
The influence of cast shadows on visual search

Ronald A Rensink

Departments of Psychology and Computer Science, 2136 West Mall, University of British Columbia, Vancouver, BC V6T 1Z4, Canada; e-mail: rensink@psych.ubc.ca

Patrick Cavanagh

Department of Psychology, Harvard University, 33 Kirkland Street, Cambridge, MA 02138, USA

Received 9 December 2003, in revised form 28 May 2004; published online on 8 December 2004

Abstract. We show that cast shadows can have a significant influence on the speed of visual search. In particular, we find that search based on the shape of a region is affected when the region is darker than the background and corresponds to a shadow formed by lighting from above. Results support the proposal that an early-level system rapidly identifies regions as shadows and then discounts them, making their shapes more difficult to access. Several constraints used by this system are mapped out, including constraints on the luminance and texture of the shadow region, and on the nature of the item casting the shadow. Among other things, this system is found to distinguish between line elements (items containing only edges) and surface elements (items containing visible surfaces), with only the latter deemed capable of casting a shadow.

1 Introduction

Studies of human vision have long distinguished between operations carried out rapidly and in parallel across the visual field, and operations carried out sequentially on an item-by-item basis (Al-Haytham c 1000/1989; Helmholtz 1867/1962; Neisser 1967). In more recent treatments, this division usually takes the form of an early stage that rapidly extracts features from the incoming image, and a subsequent attentional stage that assembles these features into more coherent descriptions of the external world (eg Beck 1982; Julesz 1984).

Although early-level features were originally believed to describe relatively simple aspects of the image [such as color and two-dimensional (2-D) orientation in the plane], later work showed they can also describe more complex aspects of the three-dimensional (3-D) scene, such as surface convexity (Kleffner and Ramachandran 1992; Ramachandran 1988), direction of lighting (Enns and Rensink 1990), and occluded surface structure (Nakayama et al 1995; Rensink and Enns 1998). These results suggest that early vision attempts to extract as much scene structure as possible, and that even when a complete and accurate interpretation is not possible, it is still worth forming a 'quick and dirty' initial estimate (Rensink and Enns 1995, 1998).

The issue examined here is whether cast shadows are also handled by such rapid processes at early levels.⁽¹⁾ Shadows do not correspond to structure in the world itself, arising instead from interactions between the world and the sources that illuminate it (see eg Baxandall 1995). Thus, although shadows can carry information about the scene (eg that a post is upright), there is little need to represent the shadows themselves. Indeed, if the lighting-dependent border of a shadow were mistaken for the lighting-independent border of an object, serious errors in scene interpretation could arise (Braje et al 1998; Cavanagh and Leclerc 1989). To have the best of both worlds—to use

⁽¹⁾The term 'rapid' denotes a process carried out within the first few hundred milliseconds of processing, while 'low-level' denotes processing not involving stimulus-specific knowledge; typically this is carried out in parallel across the visual field. The term 'early' denotes a process that is both rapid and low-level (see Rensink and Enns 1998).

the information contained in shadows but not represent them in the final description of a scene—the visual system must identify regions as shadows as soon as possible and then discount⁽²⁾ them once the relevant information has been extracted (figure 1). But, because a shadow gives rise to an image region that can always be interpreted as something else (eg a black patch on a surface), shadow interpretation is fraught with difficulties—difficulties that are compounded in the absence of knowledge about the scene (Cavanagh 1991). As such, it is not clear from a priori considerations whether shadow interpretation would or would not exist at early levels.

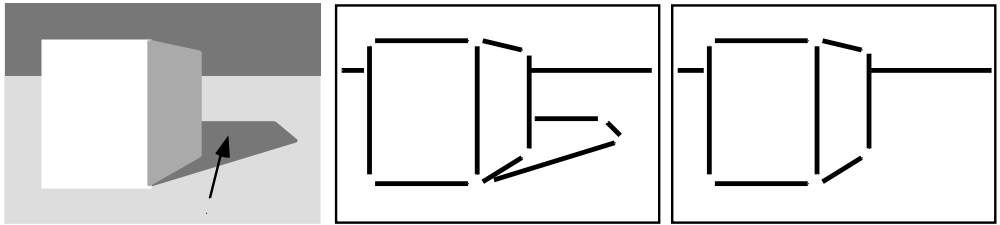


Figure 1. Discounting of shadows. If the shadow is not discounted, irrelevant edges caused by its border can interfere with recognition, for example by leading to the perception of an extraneous part. Discounting can largely reduce the effect of these edges.

A few researchers have begun to examine this issue experimentally. Rensink and Cavanagh (1993) found that an ‘anomalous’ (upside-down) shadow among a set of ‘normal’ (upright) shadows was detected more quickly than a normal shadow among anomalous ones. This was explained in terms of a rapid shadow-interpretation system, with regions corresponding to normal shadows being discounted, and regions corresponding to anomalous shadows left unaffected. Further support for such a system came from Cunningham et al (1996), who showed that cast shadows could be used to rapidly assign 3-D depth to items. In addition, Elder et al (1998) found that cast shadows could be rapidly distinguished from attached ones, and that they could be used to rapidly compute surface relief.

In this paper we look at the issue of rapid shadow interpretation in greater depth, and develop a more robust methodology to test for its existence. This methodology is based on visual-search experiments involving simple items that correspond to objects casting shadows onto a ground plane. The approach is similar to that of Rensink and Cavanagh (1993), but more direct: observers are asked to determine the presence or absence of a target item with a candidate region differing in shape from those of the distractors (figure 2). Performance is then compared to that for the same displays rotated 180° in the image plane. If shadow identification does not exist at early levels, items would be distinguished only by the two-dimensional orientation of the quadrilaterals, and so search would not be affected. But, if rapid shadow interpretation does exist, it may use some of the same assumptions as other rapid-interpretation processes; for example, it may assume that lighting is from above (Enns and Rensink 1990; Kleffner and Ramachandran 1992). If so, interpretation would occur only for those cases where lighting is from above, leading to a difference in speeds. The constraints involved could then be determined by altering various properties to find the point at which this difference appears (or disappears).

In what follows, we show that search is relatively slow for shapes of regions identified as cast shadows created by lighting from above. Results support the proposal of a system that can rapidly identify shadows and then discount them in some way, regardless

⁽²⁾The term ‘discount’ is used in two senses here: (i) the form of the shadow is no longer available, and (ii) this is the result of discounting the illuminant. For present purposes, there is no great reason to distinguish between the two, and so in the interests of simplicity a single term is employed here.

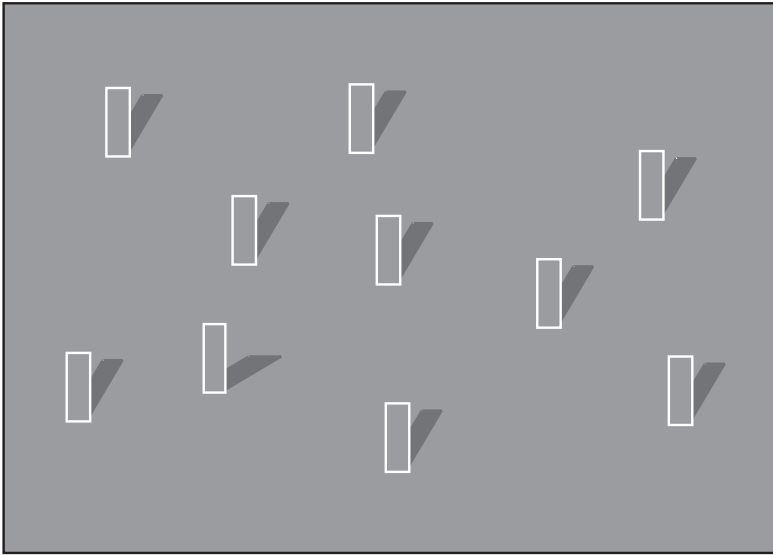


Figure 2. Example of a search display with standard stimuli (ie the stimuli of experiment 1). Figure upright: The dark-gray quadrilaterals attached to the vertically oriented rectangles correspond to shadows cast by vertical posts onto the ground plane. Figure inverted: Quadrilaterals correspond to 'upside-down' shadows on a ceiling plane.

of whether this helps or hinders performance. Several constraints are found to exist on this process, such as the constraint that shadows must be cast by surface elements, ie elements with a measure of opaque surface structure. In addition, we find that the ratios of target-absent to target-present slopes are significantly higher for stimuli corresponding to shadows than for most other search stimuli, suggesting that shadows may also influence the strategic control of the search process.

2 General design

2.1 Methodology

The experiments discussed here were based on the same methodology as for most visual-search experiments: the observer was asked to search a display as rapidly as possible for an item with a unique visual attribute (a target) among a set of other items (distractors). Targets were present on half the trials (chosen randomly) and absent on the remainder. Observers were asked to determine the presence or absence of the target item as quickly as possible, while keeping errors below 10%. Feedback was provided after each trial; average error rate over all conditions was less than 4%. Reaction time (RT) was measured as a function of the number of items present in the display (set size). The primary dependent variable in the analyses was search rate, defined as the slope of the correct RT over set size. (For more detailed discussions of visual-search methodology, see eg Treisman and Gormican 1988.)

2.2 Observers

In each experimental condition we tested ten adult observers with normal or corrected-to-normal visual acuity. About half the observers were naïve to visual-search methodology, while the others had been tested extensively in other search (or related) tasks. All observers completed 4 sets of 60 trials in each condition of each experiment. A different set of observers was generally used in each condition; to reduce the chance of long-term learning effects, any individual observer was run in no more than three conditions. When analyzing the data, the conservative assumption was made that the observers in each condition were independent of those in the others.

2.3 Stimuli

Stimuli were vertically oriented rectangles, each with an attached quadrilateral that could—at least under some conditions—be interpreted as a shadow cast by a vertical post (ie the rectangle) onto a ground plane (figure 2). Rectangles were usually of a uniform color (medium gray) and outlined in white to separate them from the background (also medium gray). Targets were usually items with quadrilaterals extending at an angle of 30° from the horizontal in the image plane; distractors had quadrilaterals at 60° . Areas of all quadrilaterals were adjusted to be approximately equal, so that observers could not base search on their area, but required access to their shape.

Displays consisted of 2, 6, or 10 items. These were positioned randomly on an imaginary 5×4 grid of possible locations, with locations selected such that the density of items was similar for all set sizes. The display area subtended approximately $10 \text{ deg} \times 8 \text{ deg}$ of visual angle, with each item less than 2 deg in size. The position of each item was jittered by $\pm 0.5^\circ$ to minimize possible effects of item collinearity.

Each observer was tested on two counterbalanced cases: *upright*, in which the quadrilaterals corresponded to shadows created via lighting from above (figure 2, upright), and *inverted*, where the displays were rotated by 180° in the image plane so that quadrilaterals corresponded to shadows created via lighting from below (figure 2, upside down).⁽³⁾ Since these differed only by a 180° rotation, image properties were similar in both cases, whereas a difference existed in inferred properties, such as lighting direction and background orientation (ground plane versus ceiling).

2.4 Analysis

Analysis was based on differences in search speeds for the upright and inverted cases of a given experimental condition. This was determined via paired *t*-tests (two-tailed) on the target-present slopes for each observer.⁽⁴⁾ In all conditions, the error data were consistent with the RT slope data, ruling out the possibility of speed–accuracy tradeoffs. A complete listing of the slopes, baselines, and error rates for the various conditions is given in the Appendix.

3 Experiment 1: Basic effect

In experiment 1, items were based on variants of a set of standard stimuli (figure 3). For these, targets and distractors both contained gray quadrilaterals that were darker than the background; these had sharp borders, corresponding to ‘hard’ shadows. Targets were distinguished by a unique orientation of the quadrilaterals. Search slopes were then compared for upright and inverted conditions.

In condition 1A we used the standard stimuli, which could be interpreted as posts casting shadows onto a ground plane (figure 2). Results are shown in figure 3. Speed in the upright case was 13.4 ms/item for target-present trials, and 40.3 ms/item for target-absent ones. Meanwhile, search in the inverted case was reliably faster: 4.8 ms/item for

⁽³⁾The experimental protocol began by showing the observer a picture of upright stimuli (ie with the quadrilaterals attached to the bottoms of the rectangles, so that the quadrilaterals might be interpreted as shadows caused by lighting from above). Observers were then told that they would be looking for a shadow that “stuck out a little bit”. For the cases where the inverted stimuli were tested, observers were told that in this condition they would see shadows that were upside down—ie the rectangles are attached to the ceiling and lit from below.

⁽⁴⁾In almost all cases, the same pattern of results was obtained for *t*-tests based on combined rates, ie rates obtained by the weighted average of the target-present and target-absent slopes, with the target-absent slopes of each observer in a particular case (upright or inverted) divided by the ratio of target-absent to target-present slopes for that case. The use of combined rates can partly compensate for the greater variability of target-absent responses (Wolfe 1998) while still involving the data from these responses in the analyses. The only condition where the two measures differed was condition 2B, in which statistical significance was achieved when using the combined rate but not the target-present rate.

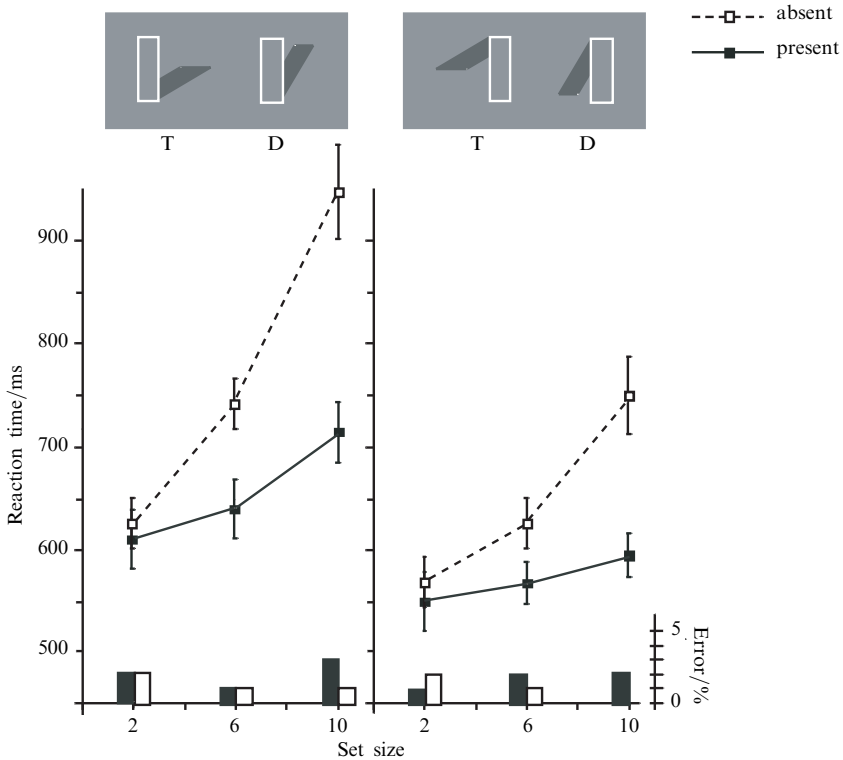


Figure 3. Results of condition 1A (the standard stimuli). Search for the unique orientation of the quadrilateral in the upright case is slower than in the inverted case. Error bars indicate standard errors of the mean. T = target; D = distractor.

target-present trials, and 22.4 ms/item for target-absent ones. This suggests that rapid shadow interpretation occurred in the upright case, with search slowed down because the shapes of regions identified as shadows were relatively difficult to access.

To show that this slowdown was not somehow due to the shape of the items, in condition 1B the intensities were reversed so that light regions were dark, and dark regions light (figure 4). Slowdown now disappeared, with speed for upright (5.1 and 12.1 ms/item) not reliably different than for inverted (5.1 and 11.4 ms/item) cases. Interestingly, these speeds were similar to those of the inverted case of condition 1A, indicating a search based on distinctive image features in all three cases. These results are consistent with a *darkness constraint*: regions are considered to be shadows only if they (or at least their borders) are darker than the surrounding background. Regions not interpreted as shadows are not discounted, their shapes remaining available as the basis of rapid search.

Although a cast shadow in the scene must correspond to a dark region in the image, the converse does not hold—a dark region need not correspond to a shadow. For example, it could correspond to a surface of low reflectance, or to a non-Lambertian surface oriented obliquely to the direction of lighting (Baxandall 1995). To test whether it was simply the dark quadrilaterals in condition 1A that somehow caused slowdown, in condition 1C we used the quadrilaterals alone, with accompanying rectangles. Slowdown now disappeared, with search for upright (2.9 and 8.2 ms/item) not reliably different from inverted (3.2 and 11.9 ms/item) cases. This suggests that the rapid identification of a shadow needs a nearby item to correspond to the shadowcaster—the structure in the world that casts the shadow.⁽⁵⁾

⁽⁵⁾ Also called 'obtruder' (eg Casati 2000) and 'casting object' (eg Mamassian et al 1998).

















Condition	Target	Distractor	Speed/ms per item		<i>p</i>
			upright	inverted	
1A			13.4 40.3	4.8 22.4	0.003
1B			5.1 12.1	5.1 11.4	0.99
1C			1.9 7.7	1.4 11.5	0.90
1D			5.5 11.2	4.7 10.2	0.65
1E			3.7 10.8	2.4 7.5	0.55
1F			2.4 12.6	2.1 10.0	0.80
1G			3.2 15.4	3.0 17.4	0.95
1H			14.3 42.9	5.5 22.0	0.003

Figure 4. Stimuli and results for the conditions of experiment 1. Target is defined by a difference in the orientation of the quadrilateral attached to the rectangle. Numbers with dark-gray backgrounds indicate a reliable difference in speeds between upright and inverted cases.

To examine whether simply obeying this *shadowcaster constraint* is sufficient for a dark patch to cause slowdown, in condition 1D the shadowcasters were restored but the quadrilaterals were surrounded with white borders. If the effects are due to a system that simply handles dark patches, slowdown should remain. But, if they are due

to a system specialized for shadows, slowdown will likely be eliminated, since white borders do not generally surround shadows. Results show that slowdown was eliminated, with no reliable differences in speeds found for upright (5.5 and 11.2 ms/item) and inverted (4.7 and 10.2 ms/item) cases. Slowdown therefore appears to exist only in conditions compatible with the identification of a region as a shadow. Moreover, the lack of slowdown found here suggests the darkness constraint requires the border of the candidate region to be darker than the background; simply having a darker interior is clearly insufficient.

In condition 1E we examined whether slowdown would occur if borders were black rather than white (figure 4). Results show that speeds for upright (3.7 and 10.8 ms/item) did not differ reliably from those for inverted (2.4 and 7.5 ms/item) cases. Evidently, it is not enough for the border of the candidate region to be darker than the background—there must also be no sudden decrease in luminance as the border is approached. This is consistent with reports based on subjective impression, where a dark surrounding line will prevent the description of a region as a shadow (Hering 1874/1964).

Looked at more closely, there are two different factors which must be disentangled here: (i) a decrease in luminance near the border of the region, and (ii) a sudden change in luminance. To examine the first of these, in condition 1F we used stimuli with a gradual darkening towards the borders of the quadrilateral (figure 4). Slowdown here disappeared, with no difference in slopes for upright (2.4 and 12.6 ms/item) and inverted (2.1 and 10.0 ms/item) cases. As such, the process appears to involve a *monotonicity constraint*: the luminance of the region must be monotonically nonincreasing (ie must not increase) as the border is approached. This is a natural requirement on a shadow-interpretation system.⁽⁶⁾

Condition 1G was designed to test the second component—the sharpness of the luminance change. Here, a narrow black region was inserted inside each quadrilateral. The monotonicity constraint was therefore obeyed, but a sharp change still existed (figure 4). No reliable difference in speed between upright (3.2 and 15.4 ms/item) and inverted (3.0 and 17.4 ms/item) cases was found. Thus, there also appears to be an *internal smoothness constraint*: any sharp luminance edge within a region is sufficient to destroy its interpretation as a shadow.

To determine if the identification of a region as a shadow requires a completely uniform luminance, in condition 1H the edges of the quadrilaterals were blurred so that darkness, monotonicity, and internal-smoothness constraints were all obeyed, while still having a nonuniform luminance (figure 4). Slowdown was now found, with search for upright (14.3 and 42.9 ms/item) being reliably slower than for inverted (5.5 and 22.0 ms/item) cases. Such behavior would be expected of a shadow-interpretation system: blurred edges correspond to ‘soft’ shadows formed by diffuse or extended light sources (eg Baxandall 1995).

Taken together, the results of experiment 1 show a slowdown exists in search for dark regions that can be interpreted as cast shadows created by lighting from above. This slowdown was found only for those conditions that simultaneously obeyed the darkness, shadowcaster, monotonicity, and internal-smoothness constraints. This pattern of results is difficult to explain in terms of simple image properties: why should a 180° rotation in the image plane create such a difference in search speeds, and only for particular conditions? But it is compatible with the quadrilaterals being rapidly identified as shadows and then discounted in some way.

⁽⁶⁾Regions in which luminance decreases as the border is approached can be found in shadows cast by transparent body with nonparallel surfaces. However, such bodies are relatively rare in nature, being limited to things such as water drops or waves on a lake.

4 Experiment 2: Conjunctions

To further investigate the nature of the search slowdown, in experiment 2 we used an alternative approach in which the target was defined by a feature conjunction instead of a feature difference. In condition 2A, the target had a 30° quadrilateral with white borders, while distractors were of two kinds: (i) 60° quadrilaterals with white borders, and (ii) 30° quadrilaterals without white borders. Targets were therefore defined by a conjunction of nonshadow status plus orientation of 30° (figure 5).










Condition	Target	Distractor 1	Distractor 2	Speed/ms per item		<i>p</i>
				upright	inverted	
2A				30.4	41.4	0.03
				68.2	96.0	
2B				54.5	47.3	0.15
				109.3	87.1	
2C				32.5	20.6	0.015
				53.1	38.2	

Figure 5. Stimuli and results for the conditions of experiment 2. Target is defined by a conjunction of properties contained in the distractors. Numbers with dark-gray backgrounds indicate a reliable slowdown in the upright case. Numbers with light-gray backgrounds indicate a reliable speedup in the upright case.

If shadows are rapidly identified and discounted, search for upright targets should be relatively fast, since one of the target properties (the 30° orientation) is now less accessible in the distractors and so should interfere less. Consequently, discounting should cause a speedup rather than a slowdown. Results show that such a speedup does indeed occur: search was reliably faster for upright (30.4 and 68.2 ms/item) than for inverted (41.4 and 96.0 ms/item) cases. Consequently, the slowdown found in experiment 1 is not due to a general slowing down whenever items identified as shadows are processed. Instead, it likely stems from a difficulty in accessing the discounted regions.

In condition 2B we examined a variant in which the shadow status of the items in condition 2A was switched, by switching the assignment of white borders (figure 5). The target is now defined by a conjunction of shadow status plus orientation of 30°. Targets are therefore distinguished by a property of their shadows, and so all shadow items must be checked carefully. This should again result in relatively slow search for the upright case. Indeed, search for upright (54.5 and 109.3 ms/item) did tend to be slower than for inverted (47.3 and 87.1 ms/item) cases, although this difference was not reliable. Note that the speeds found in this condition are much slower than in experiment 1, making it unlikely that observers could have focused on the shadows alone.

To examine whether reliable slowdown could be achieved with different stimuli, in condition 2C we replaced the distractors containing the 30° quadrilaterals with isolated rectangles at a 30° orientation in the image plane (figure 5); these violate the

shadowcaster constraint, and so are also not interpreted as shadows. Search for upright (32.5 and 53.1 ms/item) was now reliably slower than for inverted (20.6 and 38.2 ms/item) cases, showing that slowdown can also occur with this kind of test.

5 Experiment 3: Textured regions

Experiments 1 and 2 indicate that rapid shadow interpretation does exist, with the region identified as a shadow being discounted, presumably after any relevant information has been extracted. By comparing conditions for which the associated search slowdown does and does not occur, it becomes possible to map out the sensitivity of this system to various properties of the stimuli, and so determine the various constraints involved in its operation.

In experiment 3 this approach was used to examine sensitivity to texture. In condition 3A, a set of thin white stripes was placed on the quadrilaterals of the standard stimuli so that they were textured (figure 6). Slowdown now disappeared: search for upright (4.9 and 18.6 ms/item) was not reliably different than for inverted (9.2 and 19.3 ms/item) cases. This suggests that regions with stripes are not interpreted as shadows.

This effect, however, may have occurred because observers could attend to parts of the items (by accentuating the quadrilaterals or inhibiting the rectangles), making the quadrilaterals the only effective items in the display. Such guidance can be done on the basis of color (Treisman and Gelade 1980; Wolfe et al 1990), and it may have also been achieved here on the basis of texture. To test this possibility, in condition 3B we used items with the same texture in both the quadrilaterals and in the rectangles (figure 6), so that the quadrilaterals could not be simply isolated via their texture. Search did not differ reliably for upright (4.1 and 23.9 ms/item) and inverted (5.6 and 20.2 ms/item) cases, supporting the proposal that the lack of slowdown in condition 3A was indeed due to a lack of shadow interpretation.

To determine if the presence of texture in the quadrilaterals is the key factor, in condition 3C stripes were placed on the background only, leaving the quadrilaterals as solid dark regions. Slowdown again disappeared: slopes for upright (4.3 and 16.6 ms/item) did not reliably differ from those for inverted (2.6 and 9.5 ms/item) cases. For these stimuli, the dark patch may have been interpreted as a solid structure that obscured the texture of the background. More generally, the key factor may simply have been the presence of a texture boundary aligned with the border of the region.

To test this latter possibility, in condition 3D stripes were placed in both the quadrilaterals *and* the background, such that the stripes were unbroken (figure 6). Slowdown now returned, with search for upright (12.5 and 34.9 ms/item) reliably different than for inverted (4.4 and 18.1 ms/item) cases. This indicates that the key factor is the alignment of a texture boundary with the border, and not the simple presence of texture per se. If such a boundary exists, the region is treated as a structure in the scene; otherwise, the texture is seen as extending through the background, with the region interpreted as a shadow falling upon it.⁽⁷⁾

To determine if these effects were due to virtual lines formed by the terminators of the stripes, or to contrast reversals along the texture boundary (Cavanagh and Leclerc 1989), similar arrangements were used in conditions 3E–3H with a texture formed by white stipples, which had minimal contact with the borders of the regions (figure 6). Condition 3E corresponded to condition 3A, where the texture was restricted to the quadrilaterals. Consistent with condition 3A, speed for upright (13.9 and 39.6 ms/item) was no slower than for inverted (12.0 and 35.9 ms/item) cases.

⁽⁷⁾The internal smoothness constraint is not violated by the sharp edges in these conditions, since the edges belong to the background rather than to the quadrilateral, and so are not internal.

















Condition	Target	Distractor	Speed/ms per item		<i>p</i>
			upright	inverted	
3A			4.9 18.6	9.2 19.3	0.20
3B			4.1 23.9	5.6 20.2	0.50
3C			4.3 16.6	2.6 9.5	0.25
3D			12.5 34.9	4.4 18.1	0.025
3E			13.9 39.6	12.0 35.9	0.25
3F			7.4 19.9	7.2 14.5	0.95
3G			3.5 19.6	4.4 13.3	0.60
3H			12.0 34.3	6.5 22.2	0.025

Figure 6. Stimuli and results for the conditions of experiment 3. Target is defined by a difference in the orientation of the quadrilateral attached to the rectangle. Numbers with dark-gray backgrounds indicate a reliable difference in speeds between upright and inverted cases.

Possible guidance via the textures of the regions was again tested in condition 3F, where both the rectangles and quadrilaterals were filled with stippled texture. No reliable difference was found between upright (7.4 and 19.9 ms/item) and inverted (7.2 and 14.5 ms/item) cases, providing additional evidence that the disappearance of slowdown was not because of texture-based guidance.

Condition 3G paralleled condition 3C in that the background alone was textured. As before, slowdown disappeared: speed for upright (3.5 and 19.6 ms/item) was not reliably different than for inverted (4.4 and 13.3 ms/item) cases.

Finally, in condition 3H a stippled texture was used that ran through both the background and the quadrilaterals. As in condition 3D, slowdown now returned: search for upright (12.0 and 34.3 ms/item) was reliably slower than for inverted (6.5 and 22.2 ms/item) cases. Consequently, both kinds of textures yield the same pattern of results: a dark region will not be interpreted as a shadow if a texture boundary is aligned with its border.

6 Experiment 4: Textured lines

Experiment 4 was designed to investigate whether a texture boundary along the border of a dark region will prevent shadow interpretation even for a textured (or patterned) line, with no textures in the region or its surround. In condition 4A dashed lines were placed along the borders of the quadrilaterals (figure 7). Slopes for upright (6.9 and 17.2 ms/item) did not reliably differ from those for inverted (7.7 and 18.8 ms/item) cases. As such, a textured line seems sufficient to prevent shadow interpretation.









Condition	Target	Distractor	Speed/ms per item		<i>p</i>
			upright	inverted	
4A			6.9 17.2	7.7 18.8	0.75
4B			3.2 17.2	3.0 17.1	0.95
4C			6.5 18.7	6.1 11.8	0.85
4D			4.1 22.3	1.5 17.1	0.30

Figure 7. Stimuli and results for the conditions of experiment 4. Target is defined by a difference in the orientation of the quadrilateral attached to the rectangle. Numbers with dark-gray backgrounds indicate a reliable difference in speeds between upright and inverted cases.

This was further investigated in condition 4B by placing two white dots along the border at the far corners of the quadrilateral (figure 7). Again, slopes for upright (3.2 and 17.2 ms/item) did not differ from those for inverted (3.0 and 17.1 ms/item) cases.

One of the dots was removed in condition 4C, so that only one dot remained at the far corner. Slowdown again disappeared: speed for the upright case (6.5 and 18.7 ms/item) was not reliably different from that of the inverted case (6.1 and 11.8 ms/item).

Condition 4D was designed to test whether the precise position of the dot along the border was important, by placing the dot midway along the far luminance boundary

(figure 7). Although a slight indication of a slowdown was found, it was still effectively absent: slopes for upright (4.1 and 22.3 ms/item) did not reliably differ from those for inverted (1.5 and 17.1 ms/item) cases. Evidently, even a single dot located anywhere along a luminance boundary can—at least to some degree—prevent rapid shadow interpretation.

The results of experiments 3 and 4 are therefore consistent with the proposal that the rapid interpretation of a region as a shadow requires that it obey the constraints found in experiment 1 (eg the darkness constraint), as well as a *nonalignment constraint*: no texture boundaries or small markings can be aligned with its border.

7 Experiment 5: Shadowcasters

Another interesting issue concerns the structure of the shadowcaster. The standard stimuli contain a rectangle formed of a medium gray region surrounded by a white line (figure 2). Since the background is also medium gray, this could be interpreted either as a surface element (an element with an opaque surface assigned to it) or a line element (an element with edges but no assigned surface structure, such as a wire loop). Given the need for a shadowcaster in rapid interpretation, it is worth investigating whether it must also be a surface element, or whether being a line element would suffice.⁽⁸⁾

Condition 5A featured a gap in the side of the rectangle (figure 8), so that it now became a line element—a C-shaped frame. No reliable difference was found between speeds for upright (6.0 and 18.4 ms/item) and inverted (5.1 and 16.1 ms/item) cases. This indicates that a distinction between line and surface elements does exist, and has consequences for shadow interpretation. In particular, it suggests that a region is identified as a shadow only if it obeys a *surface-element constraint*: the item corresponding to the shadowcaster must be a surface element.

The generality of this constraint was tested in condition 5B by moving the gap to the top of the rectangle. Again, no reliable difference was found between upright (6.7 and 24.6 ms/item), and inverted (5.5 and 17.9 ms/item) cases, indicating that the location of the gap had no effect.

To test whether simple continuity of the outline might be the critical factor, in condition 5C the rectangles were partially filled in, so that each stimulus corresponded to a shadow as seen through a rectangular frame (figure 8). Although a hint of slowdown did appear, it nevertheless remained largely absent: no reliable difference in speeds was found between upright (6.2 and 25.0 ms/item), and inverted (4.2 and 8.3 ms/item) cases.

In condition 5D the rectangle was replaced in each item by a simple vertical line. Again, no reliable difference in search speed was found between upright (1.9 and 14.7 ms/item) and inverted (3.7 and 9.6 ms/item) cases. Taken together then, conditions 5A–5D provide considerable support for a surface-element constraint.

Condition 5E was designed to investigate whether a line element could become a surface element if its interior was filled with a texture (figure 8). Results show only slight evidence of this: search for upright (8.8 and 29.9 ms/item) was marginally slower than for inverted (4.0 and 17.0 ms/item) cases, but this difference was still not reliable. As such, textural filling does not seem to be very effective at conferring surface-element status to the item, at least with the density of texture elements used here.

The object of condition 5F was to investigate whether the effect of gaps could be ‘repaired’ by the use of grouping to form a virtual line (see Rensink and Enns 1995).

⁽⁸⁾ A line element corresponds to a structure without a noticeable surface or volume to it, eg a wire or a thin branch. Since virtually all lighting sources have a finite extent, penumbral effects would wash out any shadows cast by such structures. A stimulus with sufficient thickness would eventually be perceived as having a surface structure; whether this corresponds to the point at which it would also cast a noticeable shadow in the real world is an open issue. In any event, the lines used in these experiments are sufficiently thin that their shadows would likely be washed out.

















Condition	Target	Distractor	Speed/ms per item		<i>p</i>
			upright	inverted	
5A			6.0 18.4	5.1 16.1	0.60
5B			6.7 24.6	5.5 17.9	0.45
5C			6.2 25.0	4.2 8.3	0.40
5D			1.9 14.7	3.7 9.6	0.45
5E			8.8 33.8	4.0 17.0	0.20
5F			3.3 22.7	5.7 15.2	0.30
5G			3.8 20.7	3.1 12.4	0.80
5H			6.1 47.2	4.5 22.5	0.55

Figure 8. Stimuli and results for the conditions of experiment 5. Target is defined by a difference in the orientation of the quadrilateral attached to the rectangle. Numbers with dark-gray backgrounds indicate a reliable difference in speeds between upright and inverted cases.

Here, gaps were placed in the top and bottom of the outline, creating a pair of brackets that could be easily grouped (figure 8). However, search for upright (3.3 and 22.7 ms/item) was no different than for inverted (5.7 and 15.2 ms/item) cases.

In condition 5G we tested whether status as a surface element could be restored if the gap was part of a virtual line formed by a series of dots (figure 8). Again, no reliable difference was found between upright (3.8 and 20.7 ms/item) and inverted (3.1 and 12.4 ms/item) cases.

Finally, in condition 5H we examined whether virtual lines formed by texture boundaries were sufficient to establish the shadowcaster as a surface element. Given the importance of these boundaries in the alignment constraint (experiment 3), it may be that they can be used in place of luminance edges. However, no reliable difference in speed was found between the upright (6.1 and 47.2 ms/item) and inverted (4.5 and 22.5 ms/item) cases.

The results of this experiment therefore indicate the existence of at least two distinct kinds of structure at early levels: *line elements* (with edges but no surface structure) and *surface elements* (with an opaque surface structure). Rapid shadow interpretation is possible only when the shadowcaster is a surface element. Note that this is not simply a matter of the narrowness of the lines in the item: only a slight change of this stimulus (such as the filling-in of a small gap) is enough to create a strong slowdown. Surface elements appear to require the presence of a continuous luminance edge along their border—virtual contours and grouping do not appear to be sufficient. But, although such continuity may be necessary, it is not sufficient: if evidence exists that the interior is not filled, shadow interpretation will still fail (condition 5C).

8 General discussion

The experiments described here show that visual search is influenced by the presence of cast shadows on a ground plane. In particular, search can be slowed down (experiment 1) or speeded up (experiment 2) relative to that for simple 2-D properties of the image. Results are consistent with a difficulty in accessing the shapes of regions that correspond to cast shadows. Since slowdown can occur for conditions that would otherwise lead to fast search (on the order of a few milliseconds per item), the process responsible would appear to act rapidly and in parallel—characteristics generally taken to indicate processing at early levels of vision (see Rensink and Enns 1995, 1998).

These results support the existence of a process that can rapidly identify regions as shadows and then discount them to some extent. A region not interpreted as a shadow is not discounted, and so remains available for rapid search. Note that such discounting may not be complete—the magnitudes of the slowdowns are not what might be expected were the regions identified as shadows completely inaccessible to search. Such ‘partial’ discounting may explain why cast shadows can occasionally interfere with recognition (Braje et al 1998).

In any event, slowdown disappears if any texture boundary is aligned with the border of the region, or if the region is not near an item with a visible surface. As such, it supports a clear distinction between the rapid handling of shadows and two related processes: the rapid handling of reflectance (eg Mitsudo 2002), and the rapid handling of shading (eg Kleffner and Ramachandran 1992), neither of which have—nor should have—these particular constraints.

8.1 Constraints on rapid interpretation

The pattern of slowdown appearance/disappearance found here casts light on the set of constraints that enter into the treatment of a given region as a shadow. This treatment involves two distinct tasks: (i) identification of a region as a shadow, and (ii) discounting of the identified region. It is worth pointing out that each task may involve a different set of constraints, and perhaps even a different system. If so, the results found here would be based on the joint action of these two systems, and the constraints would be the union of the constraints for the individual tasks. Determining whether such a

separation of tasks does exist, and—if so—which constraints apply to which task must await future work.

As expected of any rapid-interpretation process (Rensink and Enns 1995, 1998), the constraints found here are simple and capable of being verified rapidly and in parallel at each location in the image:

(a) *Darkness constraint*: the border of the candidate region must be darker than its surroundings (conditions 1B, 1D).

(b) *Monotonicity constraint*: the luminance of the region must be monotonically non-increasing as the border of the region is approached (conditions 1E, 1F).

(c) *Internal smoothness constraint*: the interior of the region cannot contain a sharp change in luminance (conditions 1E, 1G).

(d) *Nonalignment constraint*: the border of the region cannot be aligned with a pattern boundary (experiment 3), or fall on a small marking in the image (experiment 4).

(e) *Shadowcaster constraint*: there must exist a nearby element deemed to be the shadowcaster (condition 1C).

(f) *Surface-element constraint*: the item corresponding to the shadowcaster must be a surface element, not a line element (experiment 5).

Without knowledge of the particular objects in the scene, it is impossible to determine whether a region in the image corresponds to a shadow, a colored patch, or an oblique surface (Cavanagh 1991). Thus, the rapid interpretation of a region as a shadow can never be guaranteed to be accurate. But the constraints found here appear to minimize the probability of both kinds of error—failing to identify a shadow as a shadow, and incorrectly identifying a nonshadow as a shadow—while still being able to identify cast shadows under a large variety of conditions.

To see this, consider first the darkness and monotonicity constraints. Together, they imply that the border and interior of the region must be darker than its surroundings, a general constraint that is (at least for opaque bodies) well-grounded in the physics of shadow formation (Baxandall 1995).

The internal-smoothness constraint stems from a slightly different source: the fact that light sources in the natural environment are not structured, and that the medium of transmission (ie the atmosphere) does not create sharp changes in illumination. Consequently, any sharp change must be due to a change in surface reflectance (Horn 1974). The question is then whether this reflectance pattern should be assigned to the background or to the candidate region. The results here indicate that, if the pattern extends beyond the region, it is assigned to the background (conditions 4D and 4H); otherwise, it is assigned to the region, which then rules out its interpretation as a shadow (condition 1G). Interestingly, the subjective impression of a shadow is destroyed if lines inside are parallel to its border (Bühler 1922); it may be that the same constraints are at play here.

The nonalignment constraint is based on yet a different consideration, viz the small likelihood that a shadow border would fall exactly on a texture boundary, or on a small element. In the case of texture, it is much more likely that such an alignment corresponds to a border between two different surfaces than to a shadow border just happening to fall on a texture boundary. Similarly, it is unlikely that a shadow border would just happen to fall on the only small element on the ground nearby. Interestingly, this constraint does not apply to the use of texture for 3-D shape, being replaced by a more liberal one of no contrast reversals along the border (Cavanagh and Leclerc 1989). It may be that the more conservative nonalignment constraint is limited to the discounting aspect of shadow handling. Such an approach would be sensible: if information is rendered inaccessible to higher-level processes, this operation will be difficult to undo, and so should not be carried out if the interpretation is at all doubtful.

The shadowcaster constraint is another requirement based on the physics of shadow formation—a cast shadow requires that something cast it. The surface-element constraint is similarly based on physical considerations: since light sources generally have some extension, they will generally give rise to penumbral effects that would wash out the thin shadow a line element would generate. Note that the distinction between line and surface elements indicates that early vision is not exclusively concerned with surfaces, as has sometimes been proposed (eg He and Nakayama 1992).

Finally, it is worth mentioning that, although the requirement of lighting from above was a motivation for the design of the experiments described here, caution must be applied in asserting that it is a constraint. A 180° rotation was the main manipulation in all experiments. However, it is not clear exactly which factors were involved with this. Lighting direction is certainly one possibility: lighting in the natural world is usually from above (Lynch and Livingston 1995), and a similar regularity also exists in artificial environments. But this constraint does not appear to apply to all types of shadow perception (eg Ostrovsky et al 2001). Furthermore, other possibilities exist, such as the constraint that items must be on a ground plane (a plane viewed from above), or that they not be on a ceiling (a plane viewed from below). Further work is needed here.⁽⁹⁾

8.2 *Locus of processing*

The existence of a process that impedes rapid access to low-level features is also found in several other aspects of rapid vision, such as rapid grouping, in which the formation of a group impedes access to constituents that support rapid search on their own (Rensink and Enns 1995), and rapid completion, in which completion results in a loss of distinctive features that would otherwise support rapid search (Rensink and Enns 1998). This latter process requires approximately 250 ms (Rauschenberger and Yantis 2001). This suggests that it is the result of a *secondary* stage of rapid processing that elaborates a *primary* stage based on an initial set of simple measurements (Rensink 2000).⁽¹⁰⁾ Such elaboration could be carried out in parallel across the visual field via recurrent connections believed to exist between low-level and higher-level representations (cf DiLollo et al 2000). At the very least, some kind of higher-level involvement is required—there are simply too many of these processes (completion, recovery of scene-based properties, grouping, etc) for them all to be implemented by dedicated hardware at low levels.

Rapid shadow interpretation is likely similar, being carried out via recurrent processes in which information circulates between lower levels that hold input from the retina and higher levels that hold knowledge about constraints. Baselines for almost all conditions (including those where shadow interpretation occurs) are around 600 ms (Appendix), consistent with the proposal that this process is rapid. Timing manipulations such as those done on rapid completion (Rauschenberger and Yantis 2001) would likely affect rapid shadow interpretation in a similar way.

8.3 *Strategic control*

The discounting of regions identified as shadows accounts for the pattern of search speeds found in these experiments. Interestingly, it may also account for another pattern that was found: in each of the four conditions where slowdown was found, the ratio of target-absent to target-present slopes was significantly higher than the ratio

⁽⁹⁾ Because the experiments described here rely on the assumption of a ground plane, this possibility cannot be directly tested here. One way of doing so might be via a set of experiments involving a different background structure, such as a frontoparallel plane. A complication here would be to ensure that a different set of constraints is not brought into play for such conditions.

⁽¹⁰⁾ This distinction is somewhat similar to that between the raw and the completed primal sketch proposed by Marr (1982).

of 2.26 typical of most search tasks (Wolfe 1998). This ratio—averaging about 3.5—is consistent with a prolongation of search that compensates for the reduced accessibility of the discounted regions. As such, the higher-level strategic system that controls search may be sensitive to the difficulties created by the discounting of cast shadows.

Interestingly, high search ratios were found even for the inverted cases in these conditions, where search is quite rapid. Indeed, search ratios in all four conditions were reliably higher for the inverted case ($p < 0.02$), suggesting that there may be a cost to the presence of shadows in these displays even when they do not obey all the constraints for rapid analysis. If so, then the constraints relevant for strategic control may differ from those for rapid interpretation.

8.4 *Generality of results*

Although the studies here provide evidence for the rapid interpretation of shadows, it is important to realize that they are based on a restricted set of stimuli. In particular, they are based on stimuli much more impoverished than what is generally encountered in the real world. Consequently, different behavior may apply when more realistic stimuli are used—for example interpretation may succeed for shadows on a ceiling, provided that a sufficient number of cues exist in the scene (eg Ostrovsky et al 2001).

This does not, however, imply that the constraints encountered here are irrelevant. The use of severely impoverished stimuli has the advantage that because cues are sparse, individual constraints can be more easily isolated. As such, the constraints found here almost certainly form much of the basis for the perception of shadows in more realistic stimuli, where these constraints would likely interact with each other and with any positive cues present to yield a final determination of shadow status. The issue of how this might be done is an interesting topic for future work.

It is also worth pointing out that the experiments here were based on the background being interpreted as a ground plane. Different constraints may apply to frontoparallel planes. Furthermore, cast shadows are only one kind of shadow—others also exist, and these may involve other sets of constraints (see eg Baxandall 1995; Cavanagh and Leclerc 1989). It is difficult to imagine why some of the constraints found here (eg nonalignment) would depend on the particular orientation of the background or type of shadow, but such dependence is nevertheless possible. Again, this is another interesting topic for future work.

8.5 *Other systems*

The experiments described here have yielded a particular set of constraints on the rapid identification and discounting of shadows. However, caution must be applied when attempting to relate these constraints to those found in other studies, or to the issue of shadow perception in general.

To begin with, shadow interpretation is not limited to early levels, but appears to play a role in several other higher-level aspects of perception, including phenomenological feel (eg Bühler 1922; Hering 1874/1964; Kardos 1934). Given that different requirements exist at different levels of processing, and that different computational resources are available at each, it is not unreasonable to suppose that different systems may be involved, each with a different mode of operation (cf Zucker 1987).

Furthermore, different systems (or at least, different constraints) may be required for different tasks. Although identification and discounting are the two shadow-related tasks considered here, others exist as well. For example, shadows can be used to establish 3-D surface shape, and to determine depth ordering between items (Baxandall 1995; Cavanagh and Leclerc 1989; Kriegman and Belhumeur 2001; Mamassian et al 1998). Contradictions can appear if a single shadow-interpretation system is assumed: for example, shadows must be dark when used to extract 3-D shape (Cavanagh and Leclerc 1989), but can be white when used for motion in depth (Kersten et al 1997). However, such

contradictions can be resolved if different tasks involve different constraints. Similar considerations may also explain why relatively conservative constraints of the rapid system found here differ from those of the strategic system, and from those for the extraction of 3-D shape from shadow (Cavanagh and Leclerc 1989).

8.6 Final remarks

Much remains to be learned about shadow perception. But as the studies described here have shown, a powerful way of improving our understanding about this aspect of vision is to examine how shadows affect visual search. New methodological markers of shadow perception are important in the long run, for shadows are of interest not only in their own right, but also provide a way to further our understanding of mechanisms used in other aspects of vision (eg the distinction between line and surface elements encountered here). To the extent that such approaches continue to be developed, new insights on human vision will continue to emerge from out of the shadows.

Acknowledgments. Much of this work was done while the first author was at Cambridge Basic Research, Nissan Research & Development, Inc., Cambridge, MA, USA. We would like to thank Janis Breeze, Jennifer Corbett, Carolyn Harp, Kristen Macuga, and Andrew Synyshyn for their help in running the experiments. We would also like to express our gratitude for the support provided by Nissan Motor Co., Ltd (Japan), and the Natural Sciences and Engineering Research Council (Canada).

References

- Al-Haytham I, 1000/1989 *The Optics of Ibn Al-Haytham* Books I–III, translated by A I Sabra (London: The Warburg Institute)
- Baxandall M, 1995 *Shadows and Enlightenment* (New Haven, CT: Yale University Press) pp 1–15
- Beck J, 1982 “Textural segmentation”, in *Organization and Representation in Perception* Ed. J Beck (Hillsdale, NJ: Lawrence Erlbaum Associates) pp 285–317
- Braje W L, Kersten D, Tarr M J, Troje N F, 1998 “Illumination effects in face recognition” *Psychobiology* **26** 371–380
- Bühler K, 1922 *Die Erscheinungsweisen der Farben* (Jena: Fischer)
- Casati R, 2000 “The structure of shadows”, in *Time and Motion of Socio-Economic Units* Eds A Frank, J Raper, J P Cheylan (London: Taylor & Francis) pp 99–109
- Cavanagh P, 1991 “What’s up in top–down processing?”, in *Representations of Vision: Trends and Tacit Assumptions in Vision Research* Ed. A Gorea (Cambridge: Cambridge University Press) pp 295–304
- Cavanagh P, Leclerc Y G, 1989 “Shape from shadows” *Journal of Experimental Psychology: Human Perception & Performance* **15** 3–27
- Cunningham R K, Beck J, Mingolla E, 1996 “Visual search for a foreground object in continuous, naturalistic displays: The importance of shadows and occlusion” *Investigative Ophthalmology & Visual Science* **37**(4) S299 (abstract)
- DiLollo V, Enns J T, Rensink R A, 2000 “Competition for consciousness among visual events: The psychophysics of reentrant visual processes” *Journal of Experimental Psychology: General* **129** 481–507
- Elder J H, Trithart S, Pintilie G, MacLean D, 1998 “Rapid processing of cast and attached shadows” *Investigative Ophthalmology & Visual Science* **39**(4) S853 (abstract)
- Enns J T, Rensink R A, 1990 “Influence of scene-based properties on visual search” *Science* **247** 721–723
- He Z J, Nakayama K, 1992 “Surfaces versus features in visual search” *Nature* **359** 231–233
- Helmholtz H von, 1867/1962 *Treatise on Physiological Optics* volume 3 (New York: Dover, 1962); English translation by J P C Southall for the Optical Society of America (1925) from the 3rd German edition of *Handbuch der physiologischen Optik* (first published in 1867, Leipzig: Voss)
- Hering E, 1874/1964 *Outlines of a Theory of the Light Sense* (Cambridge, MA: Harvard University Press, 1964); English translation by L M Hurvich, D Jameson from *Zur Lehre vom Lichtsinn* (Vienna: Gerolds Söhne, 1874)
- Horn B K P, 1974 “Determining lightness from an image” *Computer Vision, Graphics, and Image Processing* **3** 277–299
- Julesz B, 1984 “A brief outline of the texton theory of human vision” *Trends in Neurosciences* **7** 41–45

-
- Kardos L, 1934 "Ding und Schatten: Eine experimentelle Untersuchung über die Grundlagen des Farbensehens" *Zeitschrift für Psychologie* **23** entire volume
- Kersten D, Mamassian P, Knill D C, 1997, "Moving cast shadows induce apparent motion in depth" *Perception* **26** 171–192
- Kleffner D A, Ramachandran V S, 1992 "On the perception of shape from shading" *Perception & Psychophysics* **52** 18–36
- Kriegman D J, Belhumeur P N, 2001 "What shadows reveal about object structure" *Journal of the Optical Society of America A* **18** 1804–1813
- Lynch D K, Livingston W, 1995 *Color and Light in Nature* (Cambridge: Cambridge University Press) pp 1–4
- Mamassian P, Knill D C, Kersten D, 1998 "The perception of cast shadows" *Trends in Cognitive Sciences* **2** 288–295
- Marr D, 1982 *Vision* (San Francisco, CA: W H Freeman)
- Mitsudo H, 2003 "Information regarding structure and lightness based on phenomenal transparency influences the efficiency of visual search" *Perception* **32** 53–66
- Nakayama K, He Z J, Shimojo S, 1995 "Visual surface representation: a critical link between lower-level and higher-level vision", in *Frontiers in Cognitive Neuroscience* 2nd edition Eds S Kosslyn, D N Osherson (Cambridge, MA: MIT Press) pp 1–70
- Neisser U, 1967 *Cognitive Psychology* (Englewood Cliffs, NJ: Prentice Hall)
- Ostrovsky Y, Cavanagh P, Sinha P, 2001 "Perceiving illumination inconsistencies in scenes" AI Memo 2001-029, Massachusetts Institute of Technology, Cambridge, MA, November 2001; <http://cbcl.mit.edu/cbcl/publications/memos-reports.html>
- Ramachandran V S, 1988 "Perceiving shape from shading" *Scientific American* **259**(8) 76–83
- Rauschenberger R, Yantis S, 2001 "Masking unveils pre-amodal completion representation in visual search" *Nature* **410** 369–372
- Rensink R A, 2000 "The dynamic representation of scenes" *Vision Cognition* **7** 17–42
- Rensink R A, Cavanagh P, 1993 "Processing of shadows at preattentive levels" *Investigative Ophthalmology & Visual Science* **34**(4) S1288 (abstract)
- Rensink R A, Enns J T, 1995 "Preemption effects in visual search: Evidence for low-level grouping" *Psychological Review* **102** 101–130
- Rensink R A, Enns J T, 1998 "Early completion of occluded objects" *Vision Research* **38** 2489–2505
- Treisman A, Gelade G, 1980 "A feature integration theory of attention" *Cognitive Psychology* **12** 97–136
- Treisman A, Gormican S, 1988 "Feature analysis in early vision: Evidence from search asymmetries" *Psychological Review* **95** 15–48
- Wolfe J M, 1998 "What can 1 million trials tell us about visual search?" *Psychological Science* **9** 33–39
- Wolfe J M, Yu K P, Stewart M I, Shorter A D, Friedman-Hill S R, Cave K R, 1990 "Limitations on the parallel guidance of visual search: Color \times color and orientation \times orientation conjunctions" *Journal of Experimental Psychology: Human Perception and Performance* **16** 879–892
- Zucker S W, 1987 "The diversity of perceptual grouping", in *Vision, Brain, and Co-operative Computation* Eds M A Arbib, A R Hanson (Cambridge, MA: MIT Press) pp 231–261

APPENDIX

Slopes, baselines, and error rates for all conditions. Data for target-absent trials are in parentheses.

Condition	Upright			Inverted		
	Slope/ms per item	Baseline/ms	Error/%	Slope/ms per item	Baseline/ms	Error/%
1A	13.4 (40.3)	588 (572)	1.9 (1.4)	4.8 (22.4)	550 (538)	1.7 (1.0)
1B	5.1 (12.1)	564 (562)	2.6 (1.7)	5.1 (11.4)	526 (508)	2.4 (1.5)
1C	2.9 (8.2)	594 (573)	1.8 (1.4)	3.2 (11.9)	624 (593)	2.4 (2.2)
1D	5.5 (11.2)	506 (536)	2.7 (3.1)	4.7 (10.2)	505 (520)	3.1 (2.1)
1E	3.7 (10.8)	548 (573)	3.4 (1.6)	2.4 (7.5)	520 (550)	3.2 (2.0)
1F	2.4 (12.6)	565 (547)	1.9 (1.5)	2.1 (10.0)	516 (515)	2.6 (1.9)
1G	3.2 (15.4)	570 (547)	2.6 (1.1)	3.0 (17.4)	540 (535)	1.9 (1.5)
1H	14.3 (42.9)	638 (622)	2.5 (1.3)	5.5 (22.0)	594 (600)	1.8 (1.2)
2A	30.4 (68.2)	706 (668)	5.4 (2.4)	41.4 (96.0)	634 (581)	5.1 (1.6)
2B	54.5 (109.3)	739 (650)	7.7 (1.8)	47.3 (87.1)	631 (594)	5.8 (0.9)
2C	32.5 (53.1)	657 (642)	4.4 (3.2)	20.6 (38.2)	586 (585)	2.3 (2.2)
3A	4.9 (18.6)	568 (582)	2.3 (1.6)	9.2 (19.3)	545 (580)	2.4 (1.4)
3B	4.1 (23.9)	623 (598)	3.8 (4.3)	5.6 (20.2)	569 (551)	3.4 (2.6)
3C	4.3 (16.6)	546 (533)	2.2 (1.1)	2.6 (9.5)	514 (536)	3.0 (1.2)
3D	12.5 (34.9)	590 (597)	2.4 (1.8)	4.4 (18.1)	551 (570)	1.6 (1.4)
3E	13.9 (39.6)	688 (682)	2.9 (1.9)	12.0 (35.9)	630 (617)	2.7 (1.5)
3F	7.4 (19.9)	563 (560)	5.8 (3.4)	7.2 (14.5)	532 (532)	4.5 (2.5)
3G	3.5 (19.6)	536 (548)	2.2 (1.8)	4.4 (13.3)	491 (526)	1.8 (1.7)
3H	12.0 (34.3)	548 (539)	3.4 (1.1)	6.5 (22.2)	522 (530)	2.3 (1.5)
4A	6.9 (17.2)	532 (550)	2.8 (1.4)	7.7 (18.8)	535 (550)	2.3 (2.2)
4B	3.2 (17.2)	563 (555)	2.3 (1.7)	3.0 (17.1)	527 (543)	2.1 (1.4)
4C	6.5 (18.7)	596 (590)	2.6 (1.8)	6.1 (11.8)	533 (542)	2.5 (1.2)
4D	4.1 (22.3)	572 (558)	1.9 (2.0)	1.5 (17.1)	552 (544)	2.7 (1.7)
5A	6.0 (18.4)	626 (615)	3.0 (1.9)	5.1 (16.1)	577 (588)	1.6 (1.0)
5B	6.7 (24.6)	613 (605)	1.4 (1.0)	5.4 (17.9)	568 (562)	2.1 (0.5)
5C	6.2 (25.0)	618 (595)	2.3 (1.5)	4.2 (8.3)	571 (590)	3.2 (1.7)
5D	1.9 (14.7)	577 (583)	2.1 (1.1)	3.7 (9.6)	558 (570)	1.8 (1.2)
5E	8.8 (29.9)	628 (615)	5.0 (2.9)	4.0 (17.0)	577 (569)	3.7 (1.9)
5F	3.3 (22.7)	633 (669)	1.9 (0.8)	5.7 (15.2)	574 (613)	1.2 (0.8)
5G	3.8 (20.7)	617 (617)	3.2 (1.2)	3.1 (12.4)	569 (573)	2.2 (0.8)
5H	6.1 (47.2)	607 (573)	2.1 (1.5)	4.5 (22.5)	545 (548)	1.8 (1.5)

ISSN 0301-0066 (print)

ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 33 2004

www.perceptionweb.com

Conditions of use. This article may be downloaded from the Perception website for personal research by members of subscribing organisations. Authors are entitled to distribute their own article (in printed form or by e-mail) to up to 50 people. This PDF may not be placed on any website (or other online distribution system) without permission of the publisher.