of research in object recognition have tested a range of stimuli from complex, naturalistic, 3-D objects to more simple stimuli, in the latter area of object-categorisation research, most studies have used only simple, impoverished stimuli such as colour patches or geometric shapes. As a result, there are few data on the effects of view changes on the ability to discriminate shape changes for relatively complex, realistic stimuli. Recently, though, these issues have started to be examined empirically (eg Biederman and Bar 1999; Cutzu and Edelman 1998; Edelman 1995; Edelman and Duvdevani-Bar 1997; Foster and Gilson 2002; Gauthier and Tarr 1997; Hayward and Williams 2000). Note, however, that only in the studies by Edelman and colleagues, and by Foster and Gilson was a series of shape changes manipulated systematically to allow comparison of the effects of different magnitudes of shape change. For example, in Biederman and Bar (1999), participants' abilities to detect quantitative changes and changes to non-accidental properties of a target novel object were compared in a sequential matching task. Six simple, two-volume objects were chosen for which the quantitative and non-accidental changes were equally detectable when targets and distractors were presented from the same view. For these objects, quantitative changes were harder to detect than non-accidental changes, if there was a depth rotation between the target and the distractor. However, the physical similarity in shape of the target and distractor was not systematically manipulated in these studies.

View sensitivity can be defined as a difficulty or even a failure to achieve object constancy when there are view changes across stimuli which must be identified as being from the same category. View sensitivity is therefore a measure of the difficulty of achieving object constancy and so of the effect of view changes on our ability to recognise objects. It has often been assumed that view sensitivity increases as the difficulty of shape discrimination increases (for example, if an animal has to be identified as a dog versus a cat rather than as a dog versus a table). There is some evidence (though to date surprisingly little) in support of this hypothesis. For example, studies of picture – word verification of plane-rotated views of familiar objects have shown increased view sensitivity as the visual similarity between the shapes of competing response alternatives increased (Lawson and Jolicoeur 1998, 2003; Murray 1998). Here, upright and, to an even greater extent, misoriented stimuli were harder to identify when the response alternatives were names of objects which were difficult to discriminate visually. Thus the effect of view changes interacted with the ease of shape discrimination. A similar conclusion was drawn from a study of the orientation priming of novel, 2-D shapes by Gauthier and Tarr (1997). Edelman (1995) presented different, depth-rotated views of morphs which lay between two endpoint objects that were complex, animal-like, 3-D novel shapes. In training, participants learnt to discriminate between two morphs which were seen from a limited range of depth-rotated views. They were then tested on their ability to transfer this training (without feedback) to novel views of the two training morphs. Across different blocks, the shape similarity between the two training morphs was varied. In three studies, Edelman found a non-significant trend for greater view sensitivity at transfer when the training morphs were more similar in shape, ie for an interaction between view-change and shape-change effects (see also Tjan and Legge 1998). In a simultaneous picture-matching task in which novel objects like bent tubes were presented, Foster and Gilson (2002) reported an interaction between view-change and shape-change effects for long (2000 ms) though not for short (100 ms) presentation durations. At the long presentation duration, sensitivity to view changes up to 45° appeared to increase when objects were more similar in shape (personal communication). Foster and Gilson reported little effect on performance of increasing the view change beyond 45° in any condition.

In contrast, Hayward and Williams (2000) failed to find an interaction between the effects of view changes and the difficulty of discriminating shape changes. They manipulated the ease of discriminating views of novel, complex objects across blocks in a speeded sequential picture-matching task. In blocks when in mismatch trials only visually dissimilar pairs of items were presented there were increases in both reaction times (RTs) and error rates on match trials as the depth rotation between the two pictures increased. Importantly, this view sensitivity for match trials was not less than that found in blocks when in mismatch trials only pairs of similar items were presented. Hayward and Williams (2000) concluded that view sensitivity is generally not affected by varying the context of shape discriminability by manipulating shape similarity on mismatch trials (except when view insensitive, distinguishing information is readily available, as was the case in their first study, where stimuli could be discriminated by colour in the visually dissimilar context condition). Instead they suggested that the intrinsic geometry of objects determined view effects independently of the task context.

Overall, there is therefore only relatively weak and inconsistent evidence that view-change and shape-change effects interact, so this issue needs to be examined further. Furthermore, both Edelman (1995) and Hayward and Williams (2000) examined the effects on view sensitivity of manipulating the similarity of shapes which had to be discriminated across different blocks. In contrast, in the present studies (like those of Foster and Gilson 2002) we examined the effects on view sensitivity of manipulating the size of shape changes within a block. Although the tasks and stimuli in these studies may appear similar, the studies address different issues. In the former studies, view sensitivity for a given object was compared across different contexts, whilst in the latter studies view sensitivity across different objects was compared in the same context. Hayward and Williams's (2000) studies suggest that view sensitivity is not influenced by the overall difficulty within a block of detecting whether two shapes are identical, whilst Edelman's (1995) studies suggest that view sensitivity is important. The discrepancy between these findings needs to be resolved. But, in the studies reported here we investigate a complementary issue, namely how view sensitivity is influenced by shape similarity given a constant overall level of difficulty of shape discrimination within the context of the task.

Recently, 3-D morphing has been used to manipulate shape changes systematically. Usually pairs of visually similar, familiar objects have been used as two endpoints objects, and a series of intermediate, morphed objects have been produced, with each morph varying in shape in a systematic way from the others in the series. To date, this has been done most often with faces (eg Beale and Keil 1995; Campanella et al 2001), but it is possible to do it with other objects. Such stimuli can be labour intensive to produce, as the morphing process typically requires that many equivalent points across the two endpoint objects be identified by the experimenter (such as the tips of the ears and the middle of the back if morphing between a horse and a donkey) in order to map between the two shapes. However, recent developments in automatic 3-D morphing software may alleviate this problem (Shelton 1998).

An alternative approach to morphing was used by Cutzu and Edelman (1998) and was adopted in the present studies. Cutzu and Edelman tested a range of morphs, all of which were derived from a single, complex, animal-like, novel object. Each morph was defined by values of 70 parameters which specified the shape of the object parts and their spatial relations relative to the other object parts. Sets of morphs were presented to participants, with the relative shapes of the morphs selected to be arranged in a known configuration (such as a triangle) in the abstract shape space defined by these 70 stimulus-defining parameters. Participants rated the perceived similarity of different pairs of morphs which were presented simultaneously. Each morph was shown rotating continuously in depth. Nonmetric multidimensional scaling techniques were used to transform the ratings of morph similarity into psychological (not physical) distances in a 2-D similarity space. The configuration of positions of the morphs in this 2-D psychological space was consistently found to reflect their positions in the original parameter space (such as a triangle). This suggests that manipulations of the physical parameters systematically influenced the perceived similarity of the morphs, and that participants used this information to extract the parameterised shape similarities which had been imposed on the morphs.

Cutzu and Edelman (1998) also ran speeded sequential matching studies with pictures of two different views of either the same or of different morphs. Participants received no prior training with the stimuli and no feedback during the task. Nevertheless, as for the unspeeded rating task described above, the results of the matching task allowed a recovery of the configuration of the morph stimulus positions which had been set in the parameterised physical shape space from the positions of the stimuli in the psychological similarity space. Importantly, the positions of different views of the same morph were generally clustered together in the 2-D psychological space and were well separated from clusters of views of other morphs.

In addition, Cutzu and Edelman (1998) once again recovered the parameterised shape space from human psychophysical data using the results of a long-term-memory rating task. Here, participants first learnt to associate novel words to different test morphs. They then had to rate the similarity of pairs of the test morphs to each other given only the verbal labels of the morphs ("is A more similar to B than C is to D?"). In this task, the test morphs were never seen together. During training, only one test morph was seen during a given training session, presented from a range of views rotated in depth. Participants learnt to discriminate that test morph from a set of different views of non-test distractor morphs which were similar to the test morph. Participants were not told that they would subsequently have to compare the different test morphs. During testing, participants were forced to rely on their long-term visual memory of the test morphs since no pictures were presented. This task therefore required access to stable memory representations, whereas the rating and sequential matching tasks could have relied on only transient perceptual representations. Nevertheless, the three tasks produced similar results. Cutzu and Edelman (1998) argued that the consistency of results across the different tasks suggests that all three tapped common representations which were structured according to similarity.

Cutzu and Edelman (1998) also tested two models of similarity. The first was an image-based model. This was found to group the morphs by similarity in view rather than shape, indicating that for the animal-like morphs tested similarity in view was more salient in the image than similarity in shape. This model failed to replicate the human participants' perceived similarity shape space derived from multidimensional scaling. Their second model was a set of radial-basis-function networks, each of which was initially trained to associate together all views of a particular morph. When tested with novel stimuli, this model, like the human participants, produced a similarity measure which, when analysed with multidimensional scaling, resembled the parameterised morph space.

The results from Cutzu and Edelman's (1998) studies suggest that humans can represent the similarity in shape between novel morphs of complex objects in a view-insensitive manner, ie they can achieve object constancy. This was in spite of view changes apparently having at least as great an effect on the image as shape changes, as evidenced by the results of their first, image-based model which grouped morphs by similarity in view (see also Edelman 1998; Edelman and Duvdevani-Bar 1997). They did not, though, experimentally test the relationship between view changes and shape changes. In both their short-term and long-term rating studies, participants were exposed to many different views of the morphs, and only overall similarity between pairs of morphs was rated (ie averaged over all views). In the sequential matching studies, all trials presented two pictures which depicted morphs from very different views in depth. Thus in none of Cutzu and Edelman's (1998) studies were the effects of varying the extent of the view change examined. The results of their studies indirectly suggest that participants developed view-insensitive stimulus representations, but this issue needs to be tested directly. In the present studies, we examined the relation between view changes and shape changes to determine whether the size of view changes influences our ability to detect shape changes.

Like Cutzu and Edelman (1998), we presented different views of sets of morphs of novel objects. The objects were derived from familiar, everyday objects with varied, complex shapes, and were intended to resemble everyday objects more closely than most novel stimuli that have been previously used. The objects were not generated from a simple alphabet of geometric parts and varied in their number of parts, and in the size, shape, and spatial arrangement of parts. We used morphs derived from thirty objects, whereas Cutzu and Edelman (1998), and Edelman (1995) used morphs derived from just one object.

We produced thirty sets of morphed stimuli (see figure 1). Each set was derived from a single, 3-D model of a real object such as a cannon, a rollerskate, or a dog. The selected model of a real object was then morphed to produce the endpoint morph, S1. S1 did not closely resemble the original object model, although in some cases it could be identified as being derived from the model. S1 was then morphed in a different way to create a series of twelve morphs from S2 to S13, such that the nature and degree of physical shape change produced by the morphing changed incrementally from each morph to the next in the series (see figure 2a). The nature and extent of the shape change between S1 and S13 differed for each of the thirty objects. All of the object or separate parts of the object could be squashed, expanded, or twisted. Finally, for each of the thirteen morphs, S1 to S13, twelve different views in depth were produced, with each view being rotated by 30° from the next.

Since the shape changes across each set of thirteen morphs were well controlled, these stimuli allowed us to explore the effects of incremental shape changes for each of the thirty different novel objects. It is, nevertheless, important to emphasise that the size and nature of the perceived change in shape probably varied somewhat from one pair of morphs to the next in each morph series. It is likely, though, that the perceived visual similarity in shape between the S1 morph and the other morphs monotonically decreased from S2 to S13, even if the magnitude of the reduction in similarity from one morph to the next varied somewhat. Cutzu and Edelman's (1998) rating studies measuring perceived similarity for different morphs also support this assumption.

We investigated how view changes caused by depth rotation influence the sensitivity of participants to detect shape changes across the S1 to S13 morphs for this diverse set of thirty objects (see figure 1). Four picture-matching studies used both simultaneous and sequential presentation of morphs, depicted from the same and from different views. On match trials, pictures of the same morph (eg both S1) of the same novel object were presented. On mismatch trials, pictures of two different morphs (eg S1

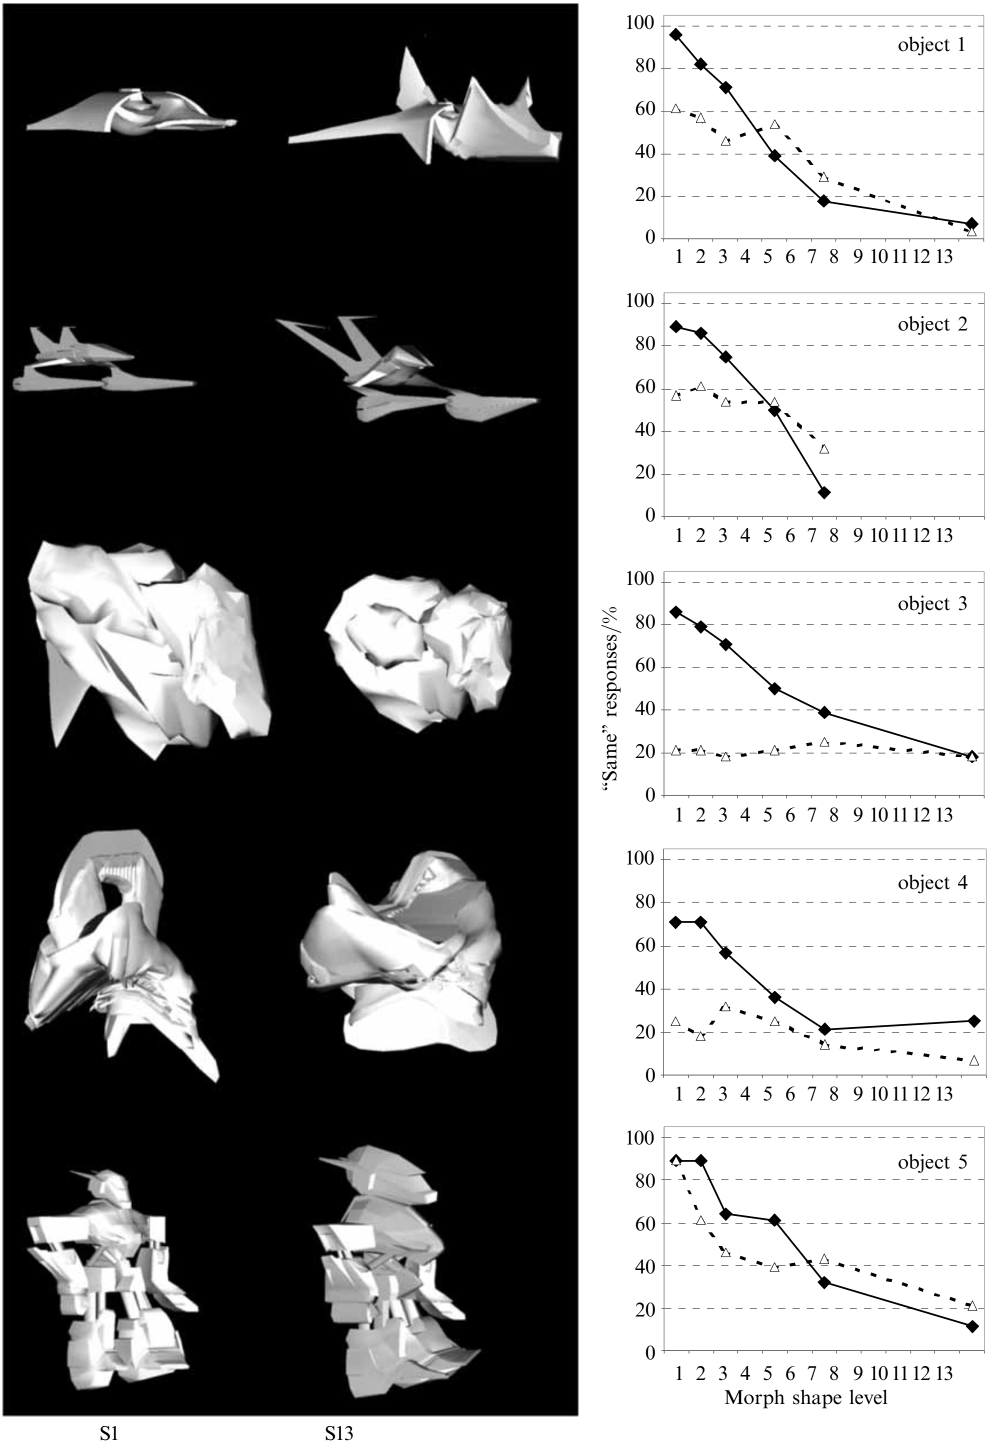

Figure 1. Pictures of the 30° view of the S1 morph (on the left) and the S13 morph (on the right) of the thirty objects used in the studies in this paper. To the right of each pair of pictures of morphs of a given object is a graph with the results from experiment 1 for that object. Each graph shows the mean percentage of "same" responses on same-view trials (0° view change, solid line) and different-view trials (90° view change, dotted line) when a picture of S1 was seen with another

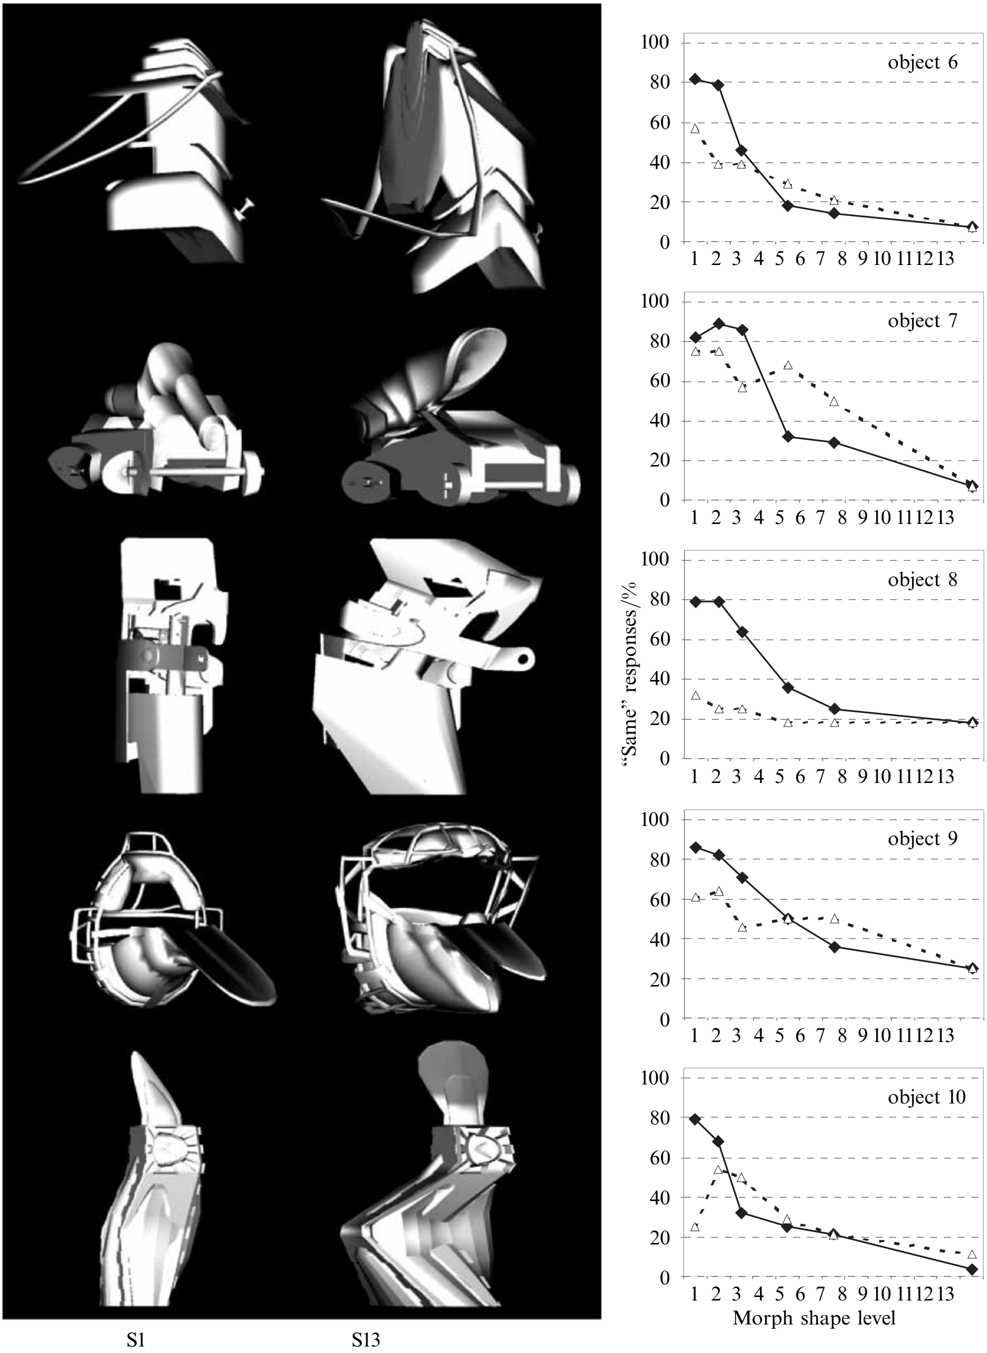

Figure 1 (continued)
picture of S1 on match trials (where "same" shape was the correct response) or with S2, S3, S5, S7, or S13 on mismatch trials (where "same" shape was the wrong response) during simultaneous picture matching. Because of a programming error, results for S1/S13 mismatch trials for object 2 only are not plotted. These graphs are in the same format as figure 3 which shows overall performance in experiment 1.

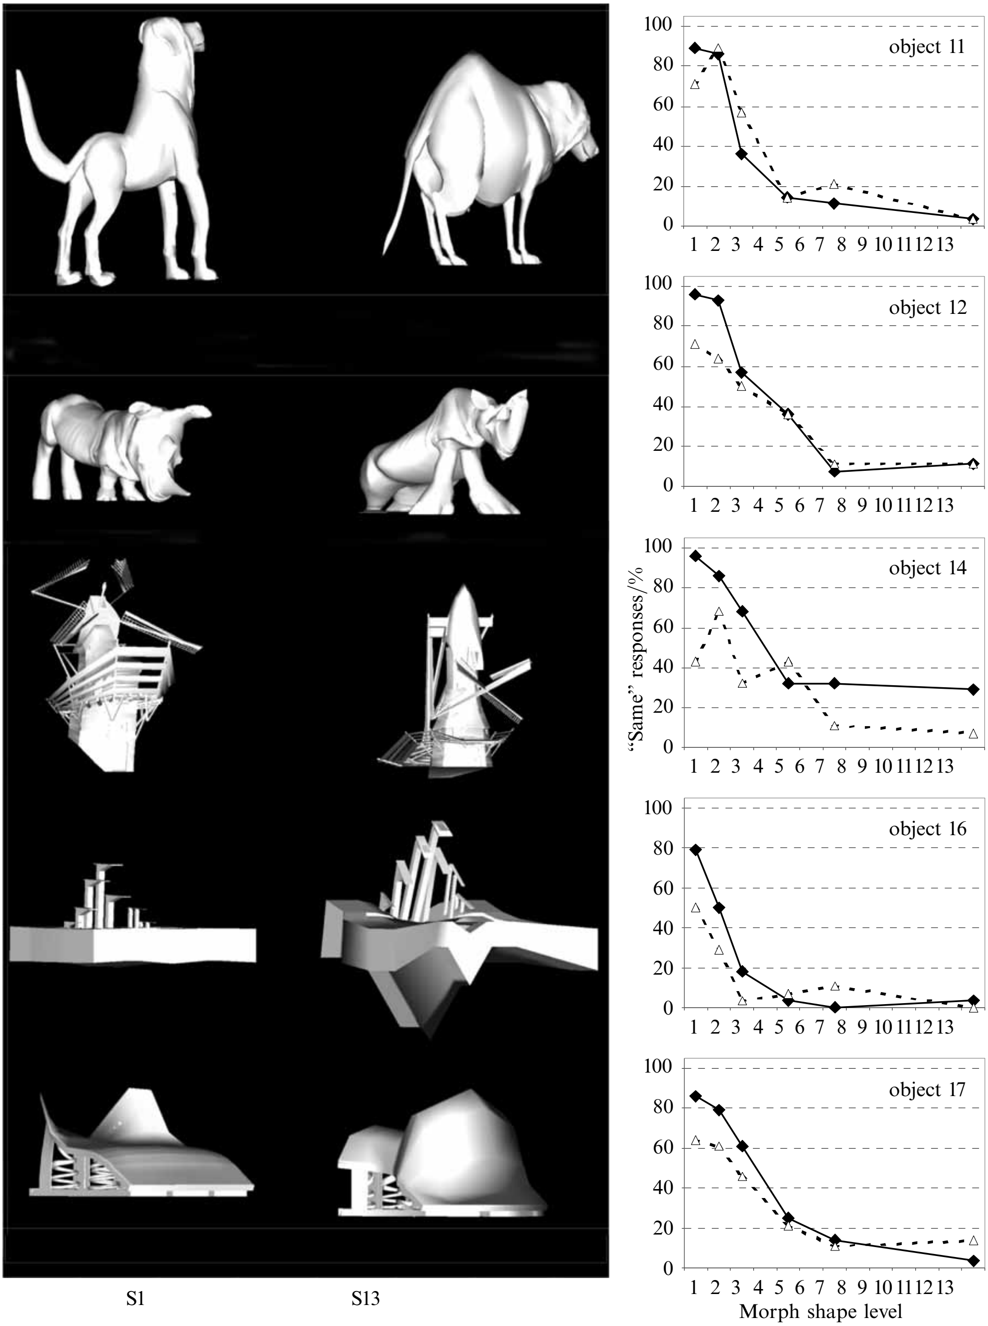

Figure 1 (continued)

and S13) of the same novel object were presented (see figures 2b and 2c). On mismatch as well as match trials, morphs derived from the same novel object were depicted. Thus, unlike most studies of the effects of view change on picture matching for familiar objects (eg Bartram 1976; Ellis and Allport 1986; Ellis et al 1989; Lawson and Humphreys 1996), results from mismatch as well as match trials are readily interpretable and theoretically interesting. Measuring view change across pictures of two different objects

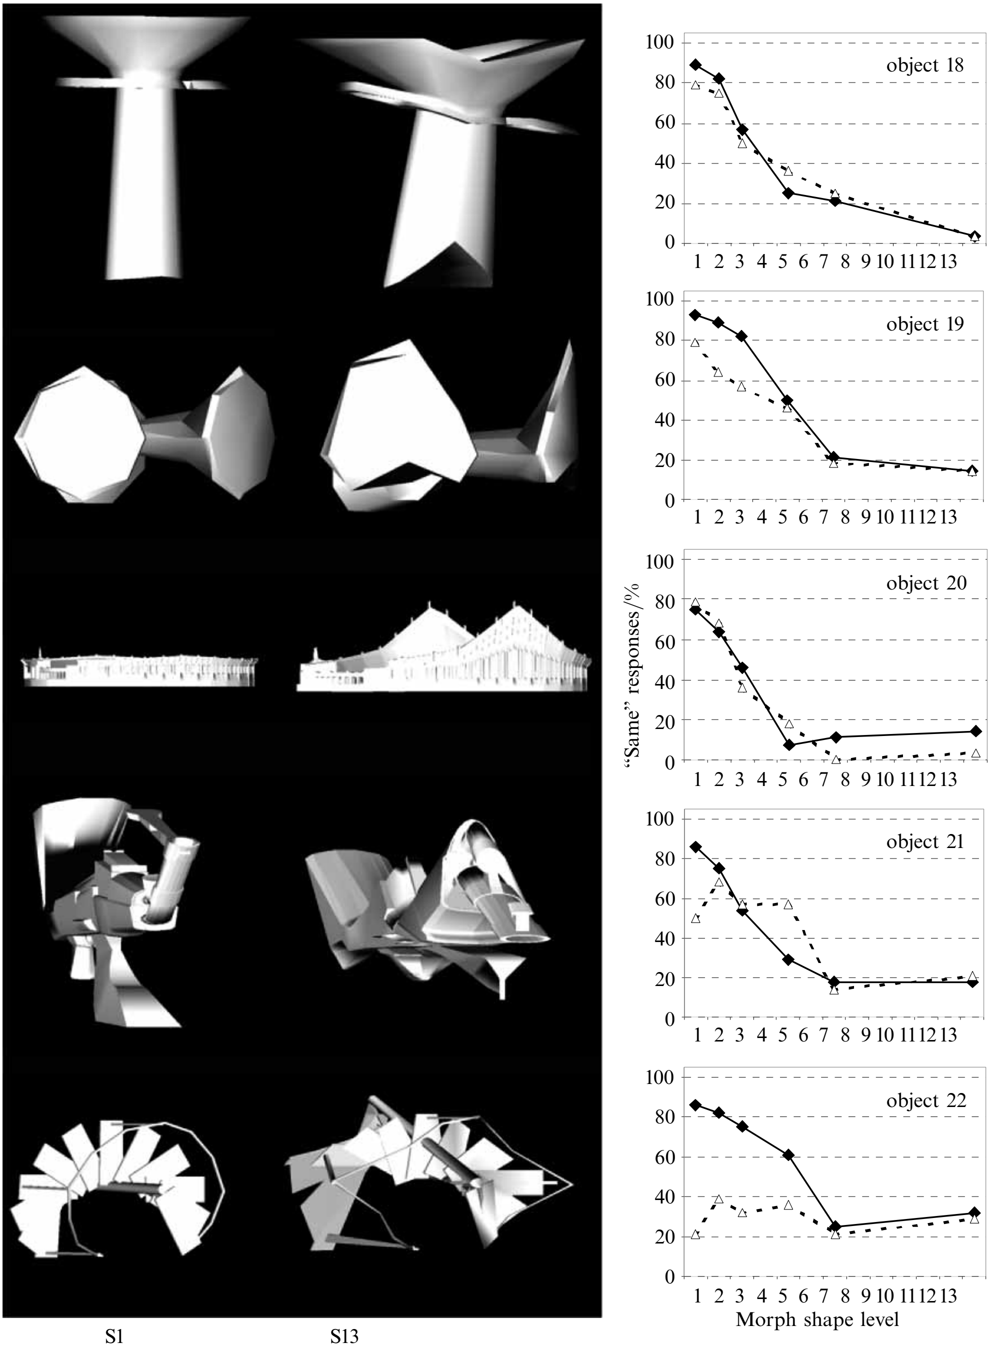

Figure 1 (continued)

is meaningful only when the depicted view of both objects can be described by using the same reference frame. This was the case for the two morphs which were presented on mismatch trials here, which always shared the same general shape (see figure 2). In contrast, in many picture-matching studies to date, on mismatch trials two objects selected at random from a diverse set of stimuli were presented, so the mean shape

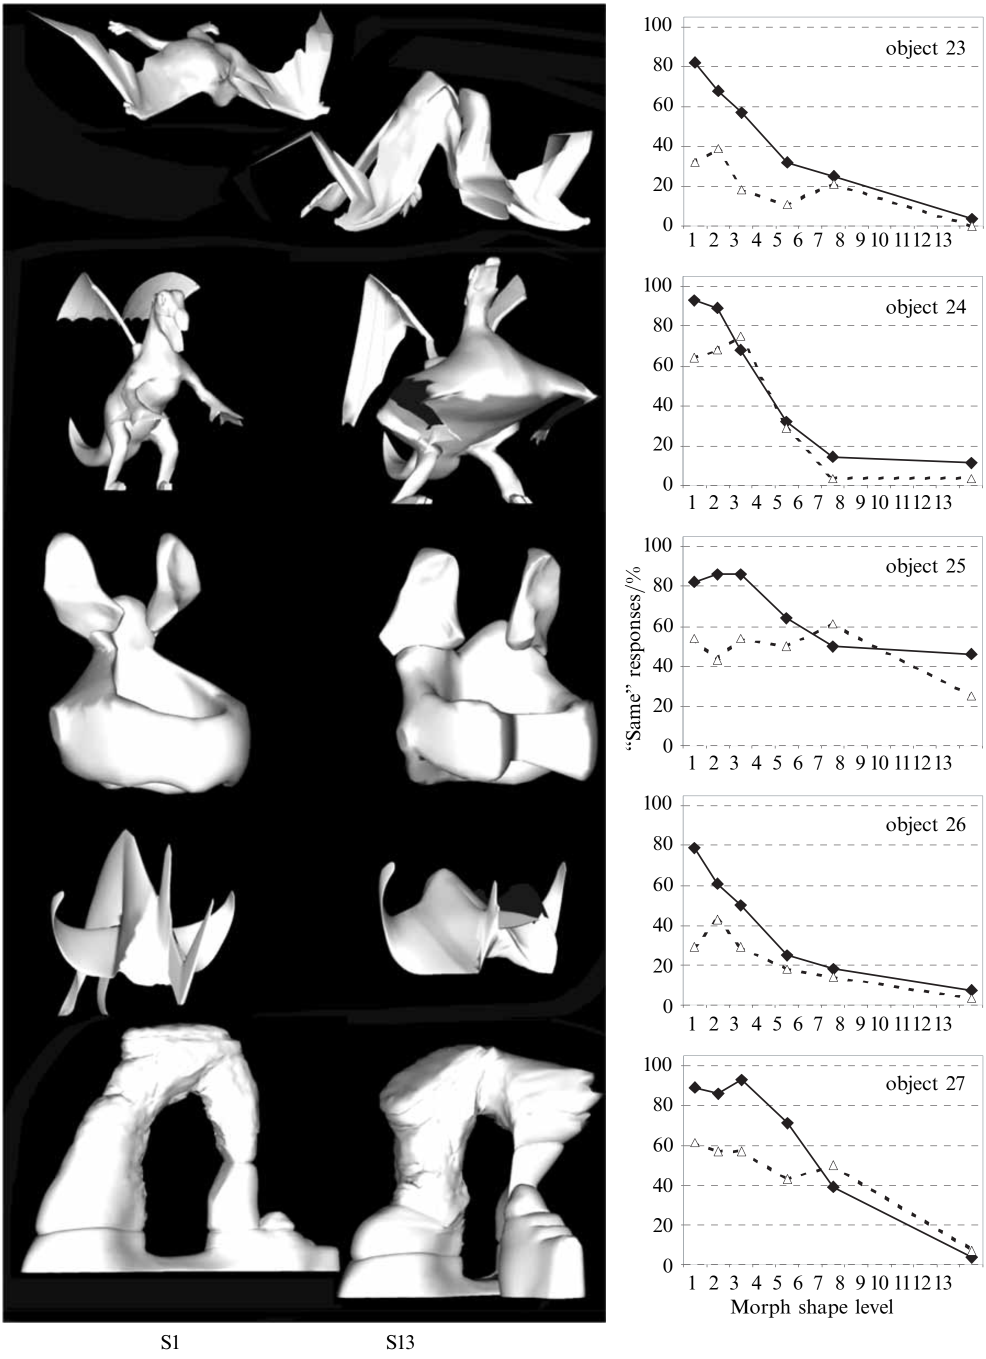

Figure 1 (continued)

similarity between the objects was low, and similarity varied greatly and in an uncontrolled way from trial to trial. For example, a mismatch trial might present a side view of a dog followed by a front view of a banana. A measure of view change on such a trial is almost meaningless, as there is little or no relation between the front views (or any other views) of two objects with such different shapes.

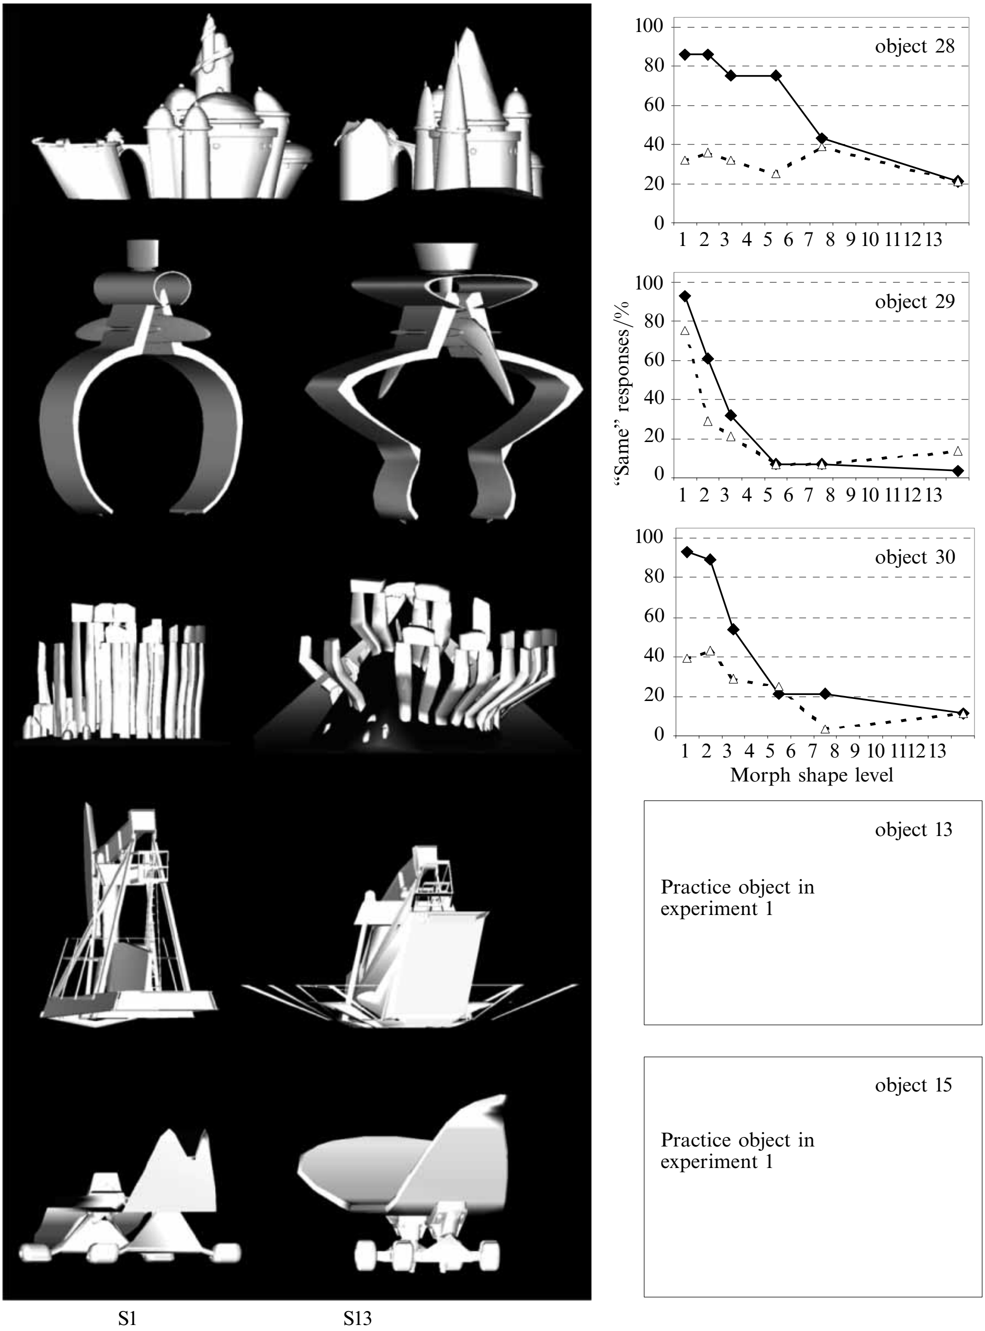

Figure 1 (continued)

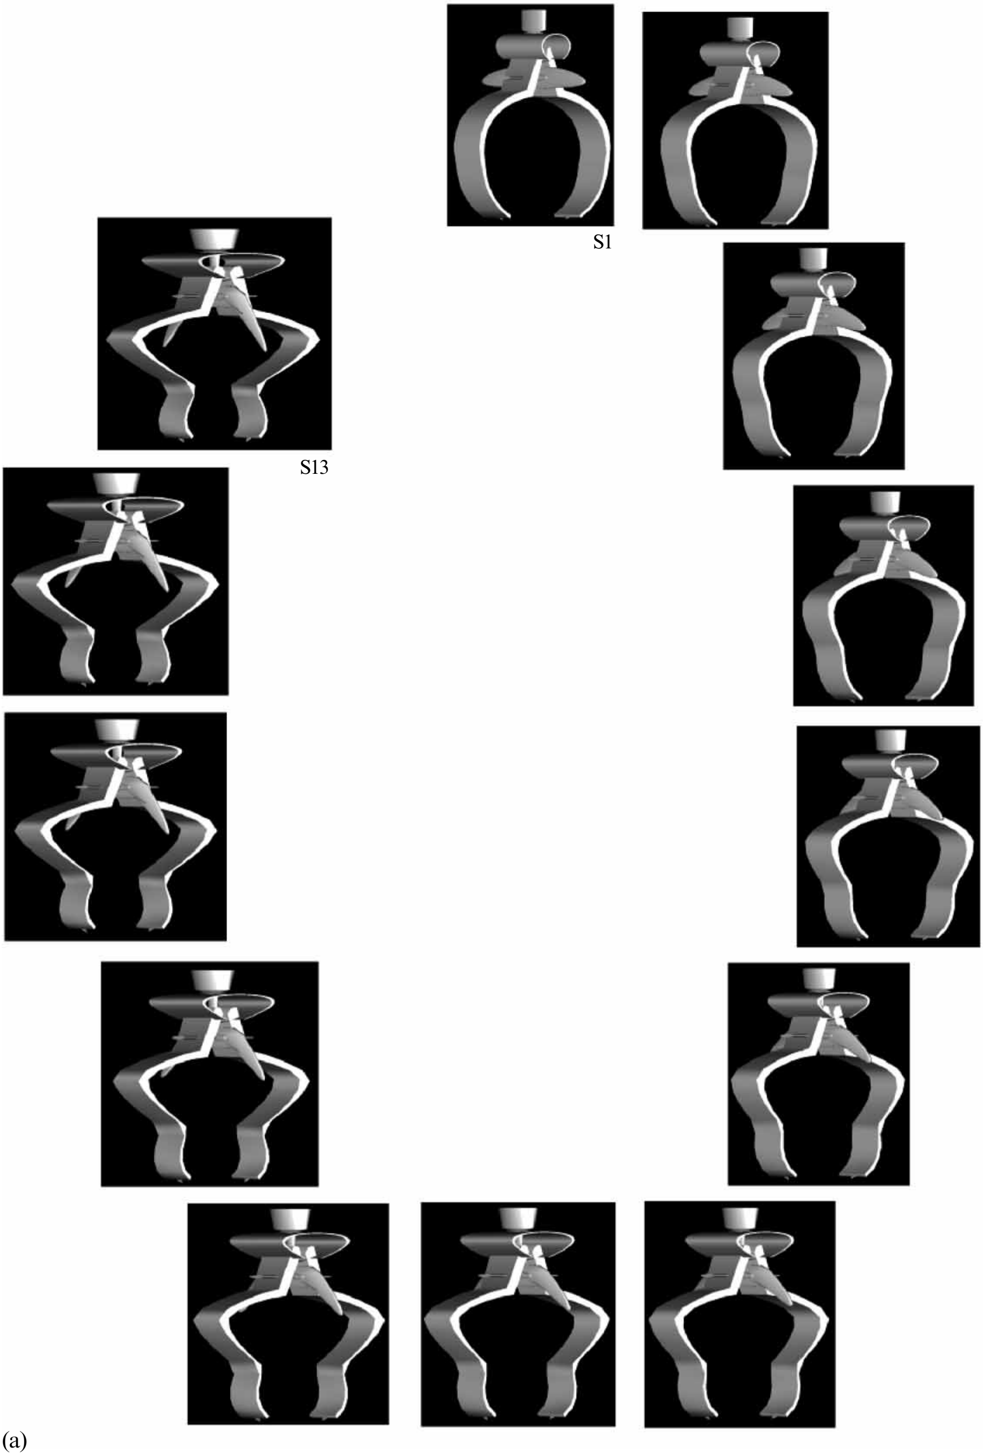

Figure 2. Pictures of three objects which were presented in all of the studies described in this paper: (a) the thirteen different morphs of object 29, all depicted from a 30° view, starting, top centre, with S1 and finishing, uppermost left, with S13; (b) object 25 and (c) object 9, depicted from 30°, 60°, 120°, and 240° views (from top to bottom, respectively) for the S1, S7, and S13 morphs of each object (in the left, centre, and right columns, respectively).

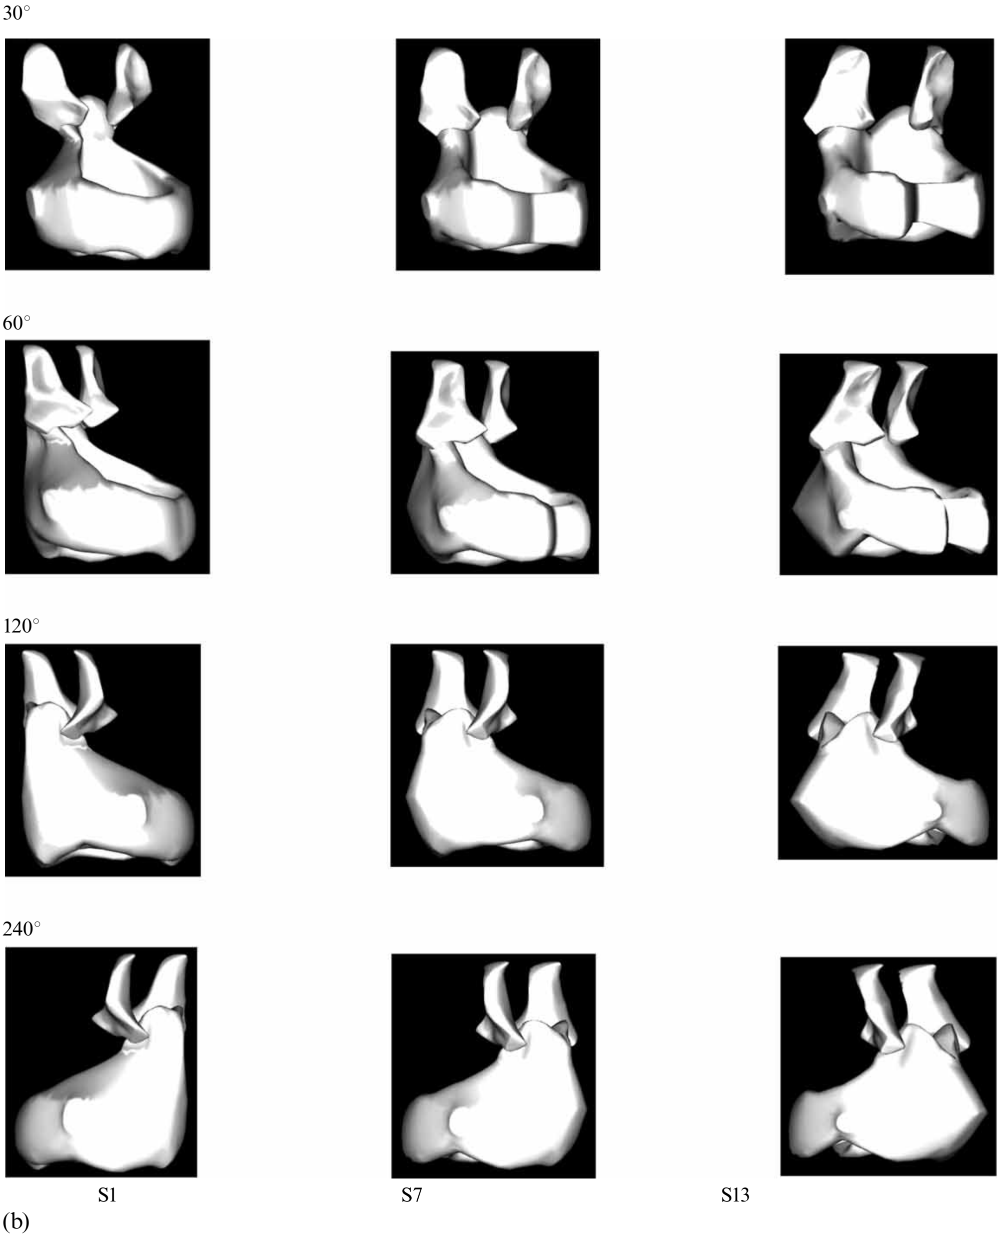

Figure 2 (continued)

2 Experiment 1

In experiment 1, we investigated the relation between the effects of view changes and shape changes on the accuracy of simultaneous picture matching. Lawson and Jolicoeur (1998, 2003) and Murray (1998) found that making shape discrimination harder by increasing visual similarity in turn increased view sensitivity. We therefore predicted that, as the shape change between two objects decreased, making shape discrimination harder, view sensitivity would increase: in experiment 1, we examined the extent of any such changes in view sensitivity over a wide range of shape changes. This prediction runs counter to the prediction which could be derived from the results of Cutzu and Edelman (1998). Their results suggested that the representation of shape similarity is

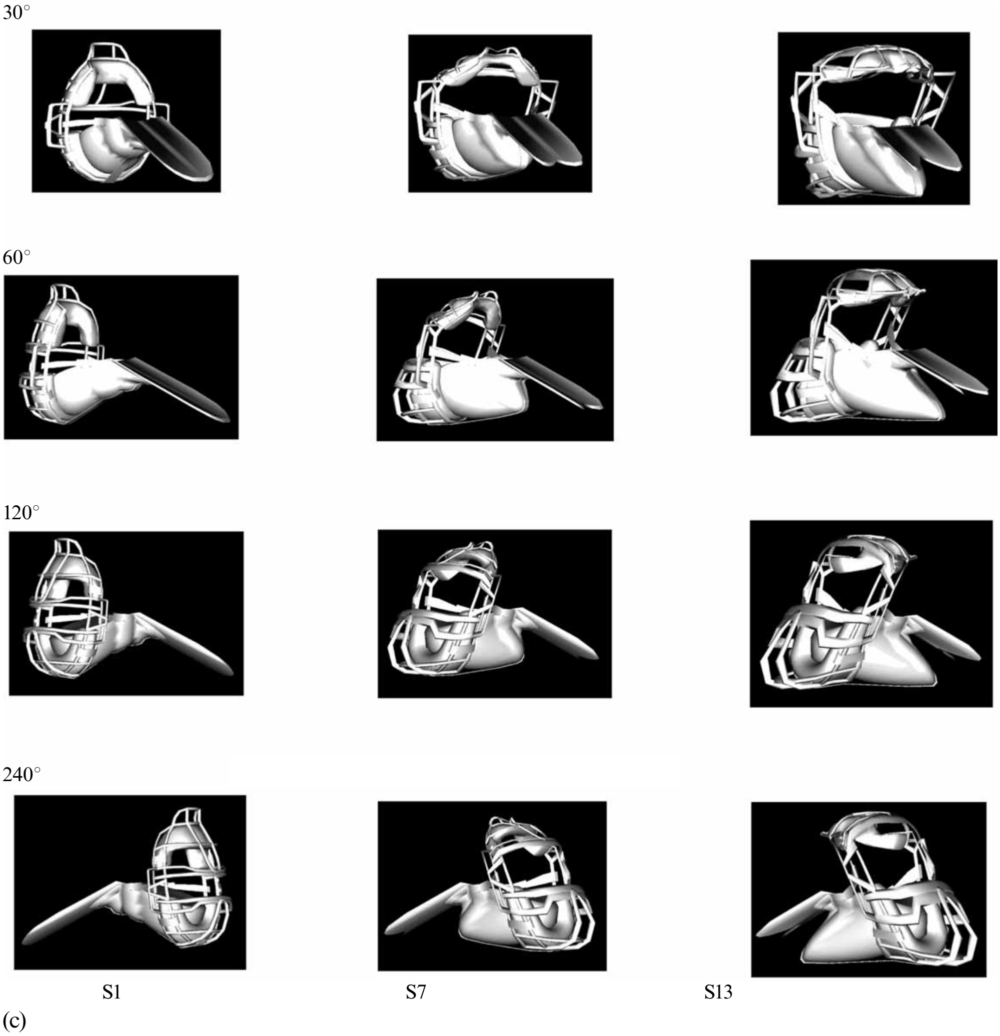

Figure 2 (continued)

view insensitive, which could lead to the prediction that view-change and shape-change effects would not interact. As discussed above, though, note that the effects of varying view change on shape discrimination was not actually tested in Cutzu and Edelman's (1998) studies (see also Edelman 1995, 1998; Edelman and Duvdevani-Bar 1997). It is therefore important to examine directly the combined effects of view changes and shape changes (see also Foster and Gilson 2002).

In experiment 1, participants were simultaneously shown two pictures, both depicting morphs derived from the same object. On "same" match trials, both pictures depicted the same S1 morph of the object. On "different" mismatch trials, one picture showed the S1 morph whilst the other picture showed the S2, S3, S5, S7, or the S13 morph. The most dissimilar pairs of stimuli were seen on S1/S13 mismatch trials (see figures 2b and 2c). On both match and mismatch trials, the two pictures could depict the object from the same view or from different views rotated 90° in depth from each other. Except during practice, participants were given no feedback as to the correct response on a given trial. We expected participants to make many errors on similar mismatch trials (S1/S2 and S1/S3), and we were concerned that feedback on such trials might confuse or dishearten participants who may have been unable to detect a difference between the morphs on these similar mismatch trials and who would not have understood why their responses were incorrect. Participants were not put under time pressure to respond, as we wanted to determine their optimal performance in trying to achieve object constancy across depth rotations whilst discriminating between similarly shaped stimuli.

2.1 *Method*

2.1.1 *Participants.* Fifty-six participants volunteered to take part in the study. In this and the following studies, participants were undergraduate students from the University of Liverpool who were naïve as to the purpose of the study.

2.1.2 *Materials.* In this and the following experiments, thirty textured, shaded, 3-D models of familiar objects were individually morphed in varied and arbitrary ways to generate a set of thirty endpoint S1 morphs. The stimuli were modelled and rendered with the SoftImage™ animation software package. Each S1 shape was defined by a large number of polygons. For a given object, its S1 shape was morphed to produce an S13 morph. Morphing was achieved by deforming a lattice which was placed around S1. Deforming the lattice altered the position of the S1 polygon vertices which changed the shape of the object to produce S13. A further eleven morphs, S2 to S12, were produced by incrementally changing the S1 shape into the S13 shape by linear interpolation (see figure 2a). This produced a total of thirteen morphs of each object, labelled from S1 to S13, where S13 had the most dissimilar shape to S1. For each morph, grey-scale pictures of twelve different views were then produced by horizontally rotating the shape in depth. The 0° view was assigned to be a foreshortened view if the object had an elongated shape and was otherwise arbitrarily assigned. Foreshortened views can be particularly difficult to identify (Lawson 1999) so we avoided presenting them. From the 0° view, each successive view was rotated by 30° about the vertical axis running through the midpoint of the shape. These views were labelled as 30°, 60°, 90°, 120°, and so on up to 330°. Altogether there were 156 pictures (thirteen morphs, each depicted from twelve views) of each of the thirty objects. The stimuli were presented against a black background inside a window measuring 450 mm by 450 mm on the computer screen and subtended a visual angle of approximately 5 deg.

Twenty-eight objects were presented in the experimental trials and the remaining two objects were presented in the practice trials (object 13 and object 15). Figure 1 shows the 30° view of S1 and S13 for all the objects. Each object was represented by six morphed versions, S1, S2, S3, S5, S7, and S13. Each morph was, in turn, depicted by two views (30° and 120°) which were rotated by 90° in depth from each other (see the first and third rows in figures 2b and 2c). Thus 12 pictures of each object (six morphs by two views) were presented. Note that the influence of a given view change is not straightforwardly related to the size of that view change, even if unusual viewpoints are avoided.

2.1.3 *Design.* Fourteen participants were randomly allocated to each of four groups. Two groups were presented with object set 1 (comprising fourteen of the novel objects) and two groups saw object set 2 (comprising the remaining fourteen novel objects), with the left picture always being a 30° view of S1 for one of each of the two groups and a 120° view of S1 for the other group. All participants completed one block of 168 trials, of which 28 were match trials and 140 were mismatch trials. There were 2 match and 10 mismatch trials for each of the fourteen objects seen by a given group. The left picture on a trial was always the 30° view or the 120° view of S1. On match trials, the right picture also showed S1, which was depicted once at the same view as the left picture (either 30° or 120°) and once at a different view (either 120° or 30°) for each object. On mismatch trials, the right picture showed S2, S3, S5, S7, or S13, each of which was depicted once at the same view as the left picture (either 30° or 120°) and once at a different view (either 120° or 30°) for each object. In both the practice and the experimental blocks, trials were presented in a different, random order for each participant.

2.1.4 *Apparatus and procedure.* The experiment was run on a Macintosh Power PC G4 computer with the Psyscope version 1.2.5 experimental presentation software. On each trial, a central fixation point appeared for 350 ms. Two pictures of morphs of the same novel object were then immediately presented to the left and to the right of fixation for 1500 ms. Participants made an unspeeded decision whether the two simultaneously presented pictures showed the same shaped stimulus or two differently shaped stimuli by making an 'm' or 'z' keypress, respectively. After the participant had responded, there was an intertrial interval of 1000 ms. Participants were told to ignore any difference in the view at which the left and right stimuli were depicted and they were warned that, on mismatch trials, the left and right pictures might depict two objects with very similar shapes, and so the task was extremely difficult. Participants received no feedback on the correct response for each trial.

Prior to starting the experimental block, all participants completed a block of 25 practice trials. The practice trials were identical to the experimental trials except that they presented morphs of two objects which were not used in the experimental block, participants received feedback on the correct response for each trial, and there were approximately equal numbers of match and mismatch trials.

2.2 *Results*

In this and all the following studies, the results for the F-values in the by-participants and by-items analyses are distinguished by superscripts F^{p} and F^{i}, respectively. The dependent measure was the percentage of trials on which "same" responses were made. On match trials, "same" responses were correct. On mismatch trials, "same" responses were incorrect and represent trials on which participants thought that two pictures of different morphs actually both depicted the same morph.

ANOVAs were conducted on the percentage of "same" responses. There were two within-participant factors: view change (whether the left and right pictures showed morphs from the same or from 90° different views) and shape change (whether the right picture showed the S1, S2, S3, S5, S7, or S13 morph; the left picture always showed the S1 morph). There were two between-participants factors, object set (1 or 2, depending on which set of fourteen objects was presented to a given participant) and left picture view (30° or 120°). The latter two counterbalancing factors of object set and left picture view were included in the analyses as they reduced variance but they were not of interest in this study, and so effects involving these factors are not reported here.

Shape change was significant ($F^{\mathrm{p}}_{5,260} = 334.852$, $p < 0.001$; $F^{\mathrm{i}}_{5,130} = 137.275$, $p < 0.001$). "Same" morph responses decreased as the shape change between the left and right pictures increased. View change was also significant ($F^{\mathrm{p}}_{1,52} = 121.131$, $p < 0.001$; $F^{\mathrm{i}}_{1,26} = 55.003$, $p < 0.001$), with more "same" morph responses (both correct and incorrect) on same-view trials (50.1%) than on different-view trials (35.9%).

Most importantly, the interaction of shape change × view change was significant ($F^{\mathrm{p}}_{5,260} = 42.885$, $p < 0.001$; $F^{\mathrm{i}}_{5,130} = 36.106$, $p < 0.001$). Matching was highly view sensitive when morphs were similar in shape and so difficult to discriminate, but was view insensitive when morphs were dissimilar in shape and so easy to discriminate (see figure 3). The difference between responses on same-view and different-view trials was largest for S1 (32.7%) match trials, less for S2 (27.0%) and S3 (19.9%) mismatch trials, and was small or eliminated for S5 (3.6%), S7 (−0.8%), and S13 (2.9%) mismatch trials.

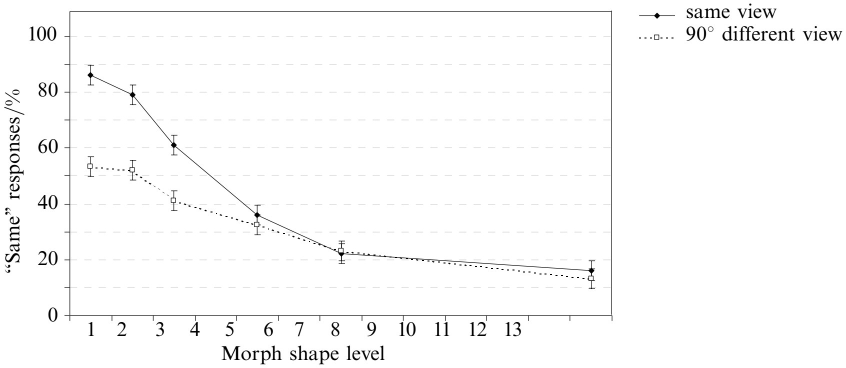

Figure 3. Results from experiment 1: mean percentage of "same" responses on same-view trials (0° view change) and different-view trials (90° view change) when a picture of S1 was seen with another picture of S1 on match trials (where "same" shape was the correct response) or with S2, S3, S5, S7, or S13 on mismatch trials (where "same" shape was the wrong response) during simultaneous picture matching, along with 95% confidence intervals based on the error term for participants for the interaction of shape change × view change (Loftus and Masson 1994).

The above by-participants and by-items analyses were repeated, separately, for each of the four participant groups. In all eight analyses, the main effects of shape change and view change and the interaction of shape change × view change were significant at $p < 0.01$. The pattern of the shape change × view change interaction was the same in all cases, and so the interaction was replicated over different groups of both participants and items, as well as for two different left-picture views. Finally, figure 1 shows the shape change × view change interaction for each individual object tested in experiment 1. The same pattern of results as is shown in figure 3 was found for over two thirds of the objects and only object 20 clearly failed to show the interaction.

Note that this interaction cannot be explained by simply assuming that participants have a bias to respond "same" on same-view trials. A simple bias explanation would predict no effect of shape change. This was not what we observed. Furthermore, participants were increasingly likely to respond "same" on different-view (as well as on same-view) trials as the two morphs depicted on a trial became increasingly similar in shape. It might still be argued that the current results are due to a bias to respond "same" on same-view trials, but that this bias is masked by a ceiling effect for performance on mismatch trials in which dissimilarly shaped morphs were presented. For example, even on same-view trials there were around only 20% incorrect "same" responses for S1/S7 mismatches and just 10% same responses for S1/S13 mismatches, so arguably there was only limited room for improvement on this performance for different-view trials. Here, though, the results from S1/S5 mismatch trials are important. There were over 30% incorrect "same" responses on same view S1/S5 mismatch trials, so performance was far from at ceiling, yet here participants also made over 30% incorrect "same" responses on different-view trials.

2.3 *Discussion*

The results from experiment 1 indicate that in an unspeeded simultaneous picture-matching task, participants could use view-insensitive information to discriminate between morphs with dissimilar shapes. On S1/S5, S1/S7, and S1/S13 mismatch trials, performance was no better when both morphs were depicted from the same view than when they were depicted from different views rotated by 90° in depth. This view-insensitive performance was not simply a result of performance being at ceiling, as the percentage of incorrect "same" responses reduced from S1/S5 to S1/S7 to S1/S13 mismatches (see figure 3).

In contrast, on trials presenting morphs with identical shapes (S1/S1 match trials) or similar shapes (S1/S2 and S1/S3 mismatch trials), participants were highly sensitive to the views of the morphs presented. When both morphs were depicted from the same view, participants were much more likely to correctly (on S1/S1 trials) or incorrectly (on S1/S2 and S1/S3 trials) respond that both pictures showed the same morph, relative to when the morphs were depicted from different views, see figure 3. Thus, when pairs of similarly shaped morphs were depicted from the same rather than different views, they were more likely to be categorised as having the same shape.

Overall, in experiment 1 sensitivity to a 90° view change increased as the shape change between two objects decreased. This result confirmed our prediction which was based on studies that showed that view sensitivity to plane misorientation increased when objects which were more visually similar had to be discriminated (Lawson and Jolicoeur 1998, 2003; Murray 1998). These results were not an artifact of the ease of recognition of the particular views selected. A clear and significant same-view advantage on S1/S1, S1/S2, and S1/S3 trials (and view-insensitive performance on S1/S5, S1/S7, and S1/S13 trials) was found for both 30°/30° and 120°/120° same-view trials compared to 30°/120° and 120°/30° different-view trials. The results also generalised across two different sets of morph objects tested with separate groups of participants and the interaction was observed for most of the individual objects tested (see figure 1). Thus, in contrast to the results of Cutzu and Edelman (1998), the present results provide direct evidence that the representation of similar objects is view sensitive and that view sensitivity increases as the difficulty of shape discrimination increases.

3 Experiment 2

In experiment 1, in a simultaneous matching task, we found a strong interaction between the effects of view change and shape change. In experiment 2, we attempted to replicate and to extend this result. As in experiment 1, we varied the size of the shape change. We also varied the ease of achieving object constancy across different view changes for 0°, 30°, and 90° depth rotations. This manipulation tested whether view sensitivity occurred only when there were large (90°) depth rotations between the two pictures, as in experiment 1, or whether view sensitivity would extend to relatively small (30°) depth rotations. Finally, we compared view-change and shape-change effects on simultaneous matching (as in experiment 1) and on sequential matching.

In a sequential picture-matching task, participants must store an internal representation of the first picture shown in order to decide whether that picture matches the second picture. The internal representation of the first picture may not be as accurate or as detailed as the representations available in a simultaneous matching task. The memory load requirement is particularly low in an unspeeded simultaneous matching task such as that used in experiment 1: here the participant could look back and forth repeatedly to compare across the two pictures and focus on the most likely areas of difference between the stimuli. The task requirements of sequential matching are closer to those of everyday object recognition (in which an input stimulus must be compared to stored memories of previously seen objects) than the requirements of simultaneous matching (for which problem-solving strategies could be used which would not typically be available in everyday viewing situations). In experiment 2, half the participants did the simultaneous matching task and half did the sequential matching task. In the simultaneous matching task, as in experiment 1, both pictures were presented for 1500 ms. In the sequential matching task, the left and right pictures were each presented individually for 1500 ms.

In experiment 1, participants were not given feedback whether their response was correct. Performance on S1/S7 and S1/S13 mismatch trials was relatively good, indicating that participants understood the requirements of the task. Nevertheless, participants may not have realised how poorly they were performing on S1/S2 and S1/S3 trials. To try to inform participants and to motivate them to improve their performance in experiment 2 we provided them with feedback throughout the study. In addition, we did not test mismatches with the most visually similar pairs of morphs, S1 and S2 (instead, we tested S10 visually dissimilar mismatches). The results of experiment 1 suggested that participants could rarely discriminate between the S1 and S2 morphs and we did not wish to confuse participants by providing feedback that different morphs had been presented when the stimuli depicted were virtually indistinguishable to them. Also, in experiment 1 there was only one match (same morph) trial to every five mismatch (different morph) trials. Since participants were receiving feedback in experiment 2, they would be sensitive to the proportion of match and mismatch trials and they may have been biased to respond "different", if most trials were mismatch trials. To avoid this, in experiment 2 we increased the proportion of match trials to equal that of mismatch trials.

In experiment 2 we examined the nature of the relationship between shape changes and view changes. First, given the results of experiment 1, we predicted that view sensitivity would increase as the difficulty of shape discrimination increased in both simultaneous and sequential matching tasks. Second, if sequential matching taps more stable, abstract representations than simultaneous matching and, if such representations are relatively view insensitive (as suggested by the results of Cutzu and Edelman 1998), then reduced view sensitivity was predicted for sequential matching relative to simultaneous matching.

3.1 *Method*

3.1.1 *Participants.* Twenty participants took part in the experiment for course credit.

3.1.2 *Materials.* Ten experimental objects were used: objects 1, 5, 6, 7, 9, 11, 12, 22, 25, and 29 (see figure 1). Each object was represented by six morphed versions: S1, S3, S5, S7, S10, and S13. Each morph was, in turn, depicted by three views in depth, at 30°, 60°, and 120°. Thus 18 pictures of each object (six morphs by three views) were presented.

3.1.3 *Design.* Ten participants were randomly allocated to the simultaneous condition and ten to the sequential condition. On each trial, two pictures of an object were presented, one on the left and one on the right side of the screen. The pictures were presented either simultaneously, or sequentially with the left picture presented first. All participants completed 300 experimental trials, of which half were match and half were mismatch trials. These comprised 15 match and 15 mismatch trials for each of the ten experimental objects. The left picture on a trial was always the 30° view of S1. On match trials, the right picture also showed S1, which was depicted at 30°, 60°, and 120° views on five trials for each object. On mismatch trials, the right picture showed S3, S5, S7, S10, or S13, each of which was depicted once at 30°, once at 60°, and once at 120° views for each object. In both the practice and the experimental blocks, the trials were presented in a different, random order for each participant.

3.1.4 *Apparatus and procedure.* This was similar to those in experiment 1 except for the following details. In the simultaneous matching task, stimulus presentation was as in experiment 1. In the sequential matching task, the left picture appeared for 1500 ms and was immediately replaced by the right picture which was also presented for 1500 ms, with the pictures appearing in the same position as in the simultaneous matching task. In both tasks, after the pictures had been presented, participants saw a written prompt to respond which reminded them of the 'm' and 'z' keypress responses for same shape (match) and different shape (mismatch) trials, respectively. After making their response, participants were provided with feedback for 600 ms which indicated the correct response for that trial. There was an intertrial interval of 500 ms.

Prior to starting the experimental block, all participants completed a block of 14 practice trials. The practice trials were identical to the experimental trials except that they presented morphed versions of seven objects which were not used in the experimental block, pictures were presented for 2500 ms, the feedback was presented for 2000 ms, and the intertrial interval was 1000 ms.

3.2 *Results*

ANOVAs (see table 1) were conducted separately on the mean percentage of "same" responses on match and mismatch trials since, unlike experiment 1, there were five times more trials at each view for match than for mismatch trials. On match trials, there was one within-participants factor—view change (the difference in view between the left and the right picture, 0°, 30°, or 90°)—and one between-participants factor—presentation (simultaneous or sequential presentation of the left and right pictures). On mismatch trials, there were the same view change and presentation factors plus an additional within-participants factor—shape change (if the right picture showed the S3, S5, S7, S10, or S13 morph, the left picture always showed the S1 morph).

Table 1. ANOVA of the mean percentage of "same" responses (%E) in experiment 2.

	Degrees of freedom	$F^{\mathrm{p}}/\%\mathrm{E}$	Degrees of freedom	$F^{\mathrm{i}}/\%\mathrm{E}$
Same-shape match trials				
view change	2, 36	82.585***	2, 18	15.417***
presentation	1, 18	0.897	1, 9	0.940
view change × presentation	2, 36	0.921	2, 18	0.671
Different-shape mismatch trials				
shape change	4, 72	211.373***	4, 26	39.856***
view change	2, 36	2.755	2, 18	0.314
presentation	1, 18	3.330	1, 9	2.118
shape change × view change	8, 144	2.475*	8, 72	1.527

Note: $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

3.2.1 *Same-shape match trials.* There were over 20% more correct "same" responses on same-view trials (95.2%) than on 30° or 90° different-view trials (73.5% and 72.4%, respectively). Performance was similar for simultaneous (81.8%) and sequential (78.9%) picture presentation. Finally, the increase in accuracy on same-view relative to different-view trials was similar for simultaneous (21.0%) and sequential (23.5%) picture presentation.

3.2.2 *Different-shape mismatch trials.* Errors reduced as the shape change increased between the two pictures presented on a trial (see figure 4), but same-view trials did not produce more incorrect "same" responses overall. As for match trials, there was a trend for better performance for simultaneous matching (23.7% incorrect "same" responses) than for sequential matching (26.7% incorrect "same" responses). The only significant interaction was that of shape change × view change (see figure 4). For S3 morphs only, there was a view-change effect in the direction predicted. Here, around 11% more incorrect "same" responses were made on same-view trials (64.0% "same" responses) than on 30° (49.5%), or 90° (56.0%) different-view trials.

3.3 *Discussion*

The results of experiment 2 replicated and extended those of experiment 1. First, for both simultaneous and sequential picture matching, performance was view insensitive when two visually dissimilar shapes (such as S1/S5, S1/S7, S1/S10, or S1/S13) were presented on mismatch trials. Performance was no better when both morphs were shown

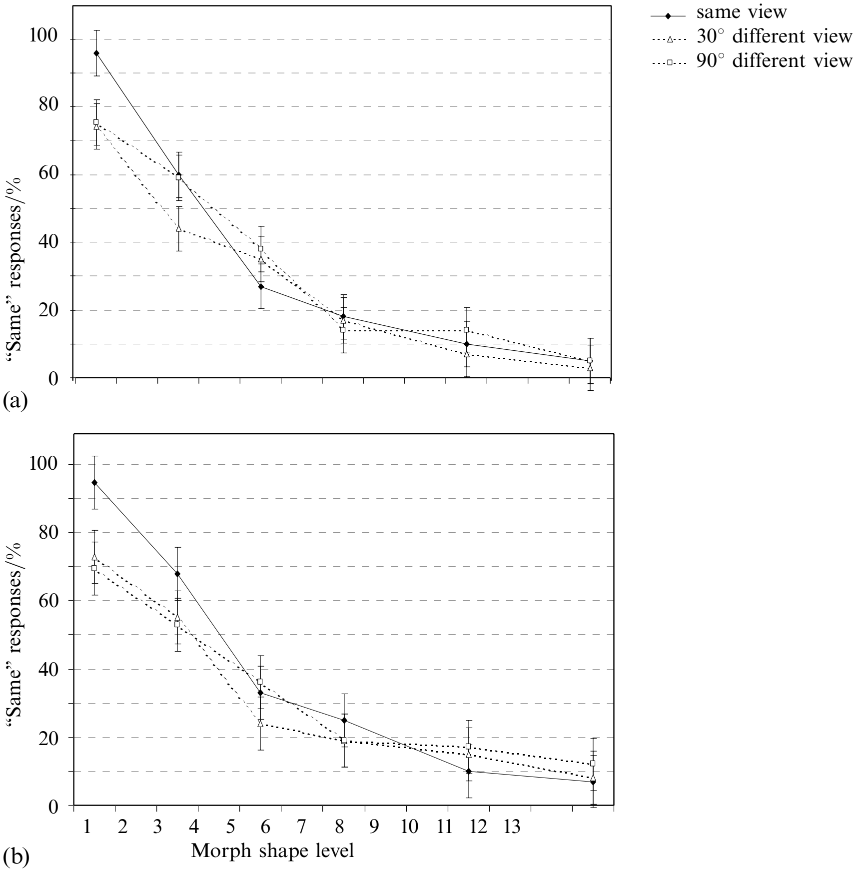

Figure 4. Results from experiment 2: mean percentage of “same” responses on same-view trials (0° view change) and different-view trials (30° or 90° view change) when a picture of S1 was seen with another picture of S1 on match trials (where “same” shape was the correct response) or with S3, S5, S7, S10, or S13 on mismatch trials (where “same” shape was the wrong response) during (a) simultaneous picture matching, and (b) sequential picture matching, along with 95% confidence intervals based on the error term for participants for the main effect of view change for match trials and on the interaction of shape change × view change for mismatch trials (Loftus and Masson 1994).

from the same view rather than from different views, although there was a strong effect of shape change with “same” responses systematically reducing across S1/S5, S1/S7, S1/S10, and S1/S13 trials. In contrast, on S1/S1 match trials, performance was highly view sensitive. When both pictures depicted S1 from the same view, there were over 20% more “same” responses than when different views of S1 were shown. Also, when two visually similar shapes (S1 and S3) were presented on mismatch trials, more “same” responses were made on same-view than on different-view trials. This view-sensitive effect mirrored, but was weaker than, that for S1/S1 match trials. The 30° and 90° different-view trials produced similar results. A physically small view change (of just 30°) disrupted shape discrimination as much as a large (90°) view change, so the view-change effects in experiment 1 were not caused by the large rotations in depth used. This final finding is consistent with the results of Foster and Gilson (2002) mentioned earlier. Overall, as in experiment 1, experiment 2 revealed a systematic decrease in the influence of view on matching from highly view-sensitive performance on S1/S1 match trials to moderate view sensitivity on S1/S3 mismatch trials to view invariance on S1/S5, S1/S7, S1/S10, and S1/S13 mismatch trials.

Second, view sensitivity was not reduced for sequential compared to simultaneous matching. Instead, there was a similar increase in "same" responses on same-view compared to different-view trials for both sequential and simultaneous matching. It therefore cannot merely be transient, low-level representations that are view sensitive. View-sensitive information must also be retained in the more stable, abstract visual representations required to mediate sequential matching performance. This issue was tested further in experiment 3 by extending the interstimulus interval (ISI) between the two pictures presented on each sequential matching trial from 0 ms (as in experiment 2 here) up to 400 ms, 1200 ms, and 3600 ms.

In experiment 2, performance on mismatch trials was largely view insensitive. However, in experiments 1 and 2, we may have underestimated the difficulty in achieving object constancy over depth rotation when visually similar shapes have to be discriminated because an unspeeded response measure was used. Unspeeded tasks allow participants to use slow, problem-solving, view-insensitive response strategies. In experiments 3 and 4 we investigated sequential picture matching with speeded response tasks. Response times may provide a more sensitive measure of the interaction between view changes and shape changes.

4 Experiment 3

In experiments 3 and 4 we investigated whether the interaction between shape-change and view-change effects, found in the unspeeded tasks of experiments 1 and 2, could be replicated in speeded response tasks. In experiment 2, the size of the view change on different-view trials (whether 30° or 90°) did not influence performance, so in experiments 3 and 4 only one view change (150°) was tested. The views tested on 90° different-view trials in experiment 2 (120° views) and those tested on 150° different-view trials in experiment 3 (240° views) were usually nearly mirror images of each other (see figures 2b and 2c). Apart from a reflection, similar visual information would be available in both views, and little or no difference between picture-matching of identical and of mirror-image views has been found in previous research (eg Lawson and Humphreys 1996). In experiments 3 and 4, only S1/S1 match and S1/S7 mismatch trials were tested. The S1 and S7 morphs had sufficiently dissimilar shapes to have produced view-insensitive performance on mismatch trials in both experiments 1 and 2.

In experiment 2, the simultaneous and sequential matching conditions produced very similar results. It is difficult, though, to make direct comparisons between sequential and simultaneous picture presentations because it is not clear how to equate the presentation duration of stimuli across the two conditions. To avoid this problem, in experiments 3 and 4 only sequential picture matching was tested and ISI was varied from 400 ms to 1200 ms to 3600 ms (in contrast to the 0 ms ISI used in sequential matching in experiment 2). Varying ISI allowed us to investigate the effect of increasing the duration over which the internal representation of the first picture had to be maintained whilst fixing the presentation duration of the first picture at 500 ms at all ISIs. The second picture was always presented until participants responded.

At longer ISIs, more information about the stimulus would be expected to be lost, so overall performance should be adversely affected. Furthermore, picture-matching studies presenting familiar objects have shown reduced view sensitivity at longer ISIs (Ellis and Allport 1986; Lawson and Humphreys 1996). As ISI lengthens, participants may increasingly rely on a more abstract description of the first object shown, for instance by activating stored, view-insensitive visual representations of the object, or by accessing a semantic or verbal representation of the object. This increasing reliance on more abstract information at longer ISIs may occur for novel as well as for familiar objects, particularly if participants are familiarised with the stimuli before the picture-matching task begins. Familiarisation may allow participants to develop and store visual, semantic, and verbal representations which are relatively view insensitive and abstract.

In experiment 3, the effects of familiarisation were tested by pre-exposing participants to the S1 stimuli. Before starting the picture-matching task, participants saw the 30° view of S1 for all the experimental objects in each of three blocks of training. In the first training block, participants were asked to rate the accuracy of a verbal description accompanying a picture of the S1 morph. In the second and third training blocks, participants had to try to recall this description whilst looking at the picture of the S1 morph, and then they had to rate the accuracy of their recall. These tasks were designed to make participants study the pictures and their descriptions carefully. In particular, we encouraged participants to attend to the 3-D structure of the objects, by using verbal descriptions that emphasised the shape and relative position of the parts of the objects.

Recently, Walker et al (2000) argued that pre-exposing participants to a novel name of a novel object ("This is a dax") encourages "attention to be directed globally, at the overall configuration of a figure, ... because this allows the types of representation more supportive of object categorisation to be encoded" (page 621). These representations were assumed to be global, 3-D-shape descriptions which are abstract and view insensitive, with privileged links to object names. They were contrasted to representations such as those derived from structurally impossible objects or derived under encoding conditions which do not emphasise 3-D object shape as an important feature. Landau et al (1998) have also suggested that providing a name and describing a function for a visually presented novel object are both important factors in determining how adults generalise from a given object to new shapes. Together, this research suggests that providing names and functions for novel objects seen during an initial training phase can change participants' subsequent ability to achieve visual object constancy for those objects by allowing view-insensitive visual, semantic, or verbal representations to be encoded. In contrast, Williams and Simons (2000) found no effect of learning to associate novel names with novel objects for participants' subsequent sensitivity at detecting shape changes to the novel objects in a sequential matching task (though participants were faster if they had learnt names for the objects).

In experiment 3, different verbal descriptions were shown to three different groups of participants:

(1) The *real group* read descriptions that explicitly named the original object from which a given morph was derived and that linked its function to specific parts of the object. Descriptions started: "This is a (machine gun) ...", see appendix.

(2) The *like group* read descriptions that named a novel object which the morph might be and that linked its possible function to specific parts of the object. Descriptions started: "This could be a (technologically advanced building machine) ...".

(3) The *part group* read descriptions that made no reference to the name of an object or to a possible function of the object. Only the parts of the object and their spatial arrangement were described. Descriptions started: "This object (has a long, complex, central horizontal section) ...".

If participants could develop abstract, view-insensitive representations of S1 morphs, then the real group were predicted to be the most likely to do so. Here, the description of S1 could usually be related directly to pre-existing, abstract, visual, semantic, and verbal representations of the familiar object which was specified in the description and from which the S1 morph had originally been derived. If such abstract representations were useful, then subsequent matching performance by the real group should be good overall and also relatively insensitive to view changes (see Laudau et al 1998; Walker et al 2000), particularly at longer ISIs. The part group was predicted to be least likely to develop abstract representations, since the verbal descriptions they were given did not include an object name or any semantic or functional information about the S1 morph. Indeed, the part group might instead develop a highly view-specific representation of S1, since they saw S1 repeatedly from the same 30° view and they were not encouraged to develop abstract, view-insensitive representations. If so, then the part group might show even greater view sensitivity, than if they had had no training. Finally, the like group were given an object name, and semantic and functional information which could be linked to the S1 morph, but this information could not usually be mapped to pre-existing, abstract representations. However, these participants may have developed view-insensitive, abstract representations during training. The like group was therefore predicted to reveal intermediate view sensitivity relative to the real and the part groups.

In summary, in experiment 3 we investigated whether, first, the view change by shape-change interaction found for unspeeded tasks in experiments 1 and 2 would be replicated in a speeded task; second, whether any view sensitivity would be reduced as ISI increased; and third, whether any view sensitivity would be reduced if participants were encouraged to encode more abstract, view-insensitive representations by being familiarised with real and like (relative to part) descriptions prior to starting the matching task.

4.1 *Method*

4.1.1 *Participants.* Thirty-six participants took part in the study for course credit.

4.1.2 *Materials.* Eighteen experimental objects were used: objects 1, 4, 6, 7, 8, 9, 10, 13, 14, 15, 16, 17, 18, 20, 21, 25, 26, and 29 (see figure 1). Each object was represented by two morphed versions, S1 and S7. Each morph was, in turn, depicted by two views in depth, at 30° and 240°. Thus four pictures of each object (two morphs by two views) were presented. Real, like, and part verbal descriptions of each of the eighteen novel objects were also produced, see appendix.

4.1.3 *Design.* Twelve participants were randomly allocated to each of the real, like, and part groups. All participants completed three training blocks of 18 rating trials and one test block of 216 picture-matching trials. The 30° view of the S1 morph of each of the eighteen experimental objects was shown once in each training block. Each participant was presented with the same verbal description of a given object each time that he/she saw that object during training. In the test block, each of the eighteen objects was presented on 6 match trials and 6 mismatch trials. On all trials, the first picture presented was the 30° view of S1 which had been seen during training. The second picture was S1 on match trials and S7 on mismatch trials. On both match and mismatch trials, relative to the first picture, the second picture could depict the object from the same, 30° view (on 3 trials in the test block for each object, with ISIs of 400 ms, 1200 ms, and 3600 ms) or from a different, 240° view (again, on 3 trials for each object, with ISIs of 400 ms, 1200 ms, and 3600 ms). Prior to starting the experimental matching block, all participants completed a block of 20 practice trials. These practice trials were selected at random from the set of 216 trials used in the experimental natching block. In all blocks, trials were presented in a different, random order for each participant.

4.1.4 *Apparatus and procedure.* The experiment was run on a Macintosh Power PC G4 computer with the Psyscope version 1.2.5 experimental presentation software. In the first training block, participants first read a description of an object, then pressed the space bar. The 30° view of the S1 morph of the object that had been described was then presented. Participants made an unspeeded rating of the accuracy and informativeness of the object description, from 1 (low) to 9 (high). In the second and third training blocks, participants first saw the 30° view of the S1 morph of a given object.

Whilst looking at the picture, they tried to remember the description of that object that they had read in the initial training block. They then pressed the space bar and the description that they had seen previously was presented. They reread this description then made an unspeeded self-rating of their recall accuracy, from 1 (low) to 9 (high).

On each trial in the training blocks, a central fixation point appeared for 300 ms and was then replaced by the object description in the first training block, and the object picture in the second and third training blocks. When participants pressed the space bar, both the description and the picture appeared together. The picture was always presented at fixation and the description appeared centrally and above fixation. Both the description and the picture disappeared once the participant made a keypress to select a rating. There was an intertrial interval of 500 ms. The descriptions and pictures were presented for an unlimited time and participants were not told to respond rapidly.

In the test block of picture-matching trials, participants made a speeded decision whether two successive pictures showed the same or different objects. Participants made 'm' or 'z' keypresses to respond "same" or "different", respectively. On each trial, a central fixation point appeared for 750 ms, then after 500 ms the first picture was presented for 500 ms. After a blank ISI of 400 ms, 1200 ms, or 3600 ms, the second picture was presented until the participant responded. The first and second pictures were both presented at fixation. After a response was made, the correct response on that trial was given as feedback for 500 ms, by presenting the letter "m" or "z" at fixation. There was an intertrial interval of 750 ms. Participants were told to ignore any difference in the view depicted in the first and second pictures. They were warned that on mismatch trials, the pictures might depict two objects with very similar shapes, and so the task was extremely difficult. Participants were encouraged to respond as fast and as accurately as possible.

4.2 *Results*

Response latencies in picture-matching trials less than 300 ms or exceeding 2300 ms were discarded as errors (less than 2% of trials). No participants were replaced. There was one missing cell in the by-items analyses which was replaced by the mean for that condition. ANOVAs were conducted on the mean correct matching RTs and on the percentage of "same" responses for match and mismatch trials separately (see table 2). There were two within-participants factors—view change (same or different)—and ISI (400 ms, 1200 ms, or 3600 ms) and one between-participants factor—condition (like, part, or real training descriptions).

Table 2. ANOVA of the mean correct matching RTs and percentage of "same" responses (%E) in experiment 3.

	Degrees of freedom	F^{p}/RTs	F^{p}/%E	Degrees of freedom	F^{i}/RTs	F^{i}/%E
Same-shape S1/S1 match trials						
condition	2, 33	1.503	0.029	2, 34	22.071***	0.087
view change	1, 33	237.904***	280.882***	1, 17	200.064***	57.547***
ISI	2, 66	19.961***	6.858**	2, 34	20.570***	4.518*
view change × ISI	2, 66	7.759***	0.487	2, 34	8.225**	0.397
Different-shape S1/S7 mismatch trials						
condition	2, 33	2.279	0.113	2, 34	30.621***	0.199
view change	1, 33	31.994***	33.438***	1, 17	14.265**	11.532**
ISI	2, 66	7.661**	7.266**	2, 34	5.292*	7.062**
view change × ISI	2, 66	0.273	9.125***	2, 34	1.326	3.868*

Note: $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

4.2.1 *Same-shape S1/S1 match trials.* Same-view matches (756 ms, 93.2% correct "same" responses) were 322 ms faster and 35.7% more accurate than different-view matches (1078 ms, 57.5%). As the ISI increased from 400 ms to 1200 ms to 3600 ms, RTs slowed from 869 ms to 922 ms to 961 ms and correct "same" responses decreased from 76.9% to 76.5% to 72.4%, respectively. The like group was slower (983 ms) but no less accurate (75.6% correct "same" responses) than the real group (891 ms, 74.9%) and the part group (877 ms, 75.4%). The interaction of view change × ISI is shown in figure 5. As the ISI increased from 400 ms to 1200 ms to 3600 ms, the benefit for same-view over different-view matches reduced but was far from eliminated for RTs (385 ms to 322 ms to 260 ms), whilst the same-view benefit for errors remained large (34.4% to 36.7% to 35.9%).

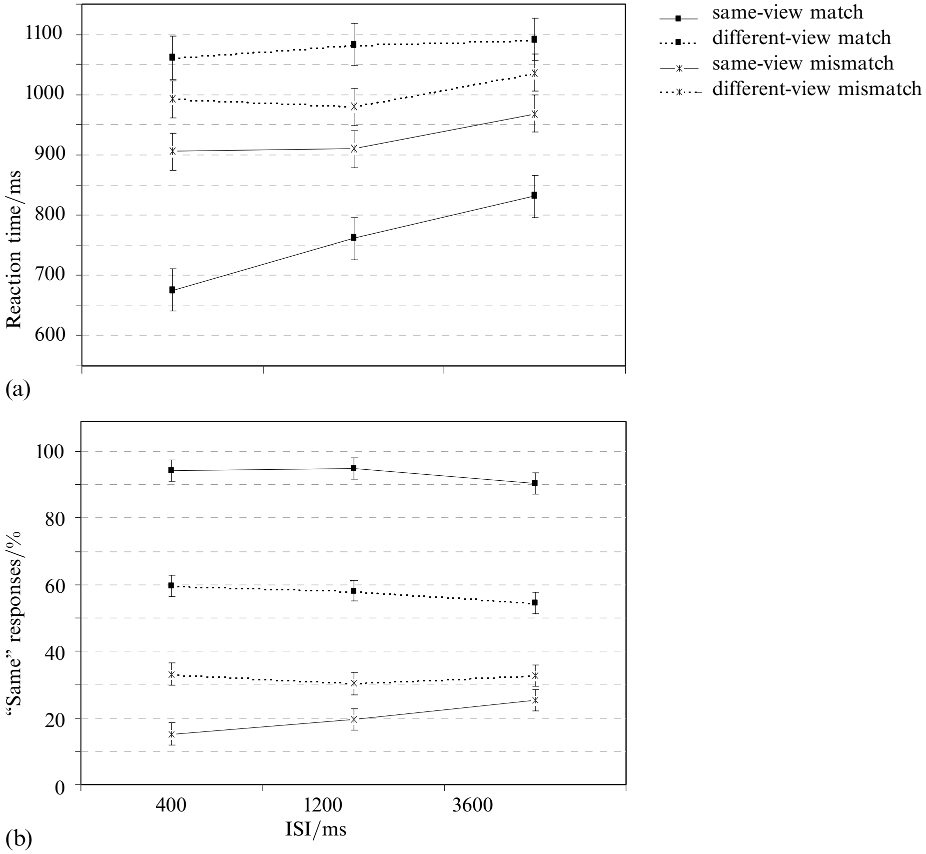

Figure 5. Results from experiment 3: (a) mean correct RTs and (b) mean percentage of "same" responses on same-view trials (0° view change) and different-view trials (150° view change) when a picture of S1 was seen with another picture of S1 on match trials (where "same" shape was the correct response) or with S7 on mismatch trials (where "same" shape was the wrong response) during sequential picture matching, with an ISI of 400 ms, 1200 ms, or 3600 ms, along with 95% confidence intervals based on the error term for participants for the interaction of view change × ISI for match and mismatch trials separately (Loftus and Masson 1994).

4.2.2 *Different-shape S1/S7 mismatch trials.* Same-view matches (928 ms, 20.0% incorrect "same" responses) were 75 ms faster and 12.1% more accurate than different-view matches (1003 ms, 32.1%). As the ISI increased from 400 ms and 1200 ms to 3600 ms, RTs increased from 949 ms and 945 ms to 1002 ms and incorrect "same" responses increased from 24.2% and 25.0% to 29.0%, respectively. As for match trials, the like group was somewhat slower (1044 ms) but no less accurate (25.3% incorrect "same" responses) than the real group (942 ms, 26.2%) or the part group (911 ms, 26.6%). The interaction of view change × ISI is shown in figure 5. As the ISI increased from 400 ms to 1200 ms to 3600 ms, the benefit for same-view over different-view mismatches was relatively small and constant for RTs (87 ms to 70 ms to 68 ms) whilst the same-view benefit reduced but was far from eliminated for errors (18.1% to 10.8% to 7.4%).

4.3 *Discussion*

In experiment 3, participants were highly sensitive to view changes. Replicating and extending the results from the unspeeded tasks used in experiments 1 and 2, performance in S1/S1 match trials was both faster and more accurate on same-view than on different-view trials. Performance on S1/S7 mismatch trials was also highly view-sensitive, with performance again being faster and more accurate on same-view than on different-view trials. This contrasts with the results for mismatch trials in experiments 1 and 2. Here, S1/S7 mismatches were view insensitive, and S1/S2 and S1/S3 mismatches were more accurate on different-view than on same-view trials. This latter difference is considered further in section 6.

At longer ISIs, overall responses were somewhat slower and less accurate. ISI also interacted with view change, such that the same-view benefit reduced somewhat at longer ISIs (see figure 5). This interaction replicates those reported by Ellis and Allport (1986), and by Lawson and Humphreys (1996) in picture–picture matching studies presenting familiar objects. Note, though, that in all conditions of experiment 3, performance remained highly view sensitive, especially on match trials. These results suggest that somewhat more abstract, less view-sensitive representations mediated performance at longer ISIs, although performance on both match and mismatch trials was still clearly view sensitive even with a 3600 ms ISI.

The real, like, and part groups tested in experiment 3 produced very similar results. On match trials, the same-view advantage for the real, like, and part groups was 340 ms, 313 ms, and 312 ms for RTs; and 35%, 36%, and 36% for errors respectively. On mismatch trials, the same-view advantage was 91 ms, 61 ms, and 72 ms for RTs; and 8%, 13%, and 15% on errors. There was no evidence that training to link pictures of S1 morphs to verbal descriptions giving names and functions for the morphs helped the real and like groups to develop abstract view-insensitive descriptions of the morphs relative to the part group. The prediction that view sensitivity would be reduced for the real and like groups compared to the part group was not supported. This result is not consistent with the claims of Landau et al (1998) and Walker et al (2000) that learning names and functions for novel objects plays an important role in achieving object constancy. Instead, our results replicate Williams and Simons (2000) in finding no effect of learning names for novel objects on participants' subsequent sensitivity to detecting shape changes to those novel objects in a sequential matching task.

It is possible that the training in experiment 3 did influence subsequent matching performance, but that differences in the content of the verbal descriptions provided to the three groups were irrelevant. For example, the only important component of the descriptions may have been the parts-based structural description of the object which was provided in all of the real, like, and part descriptions. All of these descriptions may have encouraged participants to activate abstract, view-insensitive representations of the morphs during training. In this case, reduced view sensitivity would be predicted following training for the real, like, and part groups compared to following no training. Alternatively, the content of the verbal descriptions may not have had any effect, but repeatedly seeing the 30° view of S1 during training may have encouraged participants to store a representation of that particular view. If so, then, relative to no training, greater view sensitivity would be predicted following training which presented the 30° view of S1, with faster and more accurate responses on 30° same-view trials.

These two opposing predictions were tested by comparing performance following training in experiment 3 to performance in experiment 4, where participants received no training prior to the picture-matching task.

5 Experiment 4

In experiment 4 we replicated experiment 3 except that participants were not exposed to 30° views of S1 or to verbal descriptions of the morphs in an initial training phase. The first block of picture matching in experiment 4 was identical to the block of picture matching in experiment 3. Comparing performance here to performance in experiment 3 should indicate whether there was any influence of the training phase on subsequent matching performance.

In experiment 4, we also examined the effect of practice by giving a second block of picture-matching trials to participants, which was identical to the first block. After experience with the stimuli during the first block of picture matching, view sensitivity may be reduced if participants can learn to encode and use less view-sensitive representations of the objects. Practice at the matching task may be more effective in reducing view sensitivity than the training blocks of pictures and verbal descriptions given in experiment 3, since both 30° and 240° views were presented in the first block of picture-matching trials in experiment 4, whereas only 30° views were shown in the description training in experiment 3. A reduction in view-change effects on picture matching in the second relative to the first block of trials was reported by Lawson and Humphreys (1996). A similar influence of prior experience was predicted here.

5.1 *Method*

Twelve participants took part in the study for course credit. Experiment 4 was identical to experiment 3 except that there were no training blocks and the test block of 216 matching trials was given to participants twice, with a self-timed break before the second experimental block.

5.2 *Results*

Response latencies in picture-matching trials less than 300 ms or exceeding 2300 ms were discarded as errors (less than 2% of trials). No participants were replaced. ANOVAs were conducted on the mean correct matching RTs and on the percentage of "same" responses for match and mismatch trials separately (see table 3). There were three

Table 3. ANOVA of the mean correct matching RTs and percentage of "same" responses (%E) in experiment 4.

	Degrees of freedom	F^{p}/RTs	F^{p}/%E	Degrees of freedom	F^{i}/RTs	F^{i}/%E
Same-shape S1/S1 match trials						
view change	1, 11	96.103***	88.049***	1, 17	162.947***	45.656***
ISI	2, 22	6.067**	8.950**	2, 34	16.984***	17.496***
view change × ISI	2, 22	6.450**	0.227	2, 34	5.922**	0.123
block	1, 11	18.003**	0.415	1, 17	58.158***	0.391
view change × block	1, 11	17.708**	3.443	1, 17	15.654**	3.368
Different-shape S1/S7 mismatch trials						
view change	1, 11	13.441**	9.878**	1, 17	8.525**	11.248**
ISI	2, 22	7.501**	2.595	2, 34	3.898*	3.019
view change × ISI	2, 22	2.942	3.472*	2, 34	2.756	2.891
block	1, 11	9.158*	31.201***	1, 17	10.398**	12.693**
view change × block	1, 11	0.242	2.523	1, 17	0.132	0.973

Note: $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

within-participants factors: view change (same or different), ISI (400 ms, 1200 ms, or 3600 ms), and block (1 or 2).

5.2.1 *Same-shape S1/S1 match trials.* Same-view matches (680 ms, 94.3% correct "same" responses) were 273 ms faster and 37.5% more accurate than different-view matches (953 ms, 56.8%). As the ISI increased from 400 ms to 1200 ms to 3600 ms, RTs slowed from 775 ms to 819 ms to 856 ms and correct "same" responses decreased from 78.0% to 77.0% to 71.6%, respectively. Responses in block 2 (772 ms, 76.0% correct "same" responses) were 89 ms faster but no more accurate than in block 1 (861 ms, 75.1%). The interaction of view change × ISI is shown in figure 6. As the ISI increased from 400 ms to 1200 ms to 3600 ms, the benefit for same-view over different-view matches reduced but was far from eliminated for RTs (324 ms to 287 ms to 205 ms), whilst the same-view benefit for errors remained large (38.4% to 36.8% to 37.3%). This interaction fully replicated the results of experiment 3. For the interaction of view change × block, the advantage for same-view over different-view matches was rather larger in block 1 (328 ms, 40.0%) than block 2 (217 ms, 35.0%).

5.2.2 *Different-shape S1/S7 mismatch trials.* Same-view matches (818 ms, 18.0% incorrect "same" responses) were 66 ms faster and 11.4% more accurate than different-view matches (884 ms, 29.4%). As the ISI increased from 400 ms to 1200 ms to 3600 ms, RTs increased from 824 ms to 854 ms to 874 ms and incorrect "same" responses increased from 21.8% to 24.0% to 25.3%, respectively. Responses in block 2 (818 ms, 20.1% incorrect "same" responses) were 65 ms faster and 6.1% more accurate than in block 1 (883 ms, 27.2%). The interaction of view change × ISI is shown in figure 6. As the ISI increased from 400 ms to 1200 ms to 3600 ms, the benefit for same-view over different-view mismatches reduced but was not eliminated for RTs (101 ms to 65 ms to 32 ms) and errors (16.1% to 11.3% to 6.7%). This pattern of results is similar to that of experiment 3. In contrast to the match trials, the advantage for same-view over different-view matches was similar in block 1 (73 ms, 9.1%) and block 2 (60 ms, 13.7%).

5.2.3 *Comparison of results from experiment 3 and from block 1 of experiment 4.* The following analyses were conducted in order to directly compare results from participants who were and who were not pre-exposed to pictures of the experimental objects during an initial training phase. Here, data from experiment 3 (for the three groups of twelve participants who saw pictures and descriptions of the 30° view of S1 during training, prior to starting the picture-matching task) were analysed with the data from the first block of matching in experiment 4 (for the group of twelve participants who saw no pictures or descriptions prior to starting the picture-matching task). The analyses reported in experiment 3 were then repeated, except that an additional level was included in the between-participants factor of condition, to give a total of four levels for the four groups of participants tested. The main effect of condition on RTs was marginally significant in the by-participants analysis and significant in the by-items analysis. This was because the like group were around 100 ms slower than the other three groups on both match and mismatch trials. No other main effects or interactions involving the condition factor were significant for either RTs or errors, in either the by-participants or the by-items analyses. The results of these latter analyses mirrored those reported above. Most importantly, there was no evidence to suggest that view sensitivity varied across the like, part, real, and no-training groups. This suggests that the training received by participants in experiment 3 did not influence their subsequent picture-matching performance.

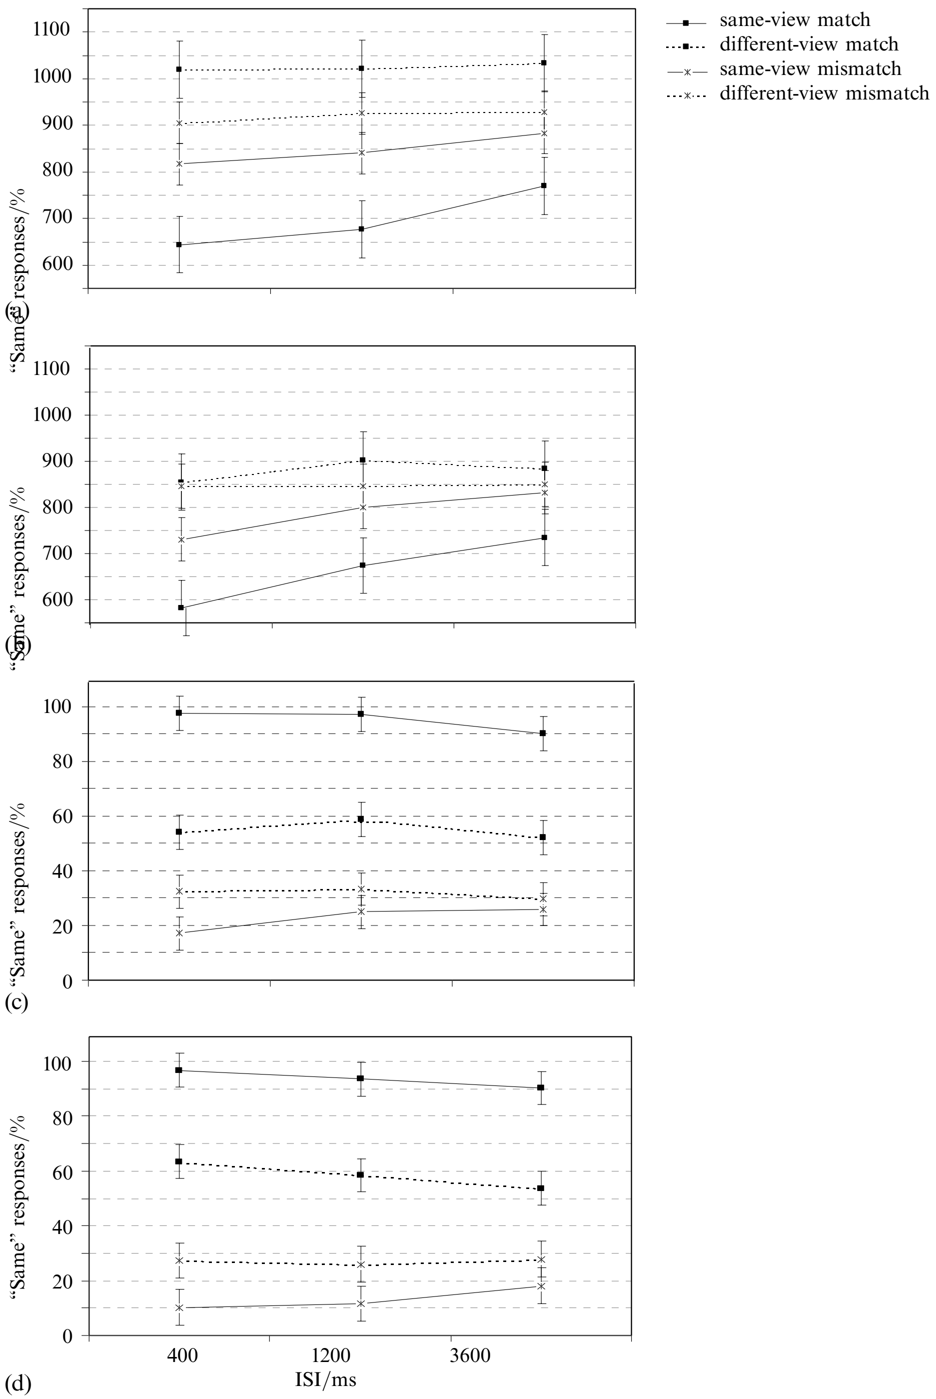

Figure 6. Results from experiment 4: mean correct RTs in (a) block 1 and (b) block 2, and mean percentage of "same" responses in (c) block 1 and (d) block 2, on same-view trials (0° view change) and different-view trials (150° view change) when a picture of S1 was seen with another picture of S1 on match trials (where "same" shape was the correct response) or with S7 on mismatch trials (where "same" shape was the wrong response) during sequential picture matching, with an ISI of 400 ms, 1200 ms, or 3600 ms, along with 95% confidence intervals based on the error term for participants for the interaction of view change × ISI for match and mismatch trials separately (Loftus and Masson 1994).

5.3 *Discussion*

The results of experiment 4 replicated and extended those of experiment 3. Same-view trials were much faster and more accurate than different-view trials, with the same-view advantage being greater on S1/S1 match trials than on S1/S7 mismatch trials. This view-change effect interacted with ISI with the same-view advantage reducing at longer ISIs (see figure 6). This replicated the interaction reported in experiment 3 (see figure 5), and is also consistent with the finding of reduced view-sensitivity at longer ISIs in picture-matching studies presenting familiar objects (eg Ellis and Allport 1986; Lawson and Humphreys 1996). Note, though, that as in experiment 3, performance in experiment 4 was stll highly view sensitive at the longest, 3600 ms, ISI. Performance on same-view trials deteriorated at longer ISIs unlike performance on different-view trials, suggesting that some view-specific information is lost over time whilst view-insensitive information appears more stable. Nevertheless, even at the longest ISI, sufficient view-specific information was available to provide a strong benefit for same-view over different-view trials. Thus the abstract representations mediating picture matching on 3600 ms ISI trials were still clearly view sensitive.

A comparison of performance in experiments 3 and 4 shows that presenting pictures and verbal descriptions that included semantic and functional information and names for the morphs during training in experiment 3 had little or no effect on subsequent matching performance. In experiment 4, participants received no training prior to the matching trials and yet their results were very similar to those of participants in experiment 3. In contrast to the lack of influence on view sensitivity of training in experiment 3, practice at the picture-matching task in experiment 4 reduced view sensitivity in block 2 relative to block 1, at least on S1/S1 match trials (see figure 6). The pattern of results in block 2 mirrored those of block 1. However, for S1/S1 match trials, while the same-view advantage in block 1 was 328 ms and 40.0% on errors, it was 217 ms and 35.0% on errors in block 2. This replicates the finding by Lawson and Humphreys (1996) of a reduced same-view advantage on match trials in the second, relative to the first, block of picture-matching trials. Note, though, that in experiment 4 practice only reduced (and was far from eliminating) view sensitivity. There was a clear same-view benefit in all conditions in block 2.

Together, the results from experiments 3 and 4 suggest that view-change effects are relatively impervious to participants' prior experience with the experimental stimuli and the task. Even at long ISIs, a robust same-view advantage remained after training with verbal descriptions and after practice at the matching task. Participants appeared unable to readily develop, store, and then use abstract, view-insensitive morph representations, which were as effective as their view-sensitive morph representations. Performance, therefore, remained highly view-sensitive across a range of experimental manipulations.

6 General discussion

In four picture-matching studies, we found that view sensitivity varied systematically with shape changes. The novel stimuli tested in our studies were derived from a relatively large and diverse set of familiar objects which varied in complexity, global shape, and the number and spatial arrangement of parts. Incrementally morphing these novel objects allowed the difficulty of shape discrimination to be manipulated systematically. On match trials, when both pictures on a trial depicted identical shapes, performance was always highly view-sensitive, with faster and more accurate responses when stimuli were presented from the same view rather than different views in depth. Performance was less view sensitive on mismatch trials, when pictures of two objects with visually similar shapes were depicted. Performance was view insensitive, with no difference between responses on same-view and different-view trials, when pictures of two objects with visually dissimilar shapes were presented. These results generalised across by-items as well as by-participants analyses. This interaction between the effects of shape changes and view changes on object recognition has often been assumed but has proved difficult to test rigorously.

View sensitivity was found reliably across a wide range of conditions: in both simultaneous matching (experiments 1 and 2) and sequential matching (for ISIs from 0 ms in experiment 2 up to 3600 ms in experiments 3 and 4); in speeded response tasks, over both RTs and errors (experiments 3 and 4), and in unspeeded response tasks (experiments 1 and 2); for different sets of objects (experiment 1); whether provided with feedback (experiments 2, 3, and 4) or not (experiment 1); for a range of view changes (30°, 90°, and 150°); with no prior training (experiments 1 and 2, and block 1 of experiment 4), after seeing pictures and verbal descriptions of the stimuli during training (experiment 3), and after practice at the picture-matching task (block 2 of experiment 4). In all of these cases, object constancy over depth rotation was difficult to achieve for visually similar shapes whereas performance with less visually similar shapes was significantly less view sensitive.

In particular, there was no evidence that participants were more reliant on abstract, view-insensitive representations in sequential (versus simultaneous) matching (experiment 2) or after pre-exposure to some of the stimuli (following training with verbal descriptions and pictures of the experimental morphs in experiment 3). View sensitivity did reduce significantly at longer (versus shorter) ISIs (experiments 3 and 4) and following practice at the task (in the second relative to the first block of trials in experiment 4). However, even in these cases there was a clear same-view advantage on matching in all conditions.

Any view sensitivity on match trials is always predicted to reveal better performance on same-view than on different-view trials, as was observed in all four studies reported here. In contrast, on mismatch trials two processes may be involved which would influence performance in different ways. First, on same-view relative to different-view mismatch trials, two pictures may be more likely to be mistakenly classified as showing the same object since the pictures appear more like each other (particularly for similarly shaped objects). Second, on same-view relative to different-view mismatch trials, two pictures may be less likely to be mistakenly classified as showing the same object because it is easier to compare across same-view pictures and so to detect changes. Interestingly, we found evidence that both of these processes occurred in the current studies. In experiments 1 and 2, the former process appeared to be more important. People made more errors on same-view than on different-view mismatch trials, particularly when similarly shaped objects were presented. This finding provides evidence for the hypothesis proposed by Gauthier and Tarr (1997) that the same view of two different but similarly shaped objects may be perceived as more similar than two different views of the same object (see figure 3). In contrast, in experiments 3 and 4 the latter process seemed to dominate. In all conditions, people made fewer errors (and were faster) on same-view than on different-view mismatch trials. Further experimentation will be required to understand how these two processes interact on mismatch trials across different tasks.

The aim of these four studies was to examine how we perceive and categorise pictures of complex, 3-D objects which are like real-world objects such as horses and cows but which, unlike most real-world objects, "grade continuously into each other" (Shepard and Cermak 1973, page 374). Our results indicate that view sensitivity is a robust, ubiquitous phenomenon which is primarily influenced by shape change. The same-view advantage in shape discrimination was both large and reliable across a range of testing conditions whenever visually similar objects had to be discriminated.

The results of experiments 1 and 2 indicated that the human visual object-recognition system can achieve view-insensitive performance, if sufficiently dissimilar objects are presented. Our results, though, emphasise the pervasive difficulty in achieving object constancy over depth rotation when similarly shaped objects must be discriminated. We suggest that such circumstances are typical under everyday viewing conditions since most objects that we can recognise could be confused with one or more similarly shaped familiar objects.

Acknowledgments. The first author gratefully acknowledges the support of the Leverhulme Trust and the Max-Planck-Institut für biologische Kybernetik. The first draft of this article was completed whilst the first author was on research leave at the Department of Psychology, the University of Otago, Dunedin, New Zealand, and she thanks Dr Janice Murray and her colleagues in the Department for their hospitality. We thank Rachael Fullam for testing participants and Isabelle Bülthoff and the members of the Visual Perception Lab at the University of Liverpool for helpful discussions.

References

Bartram D J, 1976 "Levels of coding in picture – picture comparison tasks" *Memory and Cognition* **4** 593 – 602

Beale J M, Keil C F, 1995 "Categorical effects in the perception of faces" *Cognition* **57** 217 – 239

Biederman I, Bar M, 1999 "One-shot viewpoint invariance in matching novel objects" *Vision Research* **39** 2885 – 2899

Campanella S, Chrysochoos A, Bruyer R, 2001 "Categorical perception of facial gender information: Behavioural evidence and the face – space metaphor" *Vision Cognition* **8** 237 – 262

Cutzu F, Edelman S, 1998 "Representation of object similarity in human vision: psychophysics and a computational model" *Vision Research* **38** 2229 – 2257

Edelman S, 1995 "Class similarity and viewpoint invariance in the recognition of 3D objects" *Biological Cybernetics* **72** 207 – 220

Edelman S, 1998 *Representation and Recognition in Vision* (Cambridge, MA: MIT Press)

Edelman S, Duvdevani-Bar S, 1997 "Similarity, connectionism, and the problem of representation in vision" *Neural Computation* **9** 701 – 721

Ellis R, Allport D A, 1986 "Multiple levels of representation for visual objects: A behavioral study", in *Artificial Intelligence and Its Applications* Eds A G Cohn, J R Thomas (Chichester, Sussex: John Wiley & Sons)

Ellis R, Allport D A, Humphreys G W, Collis J, 1989 "Varieties of object constancy" *Quarterly Journal of Experimental Psychology A* **41** 775 – 796

Foster D H, Gilson S J, 2002 "Recognizing novel three-dimensional objects by summing signals from parts and views" *Proceedings of the Royal Society of London, Series B* **269** 1939 – 1947

Gauthier I, Tarr M J, 1997 "Orientation priming of novel shapes in the context of viewpoint-dependent recognition" *Perception* **26** 51 – 73

Hayward W G, Williams P, 2000 "Viewpoint dependence and object discriminability" *Psychological Science* **11** 7 – 12

Landau B, Smith L, Jones S, 1998 "Object shape, object function and object name" *Journal of Memory and Language* **38** 1 – 27

Lawson R, 1999 "Achieving visual object constancy across plane rotation and depth rotation" *Acta Psychologica* **102** 221 – 245

Lawson R, Humphreys G W, 1996 "View-specificity in object processing: Evidence from picture matching" *Journal of Experimental Psychology; Human Perception and Performance* **22** 395 – 416

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

Loftus G R, Masson M E J, 1994 "Using confidence intervals in within-subjects designs" *Psychonomic Bulletin & Review* **1** 476 – 490

Murray J E, 1998 "Is entry-level recognition viewpoint-invariant or viewpoint-dependent?" *Psychonomic Bulletin & Review* **5** 300 – 304

Shelton C R, 1998 "Three-dimensional correspondence", AI Technical Report, AITR 1650, Artificial Intelligence Laboratory, MIT, 200 (545) Technology Square, MIT Building NE43, Cambridge, MA 02139, USA

Shepard R N, Cermak G W, 1973 "Perceptual-cognitive explorations of a toroidal set of free-form stimuli" *Cognitive Psychology* **4** 351 – 377

Tjan B S, Legge G E, 1998 "The viewpoint complexity of an object-recognition task" *Vision Research* **38** 2335 – 2350

Walker P, Dixon S, Smith D, 2000 "Associating object names with descriptions of shape that distinguish possible from impossible objects" *Visual Cognition* **7** 597 – 627

Williams P, Simons D J, 2000 "Detecting changes in novel, complex three-dimensional objects" *Visual Cognition* **7** 297 – 322

APPENDIX
Object descriptions used in experiment 3 for the object depicted in figure 2c

Real-group description: This is a baseball catcher's helmet seen from below. The 'tongue' protruding forward protects the back of the neck. Above this are the two ear protector cages, on the right and the left side, as well as further metal bars to prevent skull injury. There is padded skull protection at the back and front.

Like-group description: This could be a humane rat trap. The bait is placed inside the cage. The rat runs up the ramp which protrudes forward on the right side. This ramp springs up once the rat is in the cage. The metal catches at the top, left, and right sides of the structure automatically click shut to lock the ramp in place to prevent escape. The padding at top and bottom prevents injury to the rat.

Part-group description: This object is symmetrical and is hollow in the centre. At the front right is a flat, thin sheet which is angled downwards and is rounded at the far end. Behind and below the sheet is a dense, lumpy, hollow hemisphere. Above the hemisphere are a series of thin criss-crossing bars running horizontally and vertically. There is a second, smaller, dense region above the central bars. Both the dense sections are surrounded and topped by yet more thin bars.

ISSN 0301-0066 (print) ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 32 2003

www.perceptionweb.com