Aberystwyth University

The role of surface-based representations of shape in visual object recognition

Reppa, Irene; Greville, W. James; Leek, E. Charles

Published in:

The Quarterly Journal of Experimental Psychology

DOI:

[10.1080/17470218.2015.1014379](https://doi.org/10.1080/17470218.2015.1014379)

Publication date:

2015

Citation for published version (APA):

Reppa, I., Greville, W. J., & Leek, E. C. (2015). The role of surface-based representations of shape in visual object recognition. *The Quarterly Journal of Experimental Psychology*, 68(12), 2351-2369.
<https://doi.org/10.1080/17470218.2015.1014379>

General rights

Copyright and moral rights for the publications made accessible in the Aberystwyth Research Portal (the Institutional Repository) are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Aberystwyth Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Aberystwyth Research Portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

tel: +44 1970 62 2400
email: is@aber.ac.uk

The role of surface-based representations of shape in visual object recognition

Irene Reppa

Wales Institute for Cognitive Neuroscience,
Department of Psychology,
Swansea University, UK

W. James Greville

College of Medicine,
Swansea University, UK

&

E. Charles Leek

Wolfson Centre for Clinical and Cognitive Neuroscience,
School of Psychology,
Bangor University, UK

RUNNING HEAD: Surfaces in object recognition

Please address any correspondence to Irene Reppa at i.reppa@swansea.ac.uk, Wales Institute for Cognitive Neuroscience, Department of Psychology, Swansea University, Swansea, UK, SA2 8PP; or to Charles Leek at e.c.leek@bangor.ac.uk, Wolfson Centre for Clinical and Cognitive Neuroscience, School of Psychology, Bangor University, Bangor, UK, LL57 2AS.

ABSTRACT

This study contrasted the role of surfaces and volumetric shape primitives in three-dimensional object recognition. Observers ($N=50$) matched subsets of closed contour fragments, surfaces, or volumetric parts to whole novel objects during a whole-part matching task. Three factors were further manipulated: part viewpoint (either same or different between component parts and whole objects), surface occlusion (comparison parts either contained visible surfaces only, or a surface that was fully or partially occluded in the whole object), and target-distractor similarity. Similarity was varied in terms of systematic variation in non-accidental (NAP) or metric (MP) properties (NAPs) of individual parts. Analysis of sensitivity (d') showed a whole-part matching advantage for surface-based parts and volumes over closed contour fragments – but no benefit for volumetric parts over surfaces. We also found a performance cost in matching volumetric parts to wholes when the volumes showed surfaces that were occluded in the whole object. The same pattern was found for both same and different viewpoints, and regardless of target-distractor similarity. These findings challenge models in which recognition is mediated by volumetric part-based shape representations. Instead, we argue that the results are consistent with a surface-based model of high-level shape representation for recognition.

Word count: 195

The human visual system is remarkably adept at recognising complex three-dimensional (3D) objects despite the variability in sensory information brought about by changes in viewpoint, scale, translation and illumination. A fundamental question concerns the organisation and structure of the *object shape* representations mediating recognition (e.g., Attneave, 1954; Biederman, 1987; Cristino, Conlan, Patterson, & Leek, 2012; Davitt, Cristino, Wong & Leek, 2014; Edelman, 1999; Hummel & Biederman, 1992; Hummel & Stankiewicz, 1996; Leek, Cristino, Conlan, Patterson, Rodriguez & Johnston, 2012; Leek, Reppa, & Arguin, 2005; Leek, Reppa, Rodriguez, & Arguin, 2009; Marr & Nishihara, 1978; Pizlo, Sawada, Li, Kropatch & Steinman, 2010; Pizlo, 2008; Sawada, Li & Pizlo, 2011; Ullman, 2006).

Central to current ‘image- or appearance-based’ models is the hypothesis that recognition is accomplished by a hierarchical network of edge-based features from single contours to arrays of vertices forming viewpoint-dependent aspects (e.g., Poggio & Edelman, 1990; Riesenhuber & Poggio, 1999; Ullman, 2006). Other theories have proposed that higher-order primitives mediate recognition (e.g., Barr, 1981; Bergevin & Levine, 1993; Biederman, 1987; Biederman & Cooper, 1991; Guzman, 1968; Krivic & Solina, 2004; Marr & Nishihara, 1978; Pentland, 1986; Ullman, Vidal-Naquet & Sali, 2002; Zerroug & Nevatia, 1999). These have included 2D geons (Biederman, 1987), surfaces (e.g., Faugeras, 1984; Fisher, 1989; Leek, et al., 2005; Marr & Nishihara, 1978) and volumetric parts such as 3D geons (Biederman, 1987), generalised cylinders (Brooks, 1981; Marr & Nishihara, 1978) and super-quadrics (Barr, 1981; Pentland, 1986).

In this study we focus on the contribution of surface-based shape primitives to object recognition. Several lines of evidence indicate that surfaces are extracted early in visual processing and play a key role in visual perception (e.g., Norman & Todd, 1996; Norman, Todd, & Phillips, 1995; J.F. Norman, Todd, H.F. Norman, Clayton & McBride, 2006). Surfaces contribute to the binding of object shape and other attributes such as colour, texture and shadow (e.g., Cate & Behrmann; 2010; Chainay, & Humphreys, 2001; Fan, Medioni, & Nevatia, 1989; Faugeras, 1984; Fisher, 1989; Leek et al., 2005; 2009; Marr & Nishihara, 1978). They also influence facilitatory and inhibitory object-based attention (e.g., Leek, Reppa

& Tipper, 2003; Nakayama, He, & Shimojo, 1995; Nakayama & Shimojo, 1992; Reppa, Schmidt & Leek, 2012; Reppa & Leek, 2003; 2006), and play an important role in the perceptual analysis of shape for the planning and online control of reach-to-grasp movements (e.g. Cristino et al., 2012).

Some of the first evidence that surfaces can act as a high-level shape primitive for recognition was presented by Leek et al (2005). In a whole-part matching task observers viewed images of 3D novel objects made of two distinct volumetric parts. On each trial a whole-object was briefly presented, followed by an object part that either matched or mismatched shape information in the whole object. The part stimuli contained contour fragments, a non-volumetric configuration of spatially adjacent surfaces, or one of the two complete volumetric parts either from the same novel object (match trial) or a visually similar distractor object (mismatch trial). The main finding was that while whole-part matching for surface and volumetric parts was faster than for contour fragments, there was no difference in performance between surfaces and volumes. That is, arranging surfaces into volumetric components afforded no matching advantage over non-volumetric configurations of surfaces. On the basis of this finding, Leek et al (2005) argued that recognition is mediated not by volumetric structural descriptions, but by surface-based representations of shape computed from the perceptual input and matched to long-term memory.

Further evidence that surfaces contribute to recognition was reported by Leek et al (2009). Observers first memorised a sub-set of novel objects each consisting of two spatially connected volumetric parts. They then performed a primed recognition memory task in which they had to discriminate between learned and non-learned objects. Primes consisted of sub-sets of object surfaces. Some primes contained only surfaces that were visible in the whole (unsegmented) object while other primes contained surfaces that had been occluded in the whole (unsegmented) object. Priming effects were significantly lower for part primes with occluded surfaces, suggesting that visible surface information plays a key role in recognition. Taken together these findings, and those of Leek et al (2005), were used to support the hypothesis that object recognition is mediated by surface-based representations

of object shape (see also Ashbrook, Fisher, Robertson & Werghi, 1998; Fan, Medioni & Nevatia, 1989; Faugeras, 1984; Fisher, 1989; Lee & Park, 2002, for implementations of surface-based models in computer vision).

Current Study

The aim of the current study was to examine the role of surfaces as shape primitives for object recognition by addressing some important limitations of earlier work, and by testing further predictions of the surface representations hypothesis. First, one key issue that was not addressed in previous studies is the extent to which the use of surface primitives in object recognition generalises across changes in viewpoint. Viewpoint change may be an important factor in determining the kinds of shape representations that are used to support recognition (e.g., Arguin & Leek, 2003; Foster & Gilson, 2002; Harris, Dux, Benito & Leek, 2006; Leek, 1998a; 1998b; Leek & Johnston, 2006; Leek, Atherton & Thierry, 2007; Tarr & Bülthoff, 1998; Ullman, 1998). When two stimuli are shown from the same viewpoint an image-based matching strategy may be sufficient for determining their shape equivalence (e.g., in the extreme via a point-by-point comparison of pixel values). In contrast, judging the shape equivalence of objects across depth rotations is likely to require the computation of more abstract representations of shape – as a purely image-based matching strategy would fail. This is relevant because previous evidence from Leek et al (2005; 2009) supporting the use of surface-based primitives in recognition comes from studies that did not manipulate viewpoint – a factor that potentially limits the generality of the findings. More specifically, the presentation of stimuli from the same viewpoint in those studies may have resulted in the adoption of an image-based task strategy, potentially obscuring a genuine underlying advantage for volumetric representations. We tested this possibility in the current study by comparing performance in matching whole-part stimuli across same and different viewpoints.

Second, another issue arising from the earlier studies is whether the critical finding – that is, the equivalence in matching efficiency between surface and volumetric parts, might alternatively be accounted for in terms of a systematic difference in discriminability that masks a genuine underlying

advantage for volumetric parts. An important, and theoretically-relevant, distinction that defines discriminability is between metric (MP) and non-accidental (NAP) shape properties. MPs denote feature properties that require precise specification (e.g., aspect ratio, turning angle between contours and magnitude of curvature). In contrast, NAPs are binary categorical dimensions that distinguish image features (e.g., straight vs. curved, parallel vs. tapered - see Biederman, 1987; Lowe, 1985) – and which, under certain hypotheses, are fundamental to the derivation of higher-order part structure (e.g., Biederman, 1987). Other evidence has shown that observers are more efficient at discriminating shape on the basis of changes in NAPs than MPs (e.g., Amir, Biederman & Hayworth, 2012; Biederman & Gerhardstein 1993; Biederman & Bar, 1999) – consistent with the hypothesis that the rapid computation of NAPs plays an important role in shape perception. Of relevance here is whether the relative efficiency of whole-part matching for volumetric and surface-defined parts could be accounted for in terms of differences in the MP/NAP discriminability of each part type. An underlying advantage in matching volumetric parts over surface-defined parts may be obscured if the surface parts can be matched on the basis of NAPs, but the volumetric parts vary only in MPs. This possibility was not addressed in the earlier studies of Leek et al (2005; 2009).

Third, we aimed to further examine how surface occlusion/visibility influences object recognition. Consider for example the comparison parts for the objects in Figure 1A. Critically, the volumetric parts in the top row contain a surface that is either fully or partially occluded when both parts are combined in the whole object. In contrast, the volumetric parts in the bottom row contain only surfaces that are visible in the corresponding whole object. Surface occlusion/visibility provides a way to test the surface representation hypothesis. If visible surfaces mediate object recognition then we might expect a cost in whole-part matching between segmented volumetric parts that contain an occluded surface relative to volumetric parts that contain only visible surfaces. This is because the additional surface information of the occluded volume would produce a mismatch between the surfaces in the perceptual input, and those in the stored object representation (Leek et al., 2009).

These issues were investigated using a whole-part matching paradigm. Observers were shown twelve 3D novel objects made from two volumetric components and were later shown a comparison part, requiring a response as to whether they had previously been shown an object that had contained the part. Part stimuli comprised regions of closed contour fragments, complete volumetric parts or edge-defined surface polygons. Note that the last two types of comparison part contained surfaces of the whole object, while contour parts were closed regions of contour that did not correspond to any object surfaces. The parts were presented from either the same viewpoint as shown in the whole object, or from a different viewpoint. While all surface parts contained only surfaces that were visible in the whole object, some volumetric comparison parts contained an occluded surface ('occluded volumetric parts' shown in the top row of Figures 1A and 1B), while other volumetric parts contained only visible surfaces ('visible volumetric parts' shown in the bottom row of Figures 1A and 1B). This allowed us to examine the effect of surface occlusion on matching performance.

Finally, there were two different types of mismatch trials, manipulated between participants. In the MP participant group, mismatch trials contained wholes and object parts that primarily differed in terms of metric properties. In the NAP participant group, mismatch trials contained whole and object parts that primarily differed in terms of non-accidental properties. These contrasts are shown in Figure 1A/1B. This manipulation allowed us to examine whether the relative efficiency of making whole-part judgements for volumetric and surface-defined parts is determined by NAP/MP mismatch discriminability.

METHOD

Participants

Fifty participants were recruited from Swansea and Bangor Universities, twenty-five for each of the two participant groups. In the MP Mismatch Similarity Group participants had a mean age of 22 years ($SD=3.21$; 7 males) and took part in the experiment for either course credit or £3 payment. In the NAP Mismatch Similarity Group, participants had a mean age of 21.5 years ($SD=5.13$; 3 males) and took part

in the experiment for course credit. All participants reported normal or corrected-to-normal vision.

 INSERT FIGURES 1A AND 1B ABOUT HERE

Apparatus and Stimuli

The experiment was run on a Windows XP PC with a 19" RGB monitor using E-Prime. The stimuli were 12 novel and geometrically regular 3D objects, each of which consisted of two connected volumetric parts: a main base or large component, and a small component. They were rendered in externally lit, three-quarter views using Strata 3D Pro. Each object was scaled to fit within a 6 x 6-cm frame that subtended $6.86^\circ \times 6.86^\circ$ of visual angle from a viewing distance of 50 cm. For each of the 12 objects, closed contour, volumetric, and surface comparison (part)¹ stimuli were created (see Figure 1A and 1B). Two versions of each part type were created. The volumetric parts were one of the two volumes of the objects. The 'occluded volumetric' parts ($N=12$) contained one surface that was previously occluded in the unsegmented whole object, while the remaining surfaces were visible in the whole object. The 'visible volumetric' parts ($N=12$) comprised only of surfaces that were wholly visible in the unsegmented object.

For each volumetric part two types of surface parts ($N=24$) were also created ($N=24$). The constraints in the design were that the surface parts (a) did not make up a complete volume (b) and the number of surfaces in the surface part matched exactly the number of surfaces in the corresponding, occluded and visible, volumetric parts; and (c) the surfaces were spatially contiguous. For instance, if a volumetric part comprised of 2 visible surfaces and 1 previously occluded surface, the corresponding surface part would have 3 visible surfaces. Surfaces in both the volumetric and the surface parts were always spatially contiguous.

¹ Here and in the remaining text the term 'part' as a shorter term instead of comparison stimulus. We use the term part in its more general sense to indicate comparison stimuli consisting of image features or combinations of features that may match the primitives that are encoded in mental shape representations of the objects during perception.

The contour part stimuli ($N=24$) were created by selectively deleting regions of bounding and internal edge contour. The main constraints were that the resulting closed form did not correspond to any single volume or any single object surface, and contained edge contour from both volumetric components. Following creation of the closed form, surface information was removed by replacing the yellow colour with white (as the background).

Plane-rotated versions of each part were created by rotating $+90^\circ$ or -90° around the z axis perpendicular to the observer. These were used in the different viewpoint condition (see Design section and Figure 2).

 INSERT FIGURE 2 ABOUT HERE

In order to prevent a strategy of simple pixel-by-pixel matching between parts and whole objects in the same viewpoint condition, the whole object displays were enlarged to 150% the size of the images to be matched. Such moderate size transformations do not influence 3D-shape recognition (e.g., Fiser & Biederman, 1995; Norman et al., 2009). In addition, comparison part stimuli were centred on the screen so that the image pixels of the whole object and those of the comparison parts did not overlap.

As whole object and part stimuli necessarily differed in terms of low-level features (e.g. amount of visible edge contour and surface area; number of vertices and visible surfaces), their effect on matching performance is reported in the Results section (see *Influence of low-level image features* section). Tables 1 and 2 formally quantify and compare, respectively, those differences on each of these dimensions in two-dimensional space.

 INSERT TABLE 1 AND 2 ABOUT HERE

Design

The experiment was based on a 3 (Part Type: contour, volume, surface) x 2 (Part Viewpoint: same vs. different) x 2 (Mismatch Similarity: NAPs vs. MP) mixed design, with the latter factor manipulated between-subjects and the others within-subjects. We also included an additional factor of Surface Occlusion. This relates to whether or not volumetric part stimuli contained a surface that was either wholly or partially occluded in the unsegmented whole object. Thus, for each object, one of the volumetric parts contained an occluded surface – henceforth referred to as ‘occluded volumetric part’, while the other contained only visible surfaces – henceforth referred to as ‘visible volumetric part’. For each of these two types of volumetric parts corresponding contour part stimuli (‘occluded contour parts’ and ‘visible contour parts’) and surface part stimuli (‘occluded surface parts’ and ‘visible surface parts’) were made matched for contour length and N surfaces respectively.

In match trials, the comparison part stimulus comprised a sub-set of shape information from the whole novel object that was presented in the same trial. In mismatch trials, the comparison part stimulus belonged to a different object, with the non-matching parts differing from the whole object in terms of MP and NAP contrasts. The ratio of the contrasts was different depending on the Mismatch Similarity Group: In the MP group, the frequency of target objects and mismatch parts differing solely by MPs relative to NAPs was 2:1. In the NAP group the ratio of MP to NAP differences was 1:2 – see Figures 1A and 1B. The reason for using different ratios as opposed to pure sets of MP or NAP mismatch pairs was to avoid the possibility that participants might bias performance by selectively tuning to MP or NAP differences.

Each participant completed 144 match and 144 mismatch trials. There were 24 trials for each of the 12 within-subjects conditions. For each participant, each whole object was presented 24 times, and each part stimulus type was presented four times (twice in the same viewpoint as the whole object and twice in a different viewpoint). Trial order was randomized for each participant.

D prime (d') was chosen as the most appropriate measure of performance, because discriminations

required in the current study, especially in the MP participant Group, hinged on subtle differences, and errors were bound to be high. D' scores and the associated response bias (C) were calculated using the hit and false alarm rate per condition.

Procedure

Trial procedure is shown in Figure 3. Participants were seated approximately 50cm from the monitor. Each trial began with the central presentation of a visual prompt '*Ready*'? until the participant initiated the trial sequence by pressing the space bar. A blank screen was presented for 750ms before one of the whole object stimuli appeared at screen centre for 1200ms. Following a blank inter-stimulus interval of 750ms, a part stimulus was displayed in the centre of the screen until the participant made a response. Participants were informed that each part would be in the same orientation as the whole object preceding it, or plane-rotated clockwise (for half of the parts) or counter-clockwise (for the other half of the parts). The task was to decide as quickly and as accurately as possible whether or not the part stimulus came from the whole object that preceded it. Responses were made by pressing one of two keys (D or K) labelled '*Yes*' or '*No*' on a standard keyboard within 3 seconds. If a response was incorrect or timed out participants received feedback in the form of a 500ms error tone. Half of the participants in each Mismatch Similarity Group made match (Yes) responses with their dominant hand and mismatch (No) responses with their non-dominant hand. For the other half, these assignments were reversed. The experiment lasted approximately 35 minutes. D' scores were calculated using the hit rate per object and condition in the match trials, and the false alarm rate per object and condition in the mismatch trials (Brophy, 1986).

 INSERT FIGURE 3 ABOUT HERE

RESULTS

The mean error rate across all conditions was 32% ($SD=15.10\%$). Mean proportions of hit and false

alarm rates, d' scores and their associated response bias measure C' , per condition were calculated per object and are shown in Appendix 1. Analyses were carried out on d' scores and the same analyses on C (response bias) scores are reported in footnotes 2 and 4.

The goal of the analyses was to examine: (1) whether the pattern of whole-part matching between contour, volumetric and surface part types is modulated by viewpoint change and (2) surface visibility; and (3) whether matching performance is dependent on mismatch similarity in terms of NAP vs MP whole-part discriminability.

Whole-part matching as a function of part-type, viewpoint and surface visibility.

A 3 (Part Type: contour, volume, surface) X 2 (Part Viewpoint: same vs. different) X 2 (Surface Occlusion: occluded vs. visible) repeated-measures ANOVA was used on d' scores (the same analysis on C scores is reported in footnote 2). Cell means are shown in Figure 4A.

 INSERT FIGURE 4A ABOUT HERE

There was a significant main effect of Part Viewpoint, with better performance in same ($M=1.88$, $SD=.79$) vs. different viewpoint trials ($M=1.24$, $SD=.80$); $F(1, 23) = 45.74$, $MSE=30.47$, $p<.0001$, but Part Viewpoint did not interact with the other two factors (all $ps>.05$) – showing that the pattern of whole-part matching across part types was **not** modulated by viewpoint. There was also a significant main effect of Part Type, $F(2, 46) = 7.51$, $MSE=12.79$, $p=.002$ showing that whole-part matching performance was modulated by part type. Contour parts ($M=.96$, $SD=.66$) yielded the lowest d' scores compared to volumetric parts ($M=1.49$, $SD=.86$), $t(23) = 3.27$, $p=.003$, and to surface parts ($M=1.95$, $SD=1.15$), $t(23) = 4.11$, $p<.001$. Overall (collapsing across the factor of Surface Occlusion) surface parts yielded **higher** d' than volumetric parts, $t(23) = 2.41$, $p=.02$ –but only relative to volumes with occluded surfaces (see below). There was no significant main effect of Surface Occlusion, $F(1, 23) = 2.08$, $MSE=2.92$, $p>.05$

(occluded: $M=1.67$, $SD=.96$; visible: $M=1.46$, $SD=.66$), but there was a significant Surface Occlusion X Part Type interaction, $F(2, 44)=4.50$, $MSE=4.50$, $p=.02$. To explore this further, post-hoc analyses were conducted. These showed that performance with occluded volumetric parts and their corresponding contour parts was equivalent, $t(23)=.19$, $p>.05$, whereas surface parts yielded significantly higher discriminability (d') than either contour, $t(23)=2.58$, $p=.01$, or occluded volumetric parts $t(23)=3.00$, $p=.006$. In contrast, visible volumetric parts and their corresponding surface parts yielded higher d' scores than contour parts [$t(23)=3.07$, $p<.005$, and $t(23)=2.45$, $p<.01$, respectively]. The key contrast between performance with occluded (surface) volumetric parts and visible (surface) volumetric parts showed higher d' scores in the visible (surface) volumetric condition $t(23)=2.32$, $p=.03$. This suggests that matching performance with segmented volumetric parts is sensitive to surface composition: the presence of previously occluded surface information in the volumetric primes induces a performance cost².

Effects of mismatch similarity in terms NAP/MP discriminability. This analysis examined whether basing mismatch decisions on NAP or MP changes the pattern of matching performance across part types. As Part Viewpoint did not interact with Part Type, the data were collapsed across Part Viewpoint to simplify the analysis. The analyses were restricted to data from the visible volumetric part condition, and the corresponding matched contour and surface parts to ensure that the surface and volumetric parts contained the same number of visible surfaces. The analysis was carried out on d' scores (C scores'

² The equivalent analyses on C (response bias) scores revealed the same pattern of results as those on d' . The main effect of Part Viewpoint on C' was significant, $F(1, 23)=64.51$, $MSE=11.27$, $p<.0001$, with more conservative response bias when the part appeared in a different viewpoint ($M=.15$, $SD=.41$) as opposed to the same viewpoint ($M=-.25$, $SD=.08$) as the whole object. Also significant was the main effect of Part Type, $F(2, 46)=6.69$, $MSE=3.07$, $p=.003$. Contour parts ($M=.16$, $SD=.36$) yielded a more conservative response bias compared to volumetric parts ($M=-.17$, $SD=.40$), $t(23)=5.30$, $p<.0001$, and surface parts ($M=-.14$, $SD=.64$), $t(23)=2.81$, $p<.01$, while there was no difference between the latter two, $t(23)=.24$, $p>.05$.

The main effect of Surface Occlusion was not significant, $F(1, 23)=0.17$, $MSE=0.007$, $p>.05$ (occluded parts: $M=-.05$, $SD=.43$; visible parts: $M=-.04$, $SD=.44$). The Part Type X Surface Occlusion interaction was significant, $F(2, 46)=3.95$, $MSE=1.34$, $p=.03$. For occluded parts, there were no differences in C' among the three part types (all $ps>.05$). However, visible contour parts yielded significantly more conservative response bias relative to both volumetric and surface parts [$t(23)=5.66$, $p<.0001$ and $t(23)=2.40$, $p=.02$, respectively], with no difference among that latter two part types, $t(23)=1.36$, $p>.05$. Finally, pairwise comparisons of C' scores between occluded versus visible versions for each part type, did not reveal any significant differences (all $ps>.05$).

analyses appear in footnote 4). Cell means appear in Figure 4B.

 INSERT FIGURE 4B ABOUT HERE

A mixed 2 (Mismatch Similarity: NAP vs. MP) X 3 (Part Type: contour, volume, surface) ANOVA showed a significant main effect of Part Type, $F(2, 46) = 4.76$, $MSE = 3.86$, $p = .02$. Pairwise t-tests showed no difference in d' between volumetric ($M = 1.76$, $SD = 1.21$) and surface parts ($M = 1.61$, $SD = .77$), $t(23) = .53$, $p > .05$, while both part types had higher d' than contour parts ($M = 1.00$, $SD = .90$) [$t(23) = 3.07$, $p = .005$, and $t(23) = 2.45$, $p = .02$, respectively]. Neither the main effect of Mismatch Similarity (MPs Group: $M = 1.16$, $SD = .34$; NAPs Group: $M = 1.99$, $SD = .89$), $F(1, 23) = 3.69$, $MSE = 3.77$, $p = .07$, nor the Mismatch Similarity X Part Type interaction, $F(2, 46) < 1$, $MSE = .19$, $p > .05^3$, were significant⁴.

Influence of low-level image features. To examine potential contributions of percent edge contour, number of vertices, and surface area three one-way ANCOVAs were carried out with each of these factors as covariates. We collapsed across the variables of Part Viewpoint and Mismatch Similarity and combined the Part Type and Surface Occlusion factors into a single factor we called Part with six levels: contour 'occluded', contour 'visible', volumetric 'occluded', volumetric 'visible', surface 'occluded', and surface 'visible'. None of the three low-level features contributed significantly to the pattern of d' (all

³ Two further mixed model ANOVAs were carried out. One was on d' scores of occluded volumetric parts only and the other was on the collapsed d' means of occluded and visible surface volumes. The Mismatch Similarity X Part Type ANOVA on d' scores from occluded volumes only showed a significant main effect of Mismatch Similarity, $F(1, 22) = 8.15$, $MSE = 17.50$, $p = .01$, with higher d' for the NAP compared to the MP group. The main effect of Part Type was significant, $F(2, 44) = 6.66$, $MSE = 7.52$, $p < .0001$. Volumetric parts were not different from contour parts, $t(23) = .19$, $p > .05$, while surface parts had higher d' scores than both the contour and the volumetric parts [$t(23) = 2.58$, $p = .02$, and $t(23) = 2.98$, $p = .007$, respectively]. The Mismatch Similarity X Part Type interaction was not significant, $F(2, 44) = 2.06$, $MSE = 1.98$, $p > .05$.

⁴ An analyses of C scores revealed the same pattern of results as the one on d' scores. There was a significant main effect of Part Type, $F(2, 46) = 8.59$, $MSE = 1.87$, $p < .001$. Pairwise comparisons showed that contour parts ($M = .26$, $SD = .46$) yielded a more conservative response bias compared to both volumetric ($M = -.29$, $SD = .55$) and surface parts ($M = -.10$, $SD = .69$) [$t(23) = 5.66$, $p < .0001$ and $t(23) = 2.40$, $p = .02$, respectively], with no difference among that latter two part types, $t(23) = 1.36$, $p > .05$. The main effect of Mismatch Similarity Group was significant, $F(1, 23) = 4.54$, $MSE = 2.27$, $p = .04$, but the interaction was not, $F(2, 46) = .15$, $MSE = .33$, $p > .05$. The NAPs group yielded more conservative response bias ($M = .13$, $SE = .12$) compared to the MPs group ($M = -0.22$, $SE = .12$).

$p > .05$), hence differences in low-level image features did not account for the observed pattern of discriminability among part types. Separate regression analyses were carried out between each image feature and d' scores. Only percent edge contour correlated significantly with d' scores, ($r^2 = .73$), $F(1, 71) = 5.48$, $p = .02$. The correlation between d' scores and number of vertices was not significant, ($r^2 = .126$), $F(1, 71) = 1.1$, $p > .05$, and neither was the correlation between d' scores and surface area, ($r^2 = .058$), $F(1, 71) < 1$, $p > .05$.

GENERAL DISCUSSION

The current study used a whole-part matching task to examine the hypothesis that surfaces can act as primitives of shape representations mediating object recognition. Four findings in the current study support this hypothesis. First, where parts contained only surfaces that were visible in the whole object observers were equally good at matching parts comprising spatially adjacent surfaces to whole objects as they were in matching volumetric parts to whole objects. Second, this pattern of performance was found regardless of whether the part and whole object stimuli were presented at the same or at different viewpoints. Third, these findings cannot be accounted for by systematic differences in the discriminability of volumetric and surface parts in terms of MP or NAP differences. Fourth, there was a significant decrease in matching performance when volumetric parts contained a surface that was wholly or partially occluded in the whole object compared to volumes containing only visible surfaces. This cost arose from the mismatch in surface information derived from the comparison part and the whole object stimuli (due to the presence of a surface in the comparison part that was occluded in the whole object).

Taken together, these results provide new evidence that surface-based shape primitives can support object recognition. They also challenge theoretical accounts that recognition is mediated by volumetric part-based representations of shape (e.g., Barr, 1981; Biederman, 1987; Brooks, 1981; Marr & Nishihara, 1978; Pentland, 1986). Previous work by Leek et al (2005) showed, like the current study, that there is no benefit in matching parts defining volumetric primitives over parts defined by non-volumetric configurations of spatially adjacent surfaces. This finding is inconsistent with volumetric accounts, which

predict a matching advantage for volumetric parts (since those parts, by hypothesis, correspond to shape primitives that are computed during the course of perception). Here, we also showed that this pattern of results is found under conditions of viewpoint change which, in principle, favour the computation of volumetric part-based representations to support view generalisation. In contrast, we found positive evidence that recognition is sensitive to surfaces. This was shown by the surface occlusion effect. That is, we found a performance cost in whole-part matching for volumetric parts containing a previously occluded surface. This effect cannot be easily explained by hypotheses proposing that volumetric shape primitives mediate recognition. If the objects in the current study were represented in terms of volumes, then volumetric parts should have shown a matching advantage over contour-defined parts, or surfaces, regardless of whether they contained an occluded surface or not, because the occluded surface would be inferred as a result of volumetric completion. On the contrary, the presence of even a single partially occluded surface was detrimental to matching performance suggesting that visible surface shape contributes information that is necessary for object recognition. These data support those of Leek et al (2009) who found a reduction in priming when primes contained surfaces that were not visible in the whole object. Here a similar cost in performance was found using a whole-part matching task, providing converging evidence that visible surfaces underlie object recognition.

The current pattern of performance cannot be accounted for in terms of differences in low-level image properties: the pattern of differences in low-level image features did not resemble the observed pattern of discriminability among part types in either viewpoint. The current results can also not be accounted for in terms of global shape similarity between comparison parts and whole objects. For instance, one alternative explanation for the (occasional) good matching performance for surface parts over volumetric parts may be that surface parts looked more similar overall to the whole object than volumetric parts. However, if global shape similarity was driving matching performance there would *always* be a performance advantage for surface parts, which were the most similar to the whole object because they contained information about spatial configuration of the whole object. This explanation,

however, was not supported by the current data. The global similarity between surface parts and whole objects did not always give a performance advantage over volumetric parts. For instance, the ‘visible volumetric’ parts and ‘visible surface’ parts yielded statistically equivalent performance. Additionally, when global shape whole-part overlap was disrupted in the different viewpoint condition, again there was no difference in performance between visible volumetric and visible surface parts. Finally, in the analyses where comparison parts were matched in terms of surface visibility (they only contained visible surfaces), again there was no performance advantage for those parts (surface parts) that were more similar to the whole object than those that were not (volumetric parts). Overall, the pattern of results points to the critical importance of surface visibility in predicting accurate whole-part matching performance.

Another novel finding of the current experiments was that the observed pattern of results was not dependent on part discriminability in terms of NAPs or MPs. This is an important observation because in Leek et al (2005) mismatch parts all differed from the whole object in terms of metric properties. One concern was that an underlying advantage for matching volumetric parts may have been obscured by a necessity to compute mismatch judgements based solely on metric differences, which favour the use of image-based, rather than parts-based, representations. Here, even with a stimulus environment favouring the computation of volumetric primitives, no volumetric part advantage was observed, again implying that surfaces rather than volumetric primitives underlie performance.

More broadly, the current results support the hypothesis that object recognition is mediated by surface-based descriptions of object shape (Leek et al., 2005). According to this hypothesis, edge-based descriptions of 3D objects are used to define constituent surfaces (see also Fan et al., 1989; Fazl, Grossberg & Mingolla, 2009; Fisher, 1989; Phillips, Todd, Koenderink, & Kappers, 2003), and the surface-based description is used to access, or index, stored shape representations during recognition. On the original hypothesis outlined by Leek et al (2005) shape indexing is achieved by approximating surface shape, and accessing stored object representations based on pairwise spatial configurations of spatially adjacent surfaces. Thus, recognition is based on local surface configuration and does not require the

derivation of global object attributes (e.g., principal axis elongation, symmetry). However, the hypothesis does not assume, nor require, that surfaces are computed directly from perceptual input. Neither is it incompatible with other recent demonstrations from computational modelling supporting the use of edge- (rather than surface-) based reconstructions of 3D object geometry in human vision (e.g., Pizlo, 2008; Pizlo et al., 2010; Sawada et al., 2011). For example, Pizlo and colleagues have elegantly shown how veridical 3D structure can be reliably computed during perception from 2D edge-based descriptions of objects following simplicity constraints (e.g., symmetry, complexity). This is accomplished without inferring object surface structure directly from perceptual input, but instead is based on the recovery of a 3D ‘wireframe’ shape description. In turn, this representation can then serve as the basis for a surface-based model on which to base recognition by ‘wrapping’ surface structure over the wire-frame 3D shape model (Sawada et al., 2011).

In summary, this study contrasted the role of surfaces and volumetric shape primitives in three-dimensional object recognition. Observers matched subsets of closed contour fragments, surfaces, or volumetric parts to whole novel objects during a whole-part matching task. The results showed a whole-part matching advantage for surface-based parts and volumes over closed contour fragments – but no benefit for volumetric parts over surfaces. We also found a performance cost in matching volumetric parts to wholes when the volumes showed surfaces that were occluded in the whole object. The same pattern was found for both same and different viewpoints, and regardless of target-distractor similarity. These findings challenge models in which recognition is mediated by volumetric part-based shape representations. Instead, we argue that the results are consistent with a surface-based model of shape representation for recognition.

REFERENCES

- Amir, O., Biederman, I., & Hayworth, K.J. (2012). Sensitivity to nonaccidental properties across various shape dimensions. *Vision Research*, 62, 35-43.
- Ankrum, C., & Palmer, J. (1991). Memory for objects and parts. *Perception & Psychophysics*, 50, 141–156.
- Ashbrook, A.P., Fisher, R.B., Robertson, C. & Werghi, N. (1998). Finding surface correspondence for object recognition and registration using pairwise geometric histograms. *Proceedings of the 5th European Conference on Computer Vision*, 2, 674-680.
- Arguin, M., & Leek, E. C. (2003). Viewpoint invariance in visual object priming depends on prime-target asynchrony. *Perception & Psychophysics*, 65, 469–477.
- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, 61, 183-193.
- Barr, A.H. (1981). Superquadrics and angle-preserving transformations. *IEEE Computer Graphics and Applications*, 1, 1-11.
- Bergevin, R. & Levine, M.D. (1993). Generic object recognition: Building and matching course descriptions from line drawings. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15, 19-63.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115–147.
- Biederman, I., & Cooper, E. (1991). Priming contour-deleted images: Evidence for intermediate representations in visual object recognition. *Cognitive Psychology*, 23, 393–419.
- Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception & Performance*, 19, 1162–1182.

Brooks, R. (1981). Symbolic reasoning among 3-dimensional and 2-dimensional images. *Artificial Intelligence*, 17, 285-349.

Brophy, J. (1986). Where Are the Data? A Reply to Confrey. *Journal for Research in Mathematics Education*, 361-368.

Cate, D. & Behrmann, M. (2010). Perceiving parts and shapes from concave surfaces. *Attention, Perception and Psychophysics*, 72, 153-167.

Chainay, H. & Humphreys, G.W. (2001). The real object advantage in agnosia: Evidence for a role of surface and depth information in object recognition. *Cognitive Neuropsychology*, 18, 175-191.

Cristino, F., Conlan, L., Patterson, L.C. & Leek, E.C. (2012). The appearance of shape in visual perception: Eye movement patterns during recognition and reaching. *Proceedings of the 3rd International Conference on Appearance*, Lulu Press. Edinburgh, 125-127.

Davitt, L., Cristino, F., Wong, A. & Leek, E.C. (2014). Fixation preference for concave surface discontinuities during object recognition generalises across levels of stimulus classification. *Journal of Experimental Psychology: Human Perception & Performance*, 40, 451-456.

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

Fan, T-J., Medioni, G. & Nevatia, R. (1989). Recognizing 3-D objects using surface descriptions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11, 1140-1157.

Faugeras, O.D. (1984). New steps towards a flexible 3-D vision system for robotics. *Proceedings of the 7th International Conference on Pattern Recognition*, 796-805.

Fazl, A., Grossberg, S. & Mingolla, E. (2009). View-invariant object category learning, recognition and search: How spatial and object attention are coordinated using surface-based attentional shrouds. *Cognitive Psychology*, 58, 1-48.

Fisher, R. (1989). *From surfaces to objects: Computer vision and three dimensional scene analysis*. Chichester. John Wiley and Sons.

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-Biological Sciences*, 269, 1939-1947.

Guzman, A. (1968). Decomposition of a visual scene into three-dimensional bodies. *AFIPS Proceedings of the Fall Joint Computer Conference*, 33, 291-304.

Harris, I., Dux, P.E., Benito, C.T. & Leek, E.C. (2008). Orientation sensitivity at different stages of object processing: Evidence from repetition priming and naming. *PLoS ONE*, 3 (5), e2256

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

Hummel, J. E., & Stankiewicz, B. J. (1996). An architecture for rapid, hierarchical structural description. In T. Inui & J. McClelland (Eds.), *Attention and Performance XVI: On information integration in perception and communication* (pp.93-121). Cambridge, MA: MIT Press.

Krivic, J. & Solina, F. (2004). Part-level object recognition using superquadrics. *Computer Vision and Image Understanding*, 95, 105-126.

Lee, Y.-L., & Park, R.-H. (2002). A surface-based approach to 3-D object recognition using a mean field annealing neural network. *Pattern Recognition*, 35, 299–316.

Leek, E. C. (1998a). The analysis of viewpoint-dependent time costs in visual recognition. *Perception*, 27, 803–816.

Leek, E. C. (1998b). Effects of stimulus viewpoint on the identification of common polyoriented objects. *Psychonomic Bulletin & Review*, 5, 650–658.

Leek, E.C., Atherton, C.J. & Thierry, G. (2007). Computational mechanisms of object constancy for visual recognition revealed by event-related potentials. *Vision Research*, 5, 706-713.

Leek, E.C., Cristino, F., Conlan, L.I., Patterson, C., Rodriguez, E. & Johnston, S.J. (2012). Eye movement patterns during the recognition of three-dimensional objects: Preferential fixation of concave surface curvature minima. *Journal of Vision*, 12(1), 1-15.

Leek, E.C. & Johnston, S.J. (2006). A polarity effect in misoriented object recognition: The role of polar features in the computation of orientation-invariant shape representations. *Visual Cognition*, 13, 573-600.

Leek, E.C., Reppa, I., Rodriguez, E. & Arguin, M. (2009). Surface but not volumetric part structure mediates three-dimensional shape representation. *Quarterly Journal of Experimental Psychology*, 62, 814-829.

Leek, E.C. Reppa, I. & Tipper, S. (2003). Inhibition-of-return for objects and locations in static displays. *Perception and Psychophysics*, 65, 388-395

Leek, E.C., Reppa, I. & Arguin, M. (2005). The structure of 3D object shape representations: Evidence from part-whole matching. *Journal of Experimental Psychology: Human Perception & Performance*, 31, 668-684.

Lowe, D.G. (2004). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60, 91-110.

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

Nakayama, K, He, Z. & Shimojo, S. (1995). Visual surface representation: A critical link between lower-level and higher-level vision. In S.M. Kosslyn & D. Osherson (Eds.), *Visual Cognition: Vol 2. An Invitation to Cognitive Science* (2nd ed, pp 1-70). Cambridge, MA: MIT Press.

Nakayama, K. & Shimojo, S. (1992). Experiencing and perceiving visual surfaces. *Science*, 257, 1357-1363.

Norman, J. F., & Todd, J. T. (1996). The discriminability of local surface structure. *Perception*, 25(4), 381-398.

Norman, J. F., Todd, J. T., & Phillips, F. (1995). The perception of surface orientation from multiple sources of optical information. *Perception & Psychophysics*, 57(5), 629-636.

- Norman, J. F., Todd, J. T., Norman, H. F., Clayton, A. M., & McBride, T. R. (2006). Visual discrimination of local surface structure: Slant, tilt, and curvedness. *Vision Research*, 46(6), 1057-1069.
- Parron, C. & Washburn, D. (2010). Contrasting the edge- and surface-based theories of object recognition: Behavioural evidence from Macaques (*Macaca mulatta*). *Journal of Experimental Psychology: Animal Behavior Processes*, 36, 148-157.
- Pentland, A.P. (1986). Perceptual organization and the representation of natural form. *Artificial Intelligence*, 28, 293-331.
- Phillips, F., Todd, J.T., Koenderink, J.J. & Kappers, A.M.L. (2003). Perceptual representation of visible surfaces. *Perception & Psychophysics*, 65, 747-762.
- Pizlo, Z. (2008). *3D Shape: Its unique place in visual perception*. London: MIT Press.
- Pizlo, Z., Sawada, T., Li, Y., Kropatch, W.G., & Steinman, R.M. (2010). New approach to the perception of 3D shape based on veridicality, complexity, symmetry and volume. *Vision Research*, 50, 1-11.
- Poggio, T. & Edelman, S. (1990). A network that learns to recognise three-dimensional objects. *Nature*, 343, 263-266.
- Reppa, I. & Leek, E.C. (2003). The modulation of inhibition-of-return across object-internal structure: Implications for theories of object-based attention. *Psychonomic Bulletin & Review*, 10, 493-50.
- Reppa, I. & Leek, E.C. (2006). Structure-based modulation of IOR is triggered by object internal but not occluding shape features. *Quarterly Journal of Experimental Psychology*, 59, 1857-1866.
- Reppa, I., Schmidt, W. & Leek, E.C. (2011). Successes and failures in obtaining attentional object-based cueing effects. *Attention, Perception & Psychophysics*, 74 (1), 43-69
- Riesenhuber, M. & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, 2, 1019-1025.
- Sawada, T., Li, Y. & Pizlo, Z. (2011). Symmetry, shape, surfaces and objects. In: C. W. Tyler (Ed.). *Computer Vision from Surfaces to Objects*. CRC Press.

- Tarr, M. J., & Bulthoff, H. H. (1998). *Object recognition in man, monkey, and machine*. Cambridge, MA: MIT Press.
- Ullman, S. (2006). Object recognition and segmentation by a fragment-based hierarchy. *Trends in Cognitive Sciences*, 11, 58-64.
- Ullman, S., Vidal-Naquet, M. & Sali, E. (2002). Visual features of intermediate complexity and their use in classification. *Nature Neuroscience*, 5, 682-687.
- Ullman, S. (1998) Three-dimensional object recognition based on the combination of views. In Tarr, M.J. & Bülthoff, H.H. (Eds) *Object recognition in man, monkey and machine*. Cambridge, MA: MIT Press (pp 21-44).
- Zerroug, M. & Nevatia, R. (1999). Part-based 3D descriptions of complex objects from a single image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21, 835-848.

Table 1: Image properties of the contour, volumetric and surface comparison part stimuli.

		Percent (%) of total edge contour		<i>N</i> vertices		Surface area	
Surface Occlusion	Part Type	M	<i>SD</i>	M	<i>SD</i>	M	<i>SD</i>
	Contour	59	8.27	9	1.04	26	7.28
Occluded	Volumetric	69	8.28	10	1.31	22	7.96
	Surface	75	11.53	9	1.15	24	5.41
Visible	Contour	52	9.35	8	1.11	15	4.29
	Volumetric	43	13.68	6	1.42	11	4.42
	Surface	65	12.22	9	1.35	16	3.47

Table 2: Comparisons between each of the three comparison parts along each of three types of low-level feature. Comparisons are reported separately for each Surface Occlusion condition (occluded vs. visible). Note that the Surface Occlusion manipulation only applied to the volumetric comparison parts. Number of edge vertices refers to the total number of edge vertices (Y, T, and L) per comparison part stimulus. Surface area (calculated in centimetres using the ImageJ software, version 1.43) refers to the area enclosed by the bounding contour. Asterisks (*) follow each significant difference.

Surface Occlusion	Comparison	Low-level feature	<i>t</i> Statistic
Occluded	Contour vs. Volume	Percent (%) of total edge contour	$t(11)=2.98, p=.007^*$
		<i>N</i> vertices	$t(11)=4.52, p<.0001^*$
		Surface area	$t(11)=1.1, p>.05$
	Contour vs. Surface	Percent (%) of total edge contour	$t(11)=3.93, p<.001^*$
		<i>N</i> vertices	$t(11)=2.38, p=.04^*$
		Surface area	$t(11)=0.73, p>.05$
Visible	Volume vs. Surface	Percent (%) of total edge contour	$t(11)=1.48, p>.05$
		<i>N</i> vertices	$t(11)=6.27, p<.0001^*$
		Surface area	$t(11)=0.45, p>.05$
	Contour vs. Volume	Percent (%) of total edge contour	$t(11)=1.80, p>.05$
		<i>N</i> vertices	$t(11)=3.99, p<.001^*$
		Surface area	$t(11)=2.33, p=.03^*$
	Contour vs. Surface	Percent (%) of total edge contour	$t(11)=2.99, p=.007^*$
		<i>N</i> vertices	$t(11)=2.31, p=.03^*$
		Surface area	$t(11)=0.50, p>.05$
	Volume vs. Surface	Percent (%) of total edge contour	$t(11)=4.14, p<.0001^*$
		<i>N</i> vertices	$t(11)=5.74, p<.0001^*$
		Surface area	$t(11)=3.04, p=.006^*$

Appendix: Mean proportion of hit and false alarm rates, along with d' scores and (C) scores, per Part Viewpoint, Part Type, and Surface Occlusion shown separately for the MP Group (match/mismatch decisions made on the basis of metric differences) and the NAP Group (match/mismatch decisions made on the basis of NAP differences). Negative C values indicate a liberal response (willingness to respond 'Yes'), and positive C values indicate conservative bias (willingness to respond 'No').

MP Mismatch Similarity Group									
Part Viewpoint	Part Type	Surface Occlusion							
		Occluded				Visible			
		<i>Hits</i>	<i>FAs</i>	<i>d'</i>	<i>C</i>	<i>Hits</i>	<i>FAs</i>	<i>d'</i>	<i>C</i>
Same	Contour	.76	.36	1.18	-.21	.73	0.33	1.28	-.13
		(.20)	(.20)	(1.0)	(.33)	(.23)	(.16)	(1.15)	(.31)
	Volume	.81	.37	1.38	-.29	.89	.38	2.01	-.69
		(.13)	(.20)	(.71)	(.44)	(.10)	(.14)	(1.14)	(.74)
	Surface	.90	.39	1.87	-.61	.85	.35	1.93	-.58
		(.07)	(.14)	(.95)	(.54)	(.11)	(.11)	(1.23)	(.76)
Different	Contour	.59	.33	.73	3.16	.48	.33	.41	.41
		(.20)	(.14)	(.81)	(10.7)	(.21)	(.15)	(.78)	(.78)
	Volume	.64	.39	.71	.71	.71	.41	.88	.88
		(.13)	(.17)	(.53)	(.53)	(.17)	(.21)	(.78)	(.78)
	Surface	.75	.37	1.12	1.12	.63	.39	.68	.68
		(.13)	(.13)	(.68)	(.68)	(.21)	(.12)	(.50)	(.50)

NAP Mismatch Similarity Group									
Part Viewpoint	Part Type	Surface Occlusion							
		Occluded				Visible			
		<i>Hits</i>	<i>FAs</i>	<i>d'</i>	<i>C</i>	<i>Hits</i>	<i>FAs</i>	<i>d'</i>	<i>C</i>
Same	Contour	.76	.19	2.05	-.01	.58	.19	1.42	.43
		(.26)	(.16)	(1.64)	(.63)	(.28)	(.12)	(1.08)	(.78)
	Volume	.78	.25	1.56	-.06	.84	.29	2.31	-.41
		(.24)	(.15)	(.80)	(.56)	(.24)	(.25)	(1.87)	(.86)
	Surface	.88	.15	3.72	-.30	.82	.25	1.86	-.12
		(.25)	(.13)	(2.59)	(.88)	(.23)	(.22)	(.87)	(.60)
Different	Contour	.55	.25	.93	.29	.47	.18	.90	.46
		(.27)	(.17)	(1.09)	(.44)	(.26)	(.12)	(.89)	(.42)
	Volume	.65	.30	1.27	.20	.76	.25	1.83	.09
		(.23)	(.18)	(1.36)	(.73)	(.23)	(.20)	(1.37)	(.74)
	Surface	.77	.13	2.45	.40	.70	.19	1.97	.35
		(.26)	(.11)	(2.26)	(.19)	(.26)	(.16)	(1.11)	(1.08)

FIGURE LEGENDS

Figure 1A and 1B. The stimulus sets used for participants in the MP Mismatch Similarity Group (1A) and participants in the NAP Mismatch Similarity Group (1B). All the part stimuli are shown for each of the twelve objects, with ‘occluded’ volumetric parts, and corresponding contour and surface parts appearing in the top row and ‘visible’ volumetric parts along with their corresponding contour and surface parts in the bottom row. For each object its mismatch companion appears directly across from it. The NAP/MP column reports the type of difference between the volumes of each object when compared with its paired mismatch object. For instance, in Figure 1A there is an MP difference between the top volume of Object 1 and the top volume of Object 7. Similarly, in Figure 1B there is a NAP difference between the top volume of Object 2 (pyramid) and the lower volume of Object 8 (truncated pyramid). The ‘Type of difference’ column shows the type of MP or NAP differences between the left side and the right side volumes for each object. In some cases where the difference is in terms of NAP, there is more than one difference between the volumes (e.g., in Figure 1B, the volumes differ both in terms of the axis shape and in terms of their ending). Note: CS stands for cross section, and AS stands for aspect ratio.

Figure 2. An illustration, using Object 1, of the contrasting displays used for the same and different Part Viewpoint conditions across Part Types and Surface Occlusion. The term ‘occluded’ refers to the fact that the volumetric part contains a surface that is occluded (partially or fully) in the whole object. After that, occluded contour and surface parts are simply the corresponding contour and surface parts for that volumetric part. Similarly, the term ‘visible’ refers to the fact that the volumetric part only contains surfaces that are visible in the whole object, and the visible contour and surface parts are the corresponding parts to that volumetric part. See text for more details.

Figure 3. The trial procedure.

Figure 4. (A) Mean d' (discriminability) and C' (response bias) scores per Part Viewpoint, Part Type and Surface Occlusion (collapsed across Mismatch Similarity Group). (B) Mean d' and C' scores per Mismatch Similarity Group and per Part Type for visible contour volumetric, and surface parts only (see text for details). Error bars indicate standard error of the mean.

FIGURE 1 (A)

FIGURE 1 (B)

FIGURE 2

FIGURE 3

FIGURE 4A

FIGURE 4B