

This article was downloaded by: [Universite de Montreal]

On: 28 July 2010

Access details: Access Details: [subscription number 915526229]

Publisher Psychology Press

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



The Quarterly Journal of Experimental Psychology

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t716100704>

Surface but not volumetric part structure mediates three-dimensional shape representation: Evidence from part-whole priming

E. Charles Leek^a; Irene Reppa^b; Elly Rodriguez^a; Martin Arguin^c

^a Wales Institute for Cognitive Neuroscience, University of Wales, Bangor, UK ^b Wales Institute for Cognitive Neuroscience, University of Swansea, Swansea, UK ^c Centre de Recherche en Neuropsychologie et Cognition, University of Montreal, Montreal, Quebec, Canada

First published on: 21 August 2008

To cite this Article Leek, E. Charles , Reppa, Irene , Rodriguez, Elly and Arguin, Martin(2009) 'Surface but not volumetric part structure mediates three-dimensional shape representation: Evidence from part-whole priming', The Quarterly Journal of Experimental Psychology, 62: 4, 814 – 830, First published on: 21 August 2008 (iFirst)

To link to this Article: DOI: 10.1080/17470210802303826

URL: <http://dx.doi.org/10.1080/17470210802303826>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Surface but not volumetric part structure mediates three-dimensional shape representation: Evidence from part–whole priming

E. Charles Leek

Wales Institute for Cognitive Neuroscience, University of Wales, Bangor, UK

Irene Reppa

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

Elly Rodriguez

Wales Institute for Cognitive Neuroscience, University of Wales, Bangor, UK

Martin Arguin

Centre de Recherche en Neuropsychologie et Cognition, University of Montreal, Montreal, Quebec, Canada

The decomposition of three-dimensional (3-D) objects into shape primitives consisting of geometric volumes is a key proposal of some theories of object recognition. It implicitly assumes that recognition involves volumetric completion—the derivation of a three-dimensional structure that comprises inferred shape properties, such as surfaces, that are not directly visible due to self-occlusion. The goal of this study was to test this claim. In Experiment 1 participants memorized novel objects and then discriminated these from previously unseen objects. Targets were preceded by primes containing a subset of object surfaces that either matched those visible in the whole objects or that could only be inferred through volumetric completion. The results showed performance benefits through priming from visible surfaces but not from inferred surfaces. In Experiment 2, we found equivalent priming for part–primes containing two visible surfaces from the same volumetric part and for primes containing one surface from each of two volumes. These results challenge the view that 3-D object recognition is mediated by shape primitives comprising geometric volumes. Instead, the results support an alternative model that proposes that 3-D shapes are represented as a non-volumetric surface-based structural description.

Keywords: Object recognition; Shape representation; Primitives; Surfaces.

Understanding the representation of three-dimensional (3-D) objects is fundamental to theories of vision (e.g., Biederman, 1987; Bülthoff &

Edelman, 1992; Edelman, 1999; Hummel, 2001; Hummel & Biederman, 1992; Leek, Reppa, & Arguin, 2005; Marr & Nishihara, 1978;

Correspondence should be addressed to Charles Leek, Wales Institute for Cognitive Neuroscience, School of Psychology, University of Wales, Bangor, LL57 2AS, UK. E-mail: e.c.leek@bangor.ac.uk

This study was completed, in part, while E.C.L. was a visiting researcher at the University of Montreal. M.A. was supported by an operating grant from the Natural Sciences and Engineering Research Council of Canada.

Riesenhuber & Poggio, 1999; Vetter, Hurlbert, & Poggio, 1995). One unresolved issue concerns the shape primitives that are used to construct high-level object descriptions for recognition. According to some hypotheses volumetric parts play a key role in the representation of certain classes of objects, such as 3-D solid forms (e.g., Barr, 1981; Bergevin & Levine, 1993; Biederman, 1987; Brooks, 1981; Guzman, 1968; Hummel, 2001; Marr & Nishihara, 1978; Zerroug & Nevatia, 1999). This class of primitive includes generalized cylinders (e.g., Brooks, 1981; Marr & Nishihara, 1978), geons (Biederman, 1987), and superquadrics (Barr, 1981). Among the different volumetric models that have been described in human and computer vision some assume that these primitives consist of 3-D geometric volumes (Barr, 1981; Bergevin & Levine, 1993; Guzman, 1968; Marr & Nishihara, 1978; Zerroug & Nevatia, 1999), while others have proposed the use of symbolic propositional primitives that comprise a combination of elemental shape attributes such as nonaccidental properties (NAPs) of axis curvature, edge parallelism, and cross-sectional symmetry (e.g., Biederman, 1987; Hummel & Biederman, 1992).

Distinct from these accounts are models that propose only the use of nonvolumetric primitives, such as 2-D edge-based image descriptors or surfaces (e.g., Fan, Medioni, & Nevatia, 1989; Faugeras, 1984; Hoffman & Richards, 1984; Leek et al., 2005; Lowe, 2003). Given the important role played by volumetric primitives (as a general category of shape primitive) in some models, there is surprisingly little empirical support for their use in human vision (e.g., Biederman, 1987; Biederman & Cooper, 1991). Furthermore, some recent evidence casts doubt on whether the previously observed effects truly reflect volumetric image decomposition.

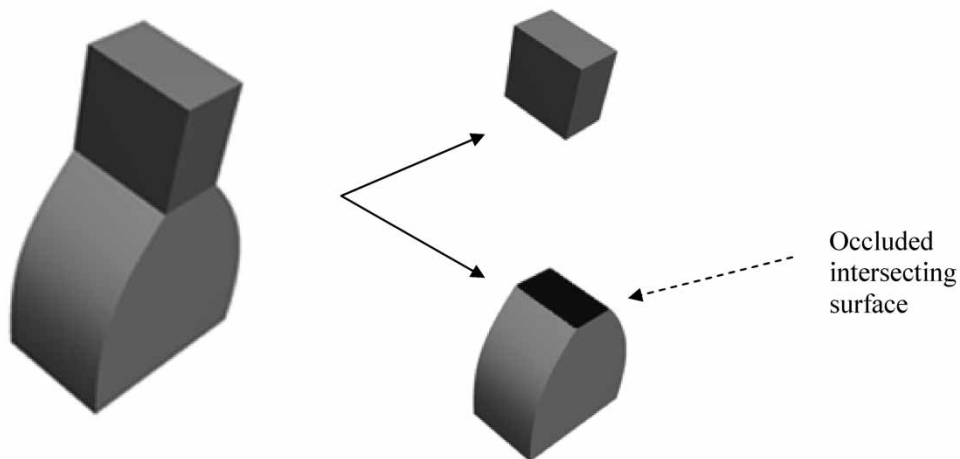
Leek et al. (2005), using a whole-part matching paradigm, found that latencies to match object parts to edge-based novel objects were equivalent for part stimuli consisting of complete volumetric parts and part stimuli containing the same number of spatially adjacent surfaces but in a nonvolumetric configuration. These findings

challenge volumetric models that predict an advantage in matching volumetrically segmented components. Instead, we proposed a surface-based structural description model in which 3-D solid objects are represented by 2-D bounded polygonal elements that are used to approximate object surfaces. The spatial configuration of visible surface elements is encoded by local pairwise relations within a surface configuration map. Of particular significance here is that the representations do not contain volumetric primitives, and the perception of object shape does not involve volumetric part segmentation.

One implicit—and hitherto untested—assumption of volumetric models that specifically propose 3-D geometric volumes (Barr, 1981; Bergevin & Levine, 1993; Guzman, 1968; Marr & Nishihara, 1978; Zerroug & Nevatia, 1999) is that recognition involves volumetric completion. This is illustrated in Figure 1a. The novel object shown could be decomposed, in principle, into two geometric volumes at regions of pairwise concave minima of curvature (e.g., Hoffman & Richards, 1984). Such decomposition involves the approximation of volumetric parts that results in the explicit representation of occluded 3-D image structure. This includes information about the previously occluded or “hidden” surface at the intersection of the two volumetric parts, as well as information about the surfaces on the occluded sides of the object. We refer to the former as occluded intersecting surfaces.

The explicit representation of occluded intersecting surface structure follows from the use of geometric volumes as primitives in shape representation and is a point of departure from other models, for example, that propose only the use of nonvolumetric primitives (e.g., Leek et al., 2005). The aim of the current study is to evaluate the role of such occluded volumetric structure in recognition in order to contrast predictions made by accounts that posit geometric 3-D volumes (e.g., Barr, 1981; Bergevin & Levine, 1993; Guzman, 1968; Marr & Nishihara, 1978; Zerroug & Nevatia, 1999) versus one model that does not—the surface-based representations model (Leek et al., 2005).

(a)



(b)



(i) Visible-surface

(ii) Occluded-surface
(Volumetric)

(iii) Visible-surface

(iv) Occluded-surface

TYPE I PRIMES

TYPE 2 PRIMES

Figure 1. Illustration of the distinction between visible and occluded surfaces that is exploited in Experiment 1. (a) From a given viewpoint volumetric image segmentation yields two volumetric components that contain a previously occluded surface at the segmentation point (black) through volumetric completion. (b) The four prime types used in the study. See text (Experiment 1, Method Section, Stimuli) for details.

EXPERIMENT 1

In Experiment 1, participants memorized the shapes of novel 3-D solid objects each containing two putative volumetric parts and then discriminated these targets from previously unseen objects. Whole objects were preceded by part primes containing a subset of object surfaces. There were two basic prime conditions—see Figure 1b. Visible-surface primes contained only

conjunctions of surfaces that were directly visible in the whole object. Occluded-surface primes also contained an occluded intersecting surface at the boundary of the two volumetric parts that would only be explicit following volumetric image segmentation. Volumetric theories that posit the representation of a 3-D geometric primitive predict a larger priming effect for primes that can be readily matched to volumetric parts than for primes that cannot—even when the former

contain occluded surface structure derived from volumetric completion. That is, occluded surface structure should contribute to the priming effect. In contrast, the surface-based representations model (Leek et al., 2005) predicts a different pattern. On this account only visible structure is encoded during perception. Thus priming should depend only on the degree of match between visible surfaces in the primes and whole objects. That is, visible-surface primes should produce a larger priming effect than occluded-surface primes.

Method

Participants

A total of 20 adults aged between 18 and 40 years from the School of Psychology, University of Wales, Bangor, UK and the Département de Psychologie, Université de Montréal, Canada,

participated in this experiment for either course credit or payment of CAN \$10. All were naïve as to the aim of the study and had normal or corrected-to-normal eyesight. None reported any history of developmental disorders.

Stimuli

The stimuli were 12 rendered novel 3-D solid objects created using Strata 3D (Strata Corp, USA). The stimuli are shown in Figure 2.

All objects were coloured in mustard yellow (RGB 197, 165, 69) in a uniform metal texture and were rendered using ray tracing at a resolution of 72 dpi and using antialiasing by super sampling and pixel averaging at a resolution of 3×3 pixels for each pixel in the final image. Images were rendered with a single illumination source. Whole objects were drawn and rescaled to fit within a 900×900 -pixel frame subtending 16.72 degrees of visual angle from the viewing distance of




Figure 2. Experiment 1. The 12 novel objects used in the test phase (not to scale).

60 cm. Each whole object was made by combining two volumetric components at a clearly defined region of paired concave minima of curvature (Biederman, 1987; Hoffman & Richards, 1984). The volumetric components were defined by variation of the following parameters: curvature of the main axis, tapering (parallelism), edges (straight vs. curved), and symmetry of the cross section (Biederman, 1987). Visual similarity among components in the object set varied according to changes in these parameters. Each volumetric part could be uniquely specified by a combination of NAP relations and aspect ratio. All stimuli shared the same spatial configuration in which one component was attached to the other by an “end-on” relation. This ensured that discrimination among stimuli required attention to the shapes of the individual components. The stimuli were depicted from a single three-quarter viewpoint that was chosen to maximize visibility of object structure. The two-component structure of each stimulus was verified in a pilot study in which naïve observers ($N = 10$) indicated the number of volumes and the segmentation point between the two volumetric parts. There was 100% agreement for all 12 stimuli about the number of volumetric parts, and the location of the volumetric part boundary.

For each object, four part priming stimuli ($n = 48$) were created (see Figure 1). These were divided into two types (see Figure 1): Type 1(i)

visible-surface primes contained a subset of complete spatially adjacent visible object surfaces but did not form a volumetric component; Type 1(ii) occluded-surface (volumetric) primes consisted of a complete volumetric component containing several spatially adjacent visible surfaces and the occluded intersecting surface. These were matched to their corresponding Type 1 visible-surface primes for the total number of surfaces shown. Type 2(iii) visible-surface primes contained two visible surfaces from the same volumetric part. Type 2(iv) occluded-surface primes contained one visible and one occluded intersecting surface from the same volumetric part. Type 2 primes served as controls for differences in information about the global spatial configuration between Type 1 occluded-surface (volumetric) and Type 1 visible-surface primes. Primes were scaled to fit within a 600×600 -pixel frame (maintaining aspect ratio) subtending 11.7 degrees of visual angle from the viewing distance of 60 cm.

The low-level image properties of the primes were controlled as these factors might also be expected to contribute to any observed priming effects (Leek et al., 2005). Table 1 shows for each prime condition: (a) The mean amount of edge contour per prime and (b) the mean number of L and Y (including arrow junctions) vertices, along with the total mean number of vertices of all types (L, Y, & T) per prime.

Table 1. Experiment 1: Image properties for the priming stimuli in each condition

Prime type		Edge contour (cm)		Vertex type (N)					
		<i>M</i>	<i>SD</i>	<i>L</i>		<i>Y</i>		Total vertices	
				<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Visible	Type 1	30.82	6.99	4.50	1.00	2.67	0.89	7.92	1.24
	Type 2	20.36	4.20	3.83	0.39	1.67	0.78	5.50	1.00
	Overall	25.58	7.76	4.16	0.81	2.16	0.96	6.70	1.65
Occluded	Type 1	26.10	5.00	2.08	1.08	4.75	1.66	6.92	1.31
	Type 2	17.54	2.78	3.92	0.29	1.92	0.29	5.92	0.51
	Overall	21.58	5.89	3.00	1.21	3.33	1.85	6.41	1.10

Note: Table shows (a) total visible edge contour, (b) mean number of L and Y vertices per prime, and (c) mean number of vertices of all types (L, Y, T) per prime.

There were no significant differences in mean edge contour between Type 1 visible-surface and Type 1 occluded-surface (volumetric) primes, or between Type 2 visible-surface and Type 2 occluded-surface primes. There was also no overall difference in mean edge contour per prime between visible-surface and occluded-surface primes collapsing across prime type. Analyses of the vertices showed that Type 1 visible-surface primes contained significantly more L vertices than Type 1 occluded-surface (volumetric) primes (Mann-Whitney U Test, $U = 7.00$, $Z = 3.75$, $p < .0001$), but there was no significant difference in L vertices between Type 2 primes. Overall, visible-surface primes contained significantly more L vertices than did occluded-surface primes collapsing across prime type ($U = 125.00$, $Z = 3.36$, $p < .0007$). Type 1 visible-surface primes contained significantly fewer Y vertices (which included both Y and arrow junctions) per prime than did Type 1 occluded-surface (volumetric) primes ($U = 20.00$, $Z = -3.00$, $p < .002$), but there was no difference between Type 2 primes. Collapsing across prime type visible-surface primes contained significantly fewer Y vertices than did occluded-surface primes ($U = 195.00$, $Z = -1.90$, $p < .05$). There were very few T vertices per prime in any condition ($M < .03$) so separate analyses are not reported. Instead, the mean total number of vertices per prime, including L, Y, and T vertices, was calculated. There was a significantly higher mean number of total vertices for Type 1 visible-surface primes than Type 1 occluded-surface (volumetric) primes ($U = 36.5$, $Z = 2.04$, $p < .04$), but no difference between Type 2 primes. There was no overall difference between visible- and occluded-surface primes collapsing across prime type.

These analyses show that the low-level image properties were matched for mean edge contour and all vertex types for Type 2 primes, but differed for L, Y, and mean total vertices between Type 1 primes. We consider the potential influence of these differences on the observed priming effects later.

Several steps were also taken to eliminate low-level feature overlap between edges in the primes

and whole objects: (a) Whole objects and primes were aligned in the centre of their respective frames—so that the locations of edges in the related primes did not correspond to the pixel (screen) locations of the same edges in the whole objects. (b) Whole objects were resized to 150% to displace feature locations between given edges in the primes and whole-object stimuli. (c) Four pattern masks were created consisting of random segments from all 12 whole-object images. Masks were presented randomly at four image plane orientations (0, 90, 180, or 270 degrees). Masks were $1,024 \times 1,024$ pixels subtending 18.92 degrees of visual angle from 60 cm.

Design

There were four within-subjects factors in a 2 (response: target, nontarget) \times 2 (prime relatedness: related, unrelated) \times 2 (prime surface visibility: visible, occluded) \times 2 (prime type: Type 1, Type 2) design.

The experiment consisted of a learning phase and a test phase. In the learning phase participants first memorized six objects. Two groups of participants were tested. Group 1 memorized six of the stimuli from the novel object set. Group 2 memorized the other six objects. For each group the remaining unlearned objects served as nontargets for the NO response trials in the test phase. Participants were randomly assigned to each group. During the learning phase participants were shown each object from a single viewpoint for an unlimited duration. The next object was shown when the participant indicated that she or he had memorized the shape. Following this, participants completed 18 learning trials in which targets and a further three distractors (not used in the experimental phase) were each shown twice. Participants indicated whether or not the object was a target or distractor. Responses were made via a key press indicating “Yes” (target) or “No” (nontarget). A criterion of 80% correct had to be obtained before participants were allowed to complete the test phase. The test phase comprised two identical blocks of trials. In each block there were a total of 108 trials consisting of 54 target and 54 nontarget trials. For both

target and nontarget trials there were 24 related prime trials, 24 unrelated prime trials, and 6 no-prime trials. These were included to allow calculation of priming effects in relation to a neutral (no-prime) condition. Across blocks there were 12 trials per prime condition, including no-prime trials (N total = 216). In related trials, the priming stimulus matched part of the whole object presented on the same trial. In unrelated trials, the priming stimulus did not match the whole object. Primes in the unrelated condition were randomly paired within each prime category with one of the 12 whole objects. There were 12 practice trials. Trial order was randomized within each block.

Procedure

Participants first completed the learning phase. In the test phase all trials began with a centrally presented fixation cross lasting 750 ms. After a blank interstimulus interval (ISI) of 500 ms, a prime (or no-prime interval) was presented in the centre of the screen for 250 ms followed by a further blank ISI of 150 ms and a centrally presented pattern mask (at 0, 90, 180, or 270 degrees) for 250 ms and a further blank ISI of 250 ms. Following this a whole object (either a target or a nontarget) was presented at the centre of the screen and remained until the participant responded or a response deadline of 5,000 ms was exceeded. Incorrect responses were immediately followed by an error message displayed in the centre of the screen. Following the response the screen was cleared, and a new trial commenced. Feedback on performance accuracy was provided

for incorrect responses in both learning and test phases. The experiment lasted 20 minutes.

Results

Trials with response times (RTs) that deviated by more than ± 2 standard deviations from the mean in each condition were removed from the data ($< 2.5\%$ of all responses). Error trials, where participants gave an incorrect response, were also discarded (4.49% of responses). Overall performance was highly accurate (see Table 1). Mean percentage error rates across participants were 4.05% ($SE = 1.10\%$) for related trials and 5.52% ($SE = 1.99\%$) for unrelated trials. A Friedman nonparametric analysis of variance (ANOVA) did not show any significant differences in error rates across prime conditions, $\chi^2(7, N = 20) = 8.48$, ns . There was a significant positive correlation between RTs and error rates; $r = .27$; $F(1, 158) = 13.16$, $p < .0003$. This indicates that there was no speed-accuracy trade-off.

Mean RTs were faster for target (YES response, $M = 700.33$ ms, $SE = 20.75$ ms) than for nontarget trials (NO response, $M = 788.19$ ms, $SE = 26.99$ ms), $t(19) = -4.02$, $p < .0007$. The remaining analyses were restricted to target present (YES response) trials. Table 2 shows the mean RTs and error rates for target present trials as a function of prime relatedness, prime surface visibility, and prime type.

A repeated measures ANOVA was carried out for target present trials using a 2 (prime relatedness: related vs. unrelated) \times 2 (prime surface visibility: visible vs. occluded) \times 2 (prime type: Type 1 vs. Type 2) design. There were significant main

Table 2. Experiment 1: Mean reaction times and percentage error rates for related and unrelated target present trials per prime condition

Prime type		Related				Unrelated			
		<i>M</i>	<i>SE</i>	% Error	<i>SE</i>	<i>M</i>	<i>SE</i>	% Error	<i>SE</i>
Visible	Type 1	596.26	21.52	6.35	1.55	664.50	21.22	5.42	2.33
	Type 2	615.25	21.25	3.17	0.94	676.22	26.08	5.83	1.78
Occluded	Type 1	639.58	31.18	3.33	0.95	677.22	23.25	4.17	1.44
	Type 2	644.46	29.24	3.33	1.43	654.90	17.08	6.67	2.43

Note: Related and unrelated target present trials: YES response.

effects of prime relatedness, $F(1, 19) = 23.26$, $p < .0001$, and prime surface visibility, $F(1, 19) = 7.88$, $p < .01$. The main effect of prime type was not significant, and this factor did not participate in any significant interaction. There was a significant interaction between prime relatedness and prime surface visibility, $F(1, 19) = 5.13$, $p < .03$. Separate 2 (prime surface visibility) \times 2 (prime type) repeated measures ANOVAs showed that for related trials there was a significant main effect of prime surface visibility, $F(1, 19) = 9.14$, $p < .001$, indicating that mean RTs were faster for visible than for occluded surface primes. There was no main effect of prime type and no interaction. For unrelated trials there were no significant effects indicating that mean RTs were equivalent across conditions. Planned contrasts revealed significant differences in mean RTs between related and unrelated trials for both types of visible-surface primes—Type 1, $t(19) = 5.57$, $p < .0001$; Type 2, $t(19) = 3.58$, $p < .001$ —and Type 1 occluded-surface (volumetric) primes, $t(19) = 2.40$, $p < .02$, but not for Type 2 occluded primes, $t(19) = 0.49$, *ns*. On the basis of the latter analysis, the Relatedness \times Surface Visibility interaction is attributed to the greater RT difference between related and unrelated trials (i.e., priming effect) for the visible-surface primes ($M = 64.61$ ms, $SE = 14.26$) than for the occluded-surface primes ($M = 24.04$, $SE = 18.34$).

Further analyses were conducted on the mean priming effects (mean unrelated RTs – mean related RTs) for target present trials across conditions. Mean priming effects per condition are shown in Figure 3. Overall priming for visible prime trials was significantly greater than that for occluded prime trials, $t(39) = 2.72$, $p < .009$. A 2 (prime surface visibility) \times 2 (prime type) repeated measures ANOVA showed a significant main effect of surface visibility, $F(1, 19) = 5.13$, $p < .03$, and no interaction.

Planned contrasts showed that the predicted difference in priming between Type 1 visible-surface and Type 1 occluded-surface (volumetric) primes was significant, $t(19) = -4.02$, $p < .04$, as was the difference between Type 2 visible-surface

and Type 2 occluded-surface primes, $t(19) = -2.03$, $p < .02$.

To establish the reliability of the priming effects and to show that they are not an artefact of the calculation method used, mean priming was also assessed by subtracting mean RTs for related trials from mean RTs for targets in the no-prime condition. The same pattern of priming effects as that above was found. A 2 (surface visibility) \times 2 (prime type) repeated measures ANOVA on mean priming effects for the related minus no-prime case showed a significant main effect of surface visibility, $F = 9.1$, $p < .007$, and no interaction. Planned contrasts showed significantly more priming for Type 1 visible-surface primes over Type 1 occluded-surface (volumetric) primes, $t(19) = 2.67$, $p < .005$, and for Type 2 visible-surface over Type 2 occluded-surface primes, $t(19) = 1.92$, $p < .03$.

Analyses of priming effects from low-level image properties

Further analyses were also undertaken to determine the influence of low-level image properties of the primes on the observed pattern of priming effects. For Type 1 primes, visible-surface and occluded-surface (volumetric) primes differed in the mean number of visible L, Y, and T vertices per prime. However, correlations between the size of the priming effects and the mean number of L, Y, and total vertices per prime were not significant ($r^2 < 1$), $F(1, 22) < 0.8$, in all cases. In addition, as noted earlier, Type 2 primes were matched for mean edge contour and for the mean number of L, Y, and total vertices, so the difference in this condition between visible-surface and occluded-surface primes cannot be accounted for in terms of low-level image differences.

Discussion

The results of Experiment 1 show that priming is dependent on the match between visible surfaces in the primes and those in the targets. The data provide no evidence that occluded intersecting surfaces derived from volumetric completion contribute to shape priming. This finding appears to


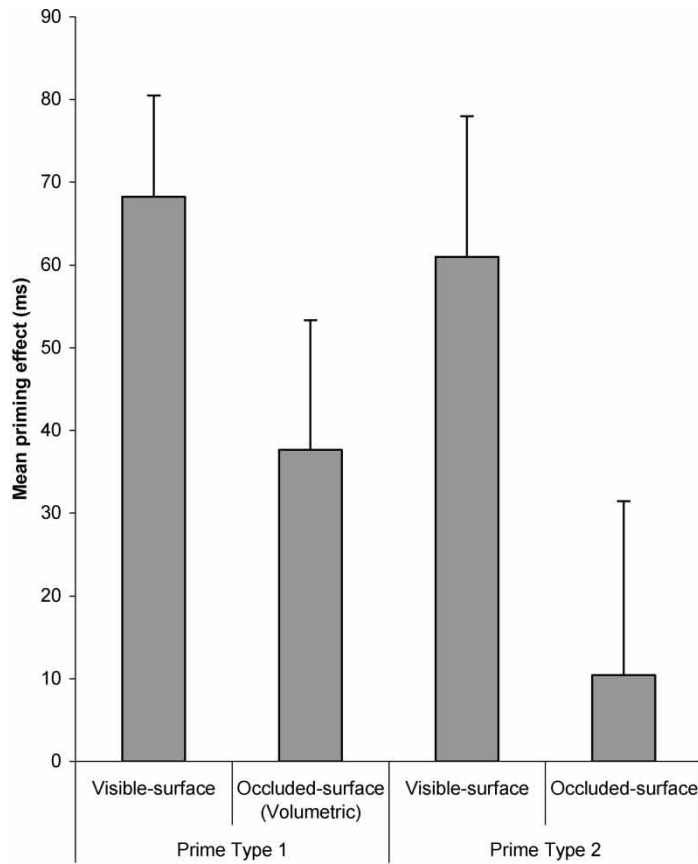
challenge the predictions of models that posit a role for 3-D geometric volumes in shape representation. However, a number of methodological points must be addressed before the theoretical implications of these findings can be fully considered.

First, it might be argued that the priming advantage for visible- over occluded-surface (volumetric) primes stems from differences between conditions in prime-target similarity. For example, although the volumetric parts of each stimulus were uniquely defined by a variation of four NAP relations and aspect ratio (Biederman, 1987), the complete volumetric segments comprising the Type 1 occluded-surface (volumetric) primes might have been more similar to components in both the target and the nontarget object sets than were the visible-surface primes. Such a bias would make occluded-surface primes less predictive of target identity because a given prime may share volumetric shape properties not only with its own target, but also with several other nontarget items as well. If so, priming effects in the Type 1 occluded-surface (volumetric) condition might be lower because unrelated primes, in this condition, also prime responses to targets, thus reducing the difference in mean RTs between related and unrelated trials. In this case one would also expect to find differences in mean RTs between visible- and occluded-surface primes in the unrelated trials, reflecting greater prime-target/nontarget similarity in the occluded-surface conditions. But contrary to this prediction, mean RTs across unrelated trial conditions were not significantly different. In addition, this argument cannot explain the difference between Type 2 visible- and occluded-surface primes, which did not contain complete volumetric parts.


Second, another possibility is that the advantage for Type 1 visible-surface primes derives from some combination of partial priming of a subset of the elemental shape features of each

volume (e.g., the recoverable NAPs) together with priming of the spatial relations between the two inferred volumetric components. For example, we might assume that the priming of volumetric parts increases as some function of the number of visible surfaces (or recoverable NAPs) in the prime, and that the recovery of partial shape information about two volumes will elicit a further gain from priming volumetric configuration. On this summation account the combination of these two sources of priming (partial shape plus volumetric configuration) might give rise to a net effect of facilitation in the Type 1 visible-surface condition (in which both volumes and their configuration were partially primed) relative to the Type 1 occluded-surface (volumetric) condition (in which an entire single volume was primed, but not volumetric configuration). This may be unlikely as all of the stimuli in the object set shared to the same configuration consisting of an "end-on" relation between the two putative volumetric components. Thus, volumetric spatial configuration was uninformative about object identity and might not be expected to contribute to a priming advantage for related over unrelated trials.¹ Even so, in order to examine this possibility further, and to better understand the source of the priming effects, we conducted a second study in which the primes always contained only two visible surfaces—see Figure 5(b). In one condition the two surfaces came from a single volumetric part (same-part primes). In the other case, primes contained one surface from each of two connected volumes (different-part primes). According to the summation account, priming for the same-part primes should be less than priming for the different-part primes, since the latter will consist of the summation of partial priming of the shape of each volume together with the priming of volumetric spatial configuration. In contrast, the surface-based representations model (Leek et al., 2005)

¹ It should be noted that while Type 2 primes contain two surfaces from the same volume (and consequently neither condition should elicit partial priming of volumetric configuration), they differ in the number of visible surfaces that are shown. As a consequence, the observed pattern of priming could be consistent with some version of the summation account. We thank an anonymous reviewer for this observation. Experiment 2 addresses this possibility.



Prime condition



Target

Figure 3. Experiment 1. Mean reaction times for priming effects across conditions in the target present (YES response) trials. Error bars show standard error of the mean.




Figure 4. Experiment 2 (a). Top panel illustrates an image segmentation point at a volumetric-part boundary along with the constituent volumetric primitives of a novel 3D object. (b) Bottom panel shows an example of the corresponding primes for this object: (i) Same-part primes contained two spatially adjacent surfaces but no volumetric part boundary; (ii) Different-part prime contained two spatially adjacent surfaces connected by a volumetric part boundary.

predicts significant but equivalent priming between the same-part and different-part primes (that is, no interaction of prime type). This is because the surface-based model attributes no special functional status to information about the spatial configuration of volumetric parts, and both primes have equivalent visible surface information. Additionally, in Experiment 2, we used a new set of novel stimuli consisting of contour-based line drawings (rather than surface rendered shapes) in which the spatial configuration of the volumetric parts was varied across items in the stimulus set. This was done both to test the generality of the previous findings to a type of contour-based stimulus set that has been used in previous studies to examine

part-based processing (e.g., Biederman & Cooper, 1991) and also to increase the likelihood that information about volume configuration would potentially contribute to the observed priming effects.

EXPERIMENT 2

Method

Participants

The participants were 22 adults (18–44 years old) recruited from the School of Psychology, University of Wales, Bangor. All were naïve about the aim of the study, and all had normal or corrected-to-normal eyesight and reported having no known history of developmental disorders.

Stimuli

The stimuli consisted of the 12 contour-based novel objects shown in Figure 4.

Each object was depicted as an opaque white line drawing with black edges and was drawn to fit within a 900×900 -pixel frame, subtending 16.72 degrees of visual angle from a viewing distance of 60 cm. As in Experiment 1, the objects were each made by combining two simple volumetric components at a clearly defined region of paired concave minima of curvature—see Figure 5(a). All of the objects were unique, sharing no volumetric parts with any other object in the stimulus set. Volumes were defined by variation among the same NAP parameters as those described in Experiment 1. Subjective impressions of volumetric part structure were confirmed by 20 naïve raters. There was 100% agreement among raters about the number of volumetric parts per stimulus and location of the part boundary. Unlike Experiment 1, stimuli varied in the spatial configuration of the two volumes. Thus, configural information was potentially informative about object identity. For each object two types of part prime were created—see Figure 5(b): (a) *same-part* primes ($n = 12$), which contained two spatially adjacent surfaces from the same volumetric part; (b) *different-part* primes ($n = 12$), which contained two spatially adjacent surfaces

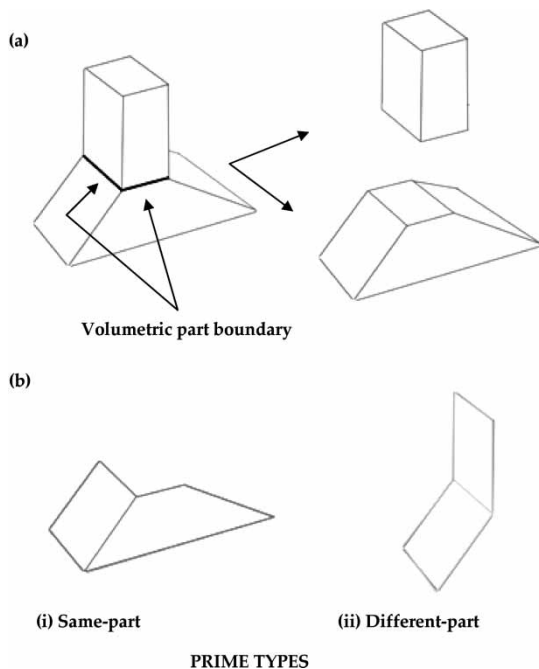


Figure 5. Experiment 2. The 12 novel objects used in the test phase (not to scale).

with one taken from each volumetric part. Same- and different-part primes were matched for low-level feature visual complexity—see Table 3.

There were no significant differences between same- and different-prime types in terms of total visible image contour, $t(23) = -1.19$, *ns*, number of L vertices (Mann-Whitney: $U = 15.00$; $Z = -0.23$, *ns*), number of Y vertices ($U = 40.00$; $Z = 0$, *ns*), or total number of vertices ($U = 40.5$; $Z = 0.36$, *ns*). Prime surface area was also equated between same- and different-part prime

conditions, $t(23) = 0.61$, *ns*. All primes were depicted from the same viewpoint as the whole object from which they were derived. Primes were scaled to fit within a 600×600 -pixel frame subtending 11.7 degrees of visual angle from the viewing distance of 60 cm. This avoided low-level contour overlap between the primes and targets.

Design

There were three within-subject factors: response (target vs. nontarget), prime relatedness (related vs. unrelated), and prime type (same-part vs. different-part). The structure was identical to that of Experiment 1. In the learning phase participants memorized 6 of the 12 objects. Two groups of participants were tested. Group 1 memorized 6 of the stimuli from the novel object set. Group 2 memorized the other 6 objects. For each group the remaining unlearned objects served as nontargets for the no response trials. Participants were randomly assigned to each group. Learning first involved copying each object on a blank sheet of paper as accurately as possible. Participants then completed a computerized learning task (two blocks of 36 trials) in which they viewed each of the 12 stimuli one at a time for a maximum time of 5,000 ms. The task was to decide whether the presented object was a previously memorized target or a distractor. Participants responded via a key press indicating “yes” (target) or “no” (nontarget). Responses were counterbalanced for handedness within each group. A criterion of 80% correct had to be obtained before participants were allowed to complete the test phase. The test phase contained 48 yes (target) and 48 no

Table 3. Experiment 2: Stimulus properties of primes in each priming condition

	Edge contour (cm)		No. surfaces		No. vertices					
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>Y</i>		<i>L</i>		<i>Total</i>	
Prime type					<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Same part	18.1	3.9	2.0	0.0	2.00	0.0	4.10	0.4	6.10	0.4
Different part	19.3	3.3	2.0	0.0	2.17	0.3	5.04	0.9	7.21	1.0

(nontarget) response trials, each consisting of 24 related prime trials (12 same-part and 12 different-part) and 24 unrelated prime trials. In related trials the priming stimulus matched part of the whole object presented on the same trial. In unrelated trials the priming stimulus did not match the whole object. For unrelated trials prime-whole-object pairs were selected on the basis of global visual similarity in the spatial configuration and shape of the volumetric components. There were 96 experimental trials in total presented in two blocks of 48. There were also 12 practice trials. Trial order was randomized within blocks.

Procedure

Participants first completed the learning phase. Test phase trials began with a centrally presented fixation cross (750 ms). Following a blank ISI (700 ms) a prime was presented in the centre of the screen. Prime duration was 150 ms. The prime was replaced by a centrally presented whole object (ISI = 500 ms), which remained on the screen until the participant made a key response. Since both prime and whole-object displays were centred, and the primes were rescaled (see above), there was no perceptual overlap in matching contours between primes and whole objects. There was a response deadline of 5,000 ms. Incorrect responses were indicated by error message and a tone. The experiment lasted about 20 minutes.

Results

Trials with RTs that deviated by more than ± 2 standard deviations from the mean in each condition were removed from the data (<2.5% of all responses). RTs for trials where participants gave

an incorrect response were also eliminated (13.07% of all responses). Mean percentage error rates across participants were 14.2% ($SE = 3.30\%$) for related trials and 11.93% ($SE = 2.83\%$) for unrelated trials. A Friedman nonparametric ANOVA on error frequency was significant, $\chi^2(3, N = 22) = 17.32, p < .01$. This reflected a higher overall error rate to nontarget NO response trials (mean percentage error = 9.13%, $SE = 2.17\%$) than to target YES response trials (mean percentage error = 4.96%, $SE = 1.42\%$); Wilcoxon, $T = 221.50; Z = 3.88, p < .0001$. There were no significant differences between conditions within the YES and NO response categories. There was a significant positive correlation between RTs and error rates; $r = .53; F(1, 86) = 33.86; p < .0001$. This suggests that there was no speed-accuracy trade-off.

Mean RTs were faster for target ($M = 676.59$ ms, $SE = 27.97$ ms) than for nontarget trials ($M = 745.89$ ms, $SE = 25.45$ ms), $t(21) = 3.01, p < .006$. The remaining analyses were restricted to target present (YES response) trials. Table 4 shows the mean RTs and error rates as a function of prime relatedness and prime type for target present (YES response) trials.

A repeated measures ANOVA was carried out using a 2 (relatedness: related vs. unrelated) \times 2 (prime type: same-part vs. different-part) design. There was a significant main effect of relatedness, $F(1, 21) = 14.20, p < .001$, but no main effect of prime type, $F(1, 21) = 1.39, ns$, and no interaction, $F(1, 21) = 0.01, ns$. This is shown in Figure 6.

Planned contrasts on target (yes) response data showed that mean RTs for related same-part primes were significantly faster than those for

Table 4. Experiment 2: Mean reaction times and percentage error rates for related and unrelated target present trials per prime condition

Prime type	Related				Unrelated			
	<i>M</i>	<i>SE</i>	% Error	<i>SE</i>	<i>M</i>	<i>SE</i>	% Error	<i>SE</i>
Same part	691.77	22.84	6.75	1.48	739.29	31.97	4.76	1.58
Different part	677.36	25.90	2.78	1.05	722.89	35.44	5.56	1.56

Note: Reaction times in ms. Related and unrelated target present trials: YES response.

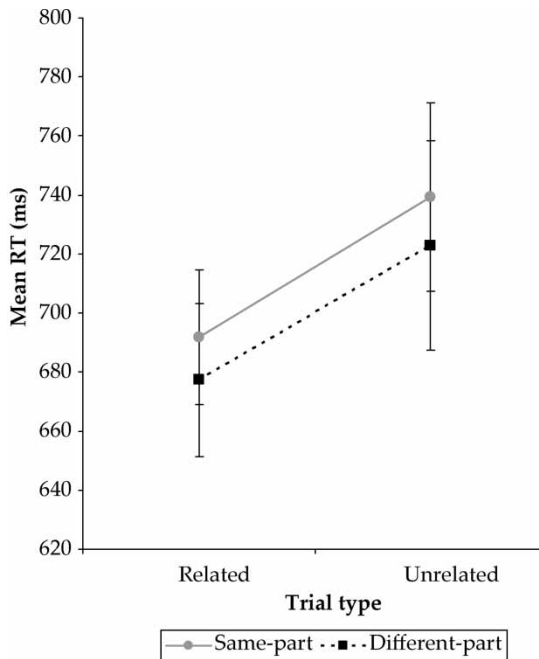


Figure 6. Experiment 2. Mean reaction times for related and unrelated target present (YES response) trials in the same- and different-part prime conditions. Bars show standard error of the mean.

unrelated same-part primes, $t(21) = 3.08$, $p < .005$, showing a mean priming effect of 47.52 ms ($SE = 15.77$ ms). Mean RTs for related different-part primes were also significantly faster than those for unrelated different-part primes, $t(21) = 2.87$, $p < .009$, with a mean priming effect of 45.53 ($SE = 16.21$).

Discussion

The results of Experiment 2 showed that both same-part and different-part primes produced significant facilitatory priming effects. Critically, both conditions contained the same number of visible surfaces, and although the different-part primes could be used to infer volumetric configuration, the size of the priming effects did not interact with prime type. That is, same-part and different-part primes produced equivalent amounts of priming. It should be noted also that primes in the two conditions were also precisely matched in

terms of low-level image properties that might be expected to influence shape priming. This finding is inconsistent with the version of the summation account described earlier, which predicted that the combined partial priming of two volumetric parts, together with information about volumetric configuration, would be greater than the sum of partial priming from the visible surfaces of a single volume alone. In contrast, the pattern of priming between the two conditions is consistent with the surface-based representations model, which predicted that priming should be equivalent between these conditions since no special functional status is attributed, on this model, to the volumetric part boundary and volumetric configuration. However, before discussing the wider implications of the data, we consider one other alternative account of the results. Experiment 2 rules out one version of the summation account based on the priming that might be expected from the combination of partial activation of volumetric shape and volumetric configuration. However, it does not address one other possible version of the summation account, which attributes priming only to the combination of partial activation of volumetric shape information alone. On this version, the equivalent priming between same-part and different-part primes in Experiment 2 can be accounted for by volumetric segmentation theories: In the same-part condition the two surfaces from a single volume provide enough information to derive one complete volume in the whole object, which facilitates its recognition. In contrast, in the different-part condition, primes contain visible surface information from both volumes that equally facilitates recognition, via summation of partial shape priming. However, this version cannot account for the difference in priming in Experiment 1 between the Type 1 visible-surface primes, which did contain partial visible shape information from two volumes, and occluded-surface (volumetric) primes, which did not. Thus, neither of the two versions of the summation hypothesis (partial shape plus volumetric configuration or partial shape alone) provides a viable account of the patterns of priming found across the two studies.

Finally, one might also argue that a genuine effect of volumetric structure is masked in Experiment 2 as a result of the apparent perceptual bistability of the two-surface primes (see Figure 4). In both conditions the primes could be seen as showing two surfaces that connect at either a convex or a concave intersection. Note that only the concave interpretation would potentially prime information about the volumetric configuration. This might be a problem if there was a perceptual bias towards interpreting the stimuli as convex rather than concave (as this might reduce the amount of priming in the different-part condition from volumetric configuration). Current evidence is unclear about whether there is a convexity bias in recognition tasks (e.g., Bertamini, 2001). However, it is relevant to note that even if one assumes such a bias during the processing of the primes this would be inconsistent with the observed pattern of results across the two experiments. In the Type 2 primes of Experiment 1 a convexity bias would benefit both the visible-surface and occluded-surface primes equally. Despite this, there was significantly more priming in the visible-surface condition. Thus, a convexity bias in the interpretation of bistable primes does not provide a viable explanation for the data across the two experiments.

We now discuss the wider implications of the results from both experiments in the General Discussion.

GENERAL DISCUSSION

The main findings from these two experiments can be summarized as follows: (a) Experiment 1 showed that part-whole shape priming is dependent on the match between visible surfaces in the primes and those in the targets. (b) Experiment 2 provided further converging evidence that priming effects are mediated by visible surface structure. This study showed that there is no additional information gain from primes containing visible surfaces that preserve the spatial configuration of volumetric parts over primes containing the same number of visible surfaces from a single volume. These results provide evidence that 3-D object recognition is

primed by visible surface structure, but not by occluded shape structure derived from volumetric completion. This was case when using surface-rendered volumetric objects (Experiment 1) and using contour-based line drawings (Experiment 2), all of which could be readily decomposed into volumetric parts at regions of curvature minima (Hoffman & Richards, 1984).

The data from Experiment 2 show that the priming effects found in Experiment 1 cannot be explained in terms of the summation of the partial priming of volumetric parts and spatial configuration. It is important to note, however, that this finding does not imply that spatial relations are not encoded during perception, nor do they undermine the potential role of the structural description as a model for shape representation in human vision (Winston, 1975). Indeed, there is evidence from numerous studies for the independent encoding of shape information and spatial configuration (e.g., Arguin & Saumier, 2004). The results do, though, suggest that volumetric configuration per se may play no role in shape representation contrary to volumetric theories of recognition (e.g., Barr, 1981; Bergevin & Levine, 1993; Guzman, 1968; Marr & Nishihara, 1978; Zerroug & Nevatia, 1999). The data also provide no evidence that occluded intersecting surfaces derived by volumetric completion contribute to shape priming.

More broadly, the data give further support for nonvolumetric models of 3-D representation (e.g., Fan et al., 1989; Faugeras, 1984; Hoffman & Richards, 1984; Leek et al., 2005; Lowe, 2003). One such model is the surface-based structural description (Leek et al., 2005). In this model 3-D solid objects are represented by planar 2-D bounded polygonal patches that are used to approximate object surfaces. This account predicts that priming should depend on the match between visible surfaces only and surface configurations of the primes and targets, without any contribution from self-occluded volumetric object structure. This model also provides an alternative account of some previous studies reporting data from contour-deletion experiments that have been interpreted within the context of volumetric shape models (e.g., Biederman, 1987; Biederman & Cooper, 1991).

The current findings suggest that in those studies it is the deletion of surfaces from objects, rather than volumetric parts per se, that determine the effects of feature deletion on recognition performance.

Although the surface-based representations model of Leek et al. (2005) emphasizes a key role for surface structure, the data could also potentially be accounted for by other hypotheses about nonvolumetric image primitives, including 2-D polygonal elements derived from relations among NAPs (e.g., Biederman, 1987) or other forms of complex contour-based parts, such as codons (e.g., Hoffman & Richards, 1984). Additionally, it should be noted that we restricted our examination of volumetric models to those that explicitly propose the encoding of 3-D geometric primitives such as generalized cylinders (and their variants) and superquadrics (e.g., Barr, 1981; Bergevin & Levine, 1993; Guzman, 1968; Marr & Nishihara, 1978). In those models there is a clear, albeit implicit, assumption about the geometric completion of 3-D volumes. It is less clear how the current findings relate to some other proposals about volumetric primitives that are based on the symbolic (rather than geometric) descriptions of shape attributes (e.g., Biederman, 1987; Hummel & Biederman, 1992). Further studies will be needed to distinguish among these different forms of representation.

Original manuscript received 4 September 2007

Accepted revision received 22 April 2008

First published online 21 August 2008

REFERENCES

- Arguin, M., & Saumier, D. (2004). Independent processing of parts and of their spatial organisation in complex visual objects. *Psychological Science, 15*, 629–633.
- 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.
- Bertamini, M. (2001). The importance of being convex: An advantage for convexity when judging position. *Perception, 30*, 1295–1310.
- Biederman, I. (1987). Recognition-by-components: A theory of human image interpretation. *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.
- Brooks, R. (1981). Symbolic reasoning among 3-dimensional and 2-dimensional images. *Artificial Intelligence, 17*, 285–349.
- Bülthoff, H. H., & Edelman, S. (1992). Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proceedings of the National Academy of Sciences, USA, 89*, 60–64.
- Edelman, S. (1999). *Representation and recognition in vision*. London: 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. In *Proceedings of the 7th International Conference on Pattern Recognition* (pp. 796–805). North Holland: IEEE Computer Society Press.
- Guzman, A. (1968). Decomposition of a visual scene into three-dimensional bodies. In *AFIPS Proceedings of the Fall Joint Computer Conference* (Vol. 33, pp. 291–304). San Francisco, CA, USA.
- Hoffman, D. D., & Richards, W. A. (1984). Parts of recognition. *Cognition, 18*, 65–96.
- Hummel, J. E. (2001). Complementary solutions to the binding problem in vision: Implications for shape perception and object recognition. *Visual Cognition, 8*, 489–517.
- Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review, 99*, 480–517.
- Leek, E. C., Reppa, I., & Arguin, M. (2005). The structure of 3D object shape representations: Evidence from whole-part matching. *Journal of Experimental Psychology: Human Perception and Performance, 31*, 668–684.
- Lowe, D. G. (2003). 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, B200*, 269–294.
- Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, 2, 1019–1025.
- Vetter, T., Hurlbert, A., & Poggio, T. (1995). View-based models of object recognition: Invariance to imaging transformations. *Cerebral Cortex*, 3, 261–269.
- Winston, P. (1975). Learning structural descriptions from examples. In P. Winston (Ed.), *The psychology of computer vision*. New York: McGraw-Hill.
- 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.