sequence processing on phoneme monitoring *Brain* 117, 671–682
48 Price, C.J. and Friston, K.J. (1997) Cognitive conjunctions: a new approach to brain activation experiments *NeuroImage* 5, 261–270
49 De Renzi, E., Zambolin, A. and Crisi, G. (1987) The pattern of neuropsychological impairment associated with left posterior cerebral artery territory infarcts *Brain* 110, 1099–1116
50 Foundas, A.L. *et al.* (1998) Anomia: case studies with lesion localisation *Neurocase* 4, 35–43
51 Menard, M.T. *et al.* (1996) Encoding words and pictures: a positron emission tomography study *Neuropsychologia* 34, 184–194
52 Law, I. *et al.* (1991) Left supramarginal/angular gyri activation during reading of syllabograms in the Japanese language *J. Neurolinguist.* 6, 243–251
53 Coltheart, M. (1980) Reading, phonological recoding and deep dyslexia, in *Deep Dyslexia* (Coltheart, M., Patterson, K.E. and Marshall, J.C., eds), pp. 197–226, Routledge
54 Vanier, M. and Caplan, D. (1985) CT scan correlates of surface dyslexia, in *Surface Dyslexia* (Patterson, K.E., Marshall, J.C. and Coltheart, M., eds), pp. 511–525, Erlbaum
55 Marin, O.S.M. (1980) CAT scans of five deep dyslexic patients (Appendix 1), in *Deep Dyslexia* (Coltheart, M., Patterson, K.E. and Marshall, J.C., eds), pp. 407–433, Routledge
56 Talairach, J. and Szikla, G. (1967) *Atlas D'anatomie Stéréotaxique du Télencéphale*, Masson, Paris
57 Vallar, G. and Shallice, T., eds (1990) *Neuropsychological Impairments To Short Term Memory*, Cambridge University Press

The perception of cast shadows

Pascal Mamassian, David C. Knill and Daniel Kersten

When an object casts its shadow on a background surface, the shadow can be informative about the shape of the object, the shape of the background surface and the spatial arrangement of the object relative to the background. Among all these roles, we found that cast shadows were perceptually most relevant for the recovery of spatial arrangement, especially when the shadow is in motion. This finding is intriguing when one considers the ambiguities in the possible ways that shadow motion can be interpreted. We reasoned that the visual system must use *a priori* constraints to disambiguate the cast shadow motion. One of these constraints is that the light source is stationary. Though simple, the stationary-light-source constraint supports rich, reliable inferences about the qualitative motions of objects in three dimensions.

P. Mamassian is at the Psychology Department, University of Glasgow, 58 Hillhead Street, Glasgow, UK G12 8QB.

tel: +44 141 330 4663
fax: +44 141 339 8889
e-mail: pascal@psy.gla.ac.uk

D. Knill is at the Psychology Department, University of Pennsylvania, 3815 Walnut Street, Philadelphia, PA 19104, USA.

tel: +1 215 898 4253
fax: +1 215 898 7301
e-mail: knill@cattell.psych.upenn.edu

D. Kersten is at the Psychology Department, University of Minnesota, 75 East River Road, Minneapolis, MN 55455, USA.

tel: +1 612 625 2589
fax: +1 612 626 2079
e-mail: kersten@cyclops.psych.umn.edu

> It is on this natural weakness of ours that the shadow painting and conjuror art and their fellows rely when they deceive us with their tricks.
>
> Plato, *The Republic*, 602d

The importance of shadows in Western cultures cannot be underestimated. Pliny the Elder placed the origin of painting at the first time the shadow of a man was outlined on a wall, and Plato used the metaphor of a cave on the walls of which shadows of the external world were projected to explain his theory of knowledge[1]. As these examples illustrate, shadows were early on appreciated for their likeness with real objects. However, the first thorough analysis of shadows had to wait until 1490 when Leonardo da Vinci described how artists could use light and shade to evince perceptions of three-dimensional relief in paintings[2]. While da Vinci naturally focused on shadows in static images, recent developments have shown that shadows are particularly salient cues to depth in dynamic scenes where objects and their cast shadows are moving relative to one another. This paper reviews recent work which has elucidated the informational structure of moving shadows and has raised important and difficult theoretical issues about the visual interpretation of depth in dynamic scenes.

In order to understand the information content of shadows, one must first recognize that shadows come in two types, depending on how they are formed on surfaces. We will refer to the two types as 'attached' and 'cast' shadows (Fig. 1). Shadows are regions of a surface which receive no illumination from a light source. Attached shadows are formed when a surface obstructs the light falling on itself. They include two contiguous sub-regions that are physically hard to separate: regions which are in shadow because they face away from a light source, and neighbouring parts of a surface that these regions occlude from the light source[3]. Cast shadows are formed when one surface occludes another surface from the light source. In this case, an 'image' of the casting surface is formed on the occluded surface, and the cast shadow is 'surrounded by light'[4]. Shadows should also be distinguished from shading. In contrast to shadows, shaded regions of a surface directly face the light source so that a change in local surface orientation will cause a gradual change in reflected light[5].

The information available from attached shadows is necessarily restricted to the shape of the surface on which they

 1364-6613/98/$19.00 PII: S1364-6613(98)01204-2

are formed. Cast shadows, on the other hand, are related to two distinct surfaces, the surface of the casting object and the surface on which the shadow is cast. As such, cast shadows are potentially informative about the shapes of either of the surfaces and about the spatial layout of the scene. Spatial layout refers to the spatial relationships between the surfaces in the scene, both qualitatively (e.g. object A is to the right of object B) and quantitatively (e.g. the distance between objects A and B is 5 cm). One main difference between the roles of attached and cast shadows is therefore that the former leads to questions regarding local scene properties (e.g. what local surface shape can produce the observed attached shadow boundary?) and the latter to questions of global scene characteristics (e.g. how far from the background an object has to be to produce the observed cast shadow?).

This paper will focus on the perceptual interpretation of cast shadows. As we will show, the human visual system does not effectively use cast shadows as cues to surface shape, despite the potential reliability of the information they provide. Conversely, cast shadows are very salient cues to the spatial layout of objects in a scene. The contrast between cast shadows as cues to surface shape and cast shadows as cues to spatial layout seems also to hold in dynamic scenes. Particularly striking is the observation that moving cast shadows can produce a very vivid impression of an object moving in depth. Moreover, simply manipulating the motion of an object's cast shadow can dramatically change the perceived trajectory of the object. We will briefly review the evidence for these claims and then discuss the particular problem of inferring spatial layout from moving cast shadows. In particular, we will analyse the nature of the *a priori* constraints needed to infer surface layout from moving cast shadows and relate these to the perceptual phenomenology.

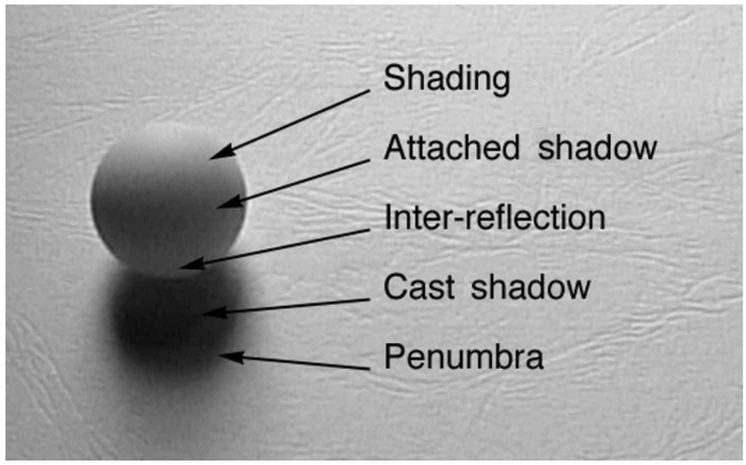

Fig. 1 Definitions of shadows. Shading is the variation of reflected light on a surface patch which faces directly the light source. Shadows are regions occluded from the light source and come in two types: attached shadows are formed on the very surface which is occluding the light whereas cast shadows are formed on remote surfaces. For extended light sources, penumbras surround the cast shadows. Finally, attached shadows sometimes include inter-reflections that result from light rays bouncing back from surrounding surfaces.

Cast shadow cues

The shape of a cast shadow is related to the shapes of both the casting object and the receiving surface, while the location of the shadow is a function of the location of these two surfaces. In this section, we shall present evidence suggesting that cast shadows are more effective as a cue for spatial layout than as a shape cue.

Static cue for surface shape

Cast shadows are potentially informative about the shapes of both the casting object and the surface on which they lay. In a constrained world where objects have only planar faces, shadows are more informative about the shape and position of the receiving surface than about the object that caused them[6,7]. In a more general world of piece-wise smooth surfaces, the ambiguity in the information about the casting surface increases, but the information provided about the shape of the shadowed surface remains strong, particularly at creases (i.e. orientation discontinuities)[3]. Intuitively, it is easy to imagine the complexity of recovering the shape of an object from its silhouette when that silhouette has been deformed by the shape of another surface. To see how the shape of a cast shadow can provide reliable information about the surface on which it is cast, consider the picture of a pencil casting a shadow on a folded card shown in Fig. 2. The shape of the card by itself is ambiguous: it can be folded like a 'W' or an 'M' viewed from above (this bistable perceptual display is known as the 'Mach card illusion'). When the pencil casts a shadow on the folded card, the shape of the card should no longer be ambiguous: the card is convex at points where the shadow touches the pencil and concave at points where the shadow is the farthest from the pencil. The presence of the shadow casting object should serve to strengthen the information provided by the cast shadow, but it is not technically required to make the information reliable, given some qualitative knowledge about the direction of the light source[3].

We have performed numerous experiments using ambiguous shapes like that shown in Fig. 2 in an attempt to show that the visual system can use the shapes of cast

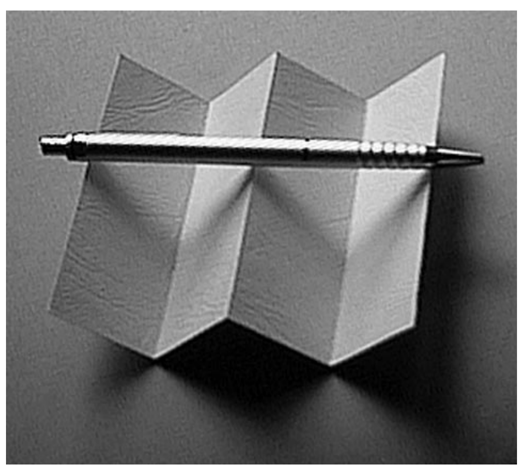

Fig. 2 Pencil over a folded card. The folds of the card are ambiguous in that they can appear convex or concave. The shape of the shadow cast by the pencil should disambiguate the fold, because when the shadow touches the pencil, the fold is necessarily convex. Nevertheless, the shape of the card remains perceptually equally ambiguous. The reader is encouraged to try this demonstration with a real pencil on a real folded card, viewing with one eye shut.

shadow boundaries to constrain the perception of surface shape. In particular, we have manipulated the opacity and penumbras of the shadows to increase the belief that the dark regions depicted realistic cast shadows. However, in no case have we found a significant effect of cast shadows on subjects' judgments. Even though the shape of the shadow in Fig. 2 is theoretically sufficient to disambiguate the shape of the card, the shape remains ambiguous perceptually. While such a negative result is necessarily limited to the types of stimuli we used in the experiments, our repeated inability to find an effect strongly suggests that the visual system does not make use of the information provided by cast shadows about surface shape, or if it does, the information is perceptually 'weak'.

Dynamic cue for surface shape

When the object casting the shadow is moving, more information can be gathered from the transformation of the shadow across time. As a source of information about the shape of the casting surface, dynamic shadows are particularly effective when formed on planar surfaces. In that case, the problem of inferring the shape of the casting surface reduces to a generalized form of the structure from motion problem[8]. Indeed, initial demonstrations of the kinetic depth effect used a shadow casting method to create stimuli[9]. These stimuli typically elicit percepts of rigid three-dimensional structures in motion when the shadows contain trackable feature points. When the surface of the object is smooth, so that the correspondence of its shadow contour across time is more difficult to establish[10,11], the object is more likely to appear non-rigid[12,13]. In these studies, however, observers interpret the shadow as the profile of an object rather than the shadow cast by an object. It is possible that shadows will be processed differently by the visual system when they are seen as mere shadows. Whether or not the shadow information would then constrain the interpretation of the casting object remains an open question.

The motion of dynamic cast shadows also provides some information about the shape of the shadowed surface when it is not flat. Consider, for example, replacing the pencil in Fig. 2 with a ball moving along the path defined by the axis of the pencil. The ball's shadow in the dynamic case would then follow the path defined by the pencil's shadow in the static case. Therefore, the information provided by dynamic cast shadows about the shape of the shadowed surface is qualitatively similar to the information provided by static cast shadows. As for static shadows, we have looked at the effect of moving cast shadows on the perceived shape of the shadowed surface. Simulating a ball flying above a folded card, we found no effects of the cast shadow trajectory on the perceived convexity of the card's folds. What we did find was a strong impression of a ball moving in three-dimensional space. This led us to a detailed investigation of the perception of dynamic surface layout from moving cast shadows. We now turn to this topic after a brief description of the static case.

Static cue for spatial layout

Qualitatively, the information provided by an object's cast shadow about its disposition in three-dimensional space is straightforward: the closer an object is to the background surface, the closer the shadow will be to the casting object. It is this property which da Vinci encouraged artists to take advantage of in his writings on cast shadows[4]. Despite da Vinci's ministrations, visual artists have been surprisingly reluctant to include shadows in their paintings[14], though they do appear prominently in production graphics, as in the use of drop-shadows. Many popular computer graphics packages include drop-shadows as a standard format for titles, legends, etc., and they appear prominently in web page graphics. The effects of static cast shadows on perceived depth relations between surfaces is phenomenologically clear, as shown in Fig. 3. Their perceptual salience has also been demonstrated experimentally, particularly in their ability to anchor surfaces in depth when they are placed in a perspective rendering[15,16].

Dynamic cue for spatial layout

Just as static cast shadows provide information about static depth relations between objects, the motions of cast shadows in dynamic scenes provides information about the relative motions of objects in scenes. We will describe the structure of this information in some detail in the next section; however, its basic features are fairly simple – the motion of an object that casts a shadow on a background surface will cause the cast shadow to move in the image. The motion of the shadow, therefore, provides information about the relative motions of the object and the background surface. Motion of a cast shadow away from the casting object suggests that the object and background surface are moving apart in space, and conversely for motion towards the object (but see next section for an elucidation of the implicit assumptions on which this description is based).

We have created a demonstration illustrating the perceptual salience of this form of dynamic shadow information. By moving only the cast shadow of a static object in a scene, we were able to induce a strong percept of object motion in depth (despite the fact that the object did not move or change size in the image[17]). If the three images in Fig. 3 represent consecutive frames of a movie, the central square will appear to move in depth as a result of the translation of the shadow beneath it. The proportion of times the square is perceived moving in depth increases when a penumbra is surrounding the shadow boundary, when the shadow displacement is consistent with a light located above the scene, and when the shadow is dark rather than light[18].

When the object is also moving, its trajectory in depth can be dramatically influenced by the trajectory of its cast shadow. For instance, a ball placed inside a box can appear to roll on the bottom of the box or rise in a frontal plane by a simple manipulation of the shadow trajectory (Fig. 4). We quantified the effect more formally by having observers estimate the height on the wall to which the ball appeared to move at the rightmost point in its trajectory. Subjects' estimates were strongly correlated with the path of the object's cast shadow, consistent with the phenomenal observations[18]. Interestingly, the perceived trajectory of the ball in space was not affected by changing the contrast, the opacity or even the shape of the shadow. This differs from

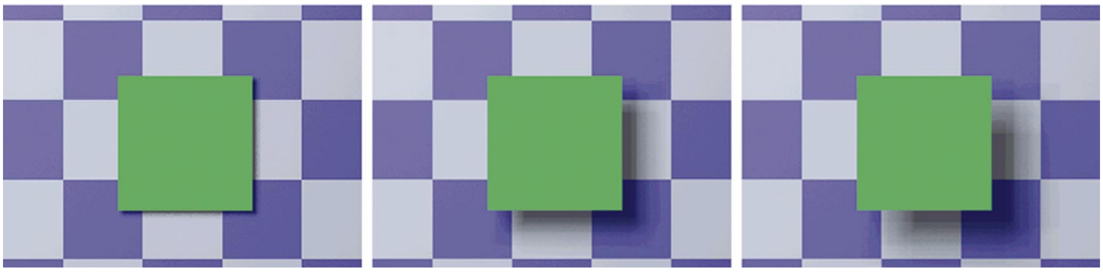

Fig. 3 Effects of cast shadow position on object location. The square appears farther from the background checkerboard when the cast shadow is detaching from the square. When these images are shown in succession, the square appears to move in depth even though the square is strictly static. (A movie demonstration of this can be seen on the web at http://vision.psych.umn.edu/www/kersten-lab/shadows.html)

the effects found for the apparently moving square (Fig. 3), a point to which we will return to in the discussion.

The vivid impressions of spatial layout induced by moving cast shadows contrasts with the limited influence they have on the perception of surface shape. In the next section, we outline some of the assumptions the visual system uses to infer the spatial layout from moving cast shadows.

Cast shadow constraints

The motion of a cast shadow is inherently ambiguous. Consider again the simple configuration of a square overlaid on a checkered background (Fig. 3). The location of the shadow cast by the square on the background could be the result of an infinite number of combinations of the positions of the viewpoint, the light source, the casting object and the background surface. As a consequence, when a shadow is moving across the retina, any combination of these four factors could be responsible for the shadow displacement (see Box 1). Nevertheless, a displacement of the cast shadow produces a percept of a square moving in depth, fairly consistently over repeated observations[17].

The selection of one particular interpretation over the multiple others indicates that the visual system uses some *a priori* constraints to resolve the ambiguity in cast shadow information. In this review, 'perceptual constraints' refer to assumptions made by the visual system unless contrary evidence is available. It is important to keep in mind that these perceptual constraints can be violated when contrary information is available in the image. For instance, we might assume that light is coming from above unless there are familiar objects in the scene (e.g. faces) that are obviously lit from below. When the scene is complex and seemingly ambiguous in many ways, perceptual constraints enable us to arrive at a quick and stable understanding of a visual scene. We consider now the constraints that justify why a moving cast shadow produces an impression of an object motion rather than a light source, background or viewpoint motion.

Light source motion

The results of our experiments are consistent with the hypothesis that the visual system assumes light sources to be stationary in the scene. Observers perceived a square moving in depth in Fig. 3 even when the image sequence was obtained by moving the light source rather than the square. Likewise, observers perceived a ball rolling on the bottom of the box or rising in a frontal plane in the animations illustrated in Fig. 4, even when the image sequence was generated by moving the light source[18]. Adding more evidence that the light was actually moving, such as adding more objects in the scene, did not affect the perceived trajectory of the ball[18]. Thus, the stationary light source constraint appears to be a strong prior assumption which observers make about scenes, rather than an inference based on image data.

To assume a stationary light source is reasonable when the average speed of objects moving around us is compared to the speed of the sun as seen from the earth. If the light source is fixed in space, where is it? It is well-known that by default, human observers assume the light to come from a single source located above their head[19–22]. From the measurements of perceived trajectories in the ball-in-the-box experiments (Fig. 4), we were able to derive more precisely the assumed light source location. For most observers, the data were consistent with an assumption of a light source located in the close vicinity of their head. Changing the shading on

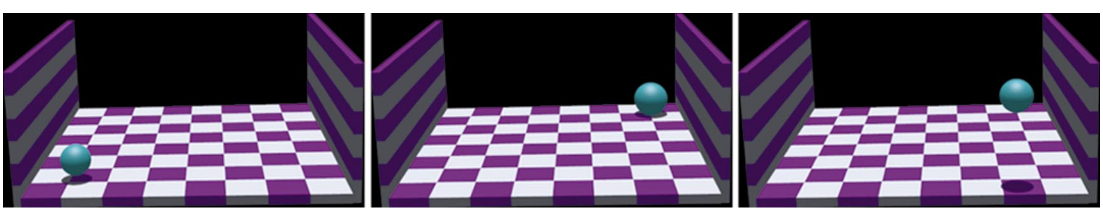

Fig. 4 Single frames of a movie showing a ball moving inside a box. The ball can be made appear to roll on the bottom of the box or rise in a frontal plane by arranging the ball's shadow to follow the same trajectory as the ball or a horizontal trajectory. The perceived location of the ball can be assessed by having observers adjust the height of a horizontal bar to the perceived height of the center of the ball at the right end of its trajectory. (See movie on the web at http://vision.psych.umn.edu/www/kersten-lab/shadows.html)

Box 1. The potential sources of a cast shadow displacement

A scene which contains cast shadows is composed of four basic elements localized in space: the light source, the background surface on which the shadow appears, the viewpoint and the object casting the shadow. It is important to single out each of these elements and to study the changes that they induce in the image when they are displaced. Here, we use the square-over-checkerboard movie to illustrate our points (see Fig. 3 in main article). We assume that the scene is viewed under perspective projection. The displacement of the shadow can be characterized as a change in visual angle from θ_1 to θ_2 for the location of the shadow relative to the casting object. Intuitively, θ_1 and θ_2 can be thought of as the distance in the image between the square and its shadow.

(1) *Light source* (L): assuming that light is coming from above, a displacement of the light source towards the background surface (L_1 to L_2) will result in a downward displacement of the shadow (Fig. A). Moving the light source will also produce a deformation of the shapes of the cast and attached shadows and a change in shading and penumbra (the penumbra shrinks as the light source gets closer to the object).

(2) *Background surface* (B): the background surface is the surface on which the shadow is cast. A displacement of the background surface away from the object (B_1 to B_2) will result in a downward displacement of the shadow (Fig. B) and an increase in the size of the penumbra. Moving the background surface will also produce a change in the projected texture (texture elements will get smaller but denser as the background gets farther away from the viewpoint).

(3) *Casting object* (O): a displacement of the casting object towards the viewpoint (O_1 to O_2) will result in a downward displacement of the shadow (Fig. C). Moving the object closer to the observer should also increase its image size and change the penumbra around the cast shadow.

(4) *Viewpoint* (V): we shall assume that the scene is viewed from a single viewpoint. Moving the viewpoint closer to the object (V_1 to V_2) will result in a downward displacement of the shadow (Fig. D). A displacement of the viewpoint should also induce a change in the image size of the object and the texture of the background.

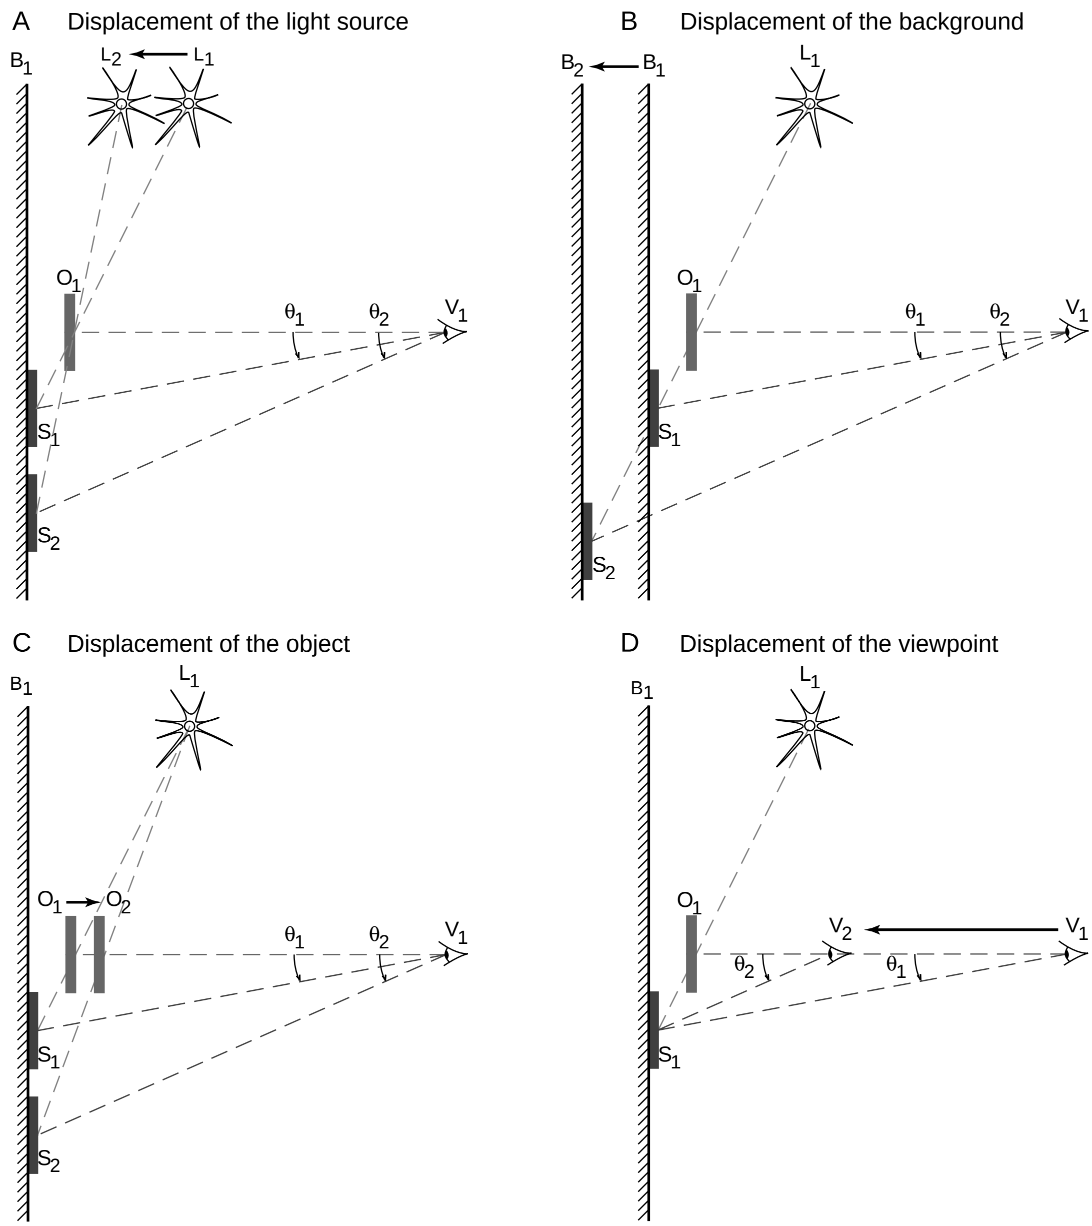

the ball to suggest illumination from alternative directions only changed the perceived trajectory of the ball slightly – not nearly as much as was predicted by the new light source location[18]. This limited ability of human observers to pick up information about the light source location is a recurrent finding[23,24].

Background motion

The background surface is not only the surface on which the shadow is cast, it is also the reference surface against which the movement of an object can be detected. If one merely looks at the relative motion of the object and the background, the absolute motions of object and background are lost. In studies of induced motion, a static object can be perceived to move when only the background is moving[25–28]. Similarly, our observers preferred to see the square moving in Fig. 3 rather than the background. In other words, it seems plausible that a stationary background constraint was used to simplify the interpretation of the scene.

Viewpoint motion

When the observer is getting closer to the object, the visual angle between the object and its cast shadow increases just like it would if it was the object that was moving instead. This ambiguity between observer and object motion was never experienced: none of the observers in our experiments ever reported a perception of self-motion. Of course, the information provided by the vestibular system was contrary to this possibility. On the other hand, it is well-known that self-motion can be experienced from visual motion alone[29–31]. The reason why self-motion was not experienced in our displays is probably linked to the fact that such experience usually requires a large optic flow, or at least a visual stimulation where the possibility of a static reference frame has been discarded. As we saw above, it appears that observers use a stationary background constraint to interpret Fig. 3, which in turn implies a constraint of stationarity for the viewpoint.

Object motion

We have seen that the displacement of the shadow in Fig. 3 was preferentially perceived as resulting from the displacement of the square in space. It is important to realize that this interpretation violates two constraints. First, the square is moving directly along the line of sight, which means that the viewpoint is in an accidental position: a small shift of the viewpoint will make the square move across the retinal plane as well as in depth. The freedom to allow the viewpoint to be anywhere in the scene rather than at accidental positions is referred to as the generic viewpoint constraint. It is interesting to note that humans behave as if they use this constraint[32,33]. In the movie, the shadow motion clearly overrides the generic viewpoint constraint. Second, even though the object is moving in depth, its image does not change size – a violation of the change-of-size cue (i.e. the image gets bigger as the object gets closer[34]). In fact, when the displacement of the object induced by the shadow motion is increased, the object appears to change size in a way that is consistent with its perceived depth (that is, to inflate as it recedes in depth so that its image is kept constant[18]). This phenomenon is quite striking, especially in the light of the evidence that the visual system prefers to assume objects to be rigid, and thus fixed in size[8]. The fact that the cast shadow cue can override the generic viewpoint constraint, the constant size cue as well as the rigidity constraint is a qualitative index of the strength of the cast shadow cue for dynamic spatial layout perception.

The cue of object motion from cast shadow motion is different from other depth cues. While most depth cues provide some information locally, cast shadows act at a distance from the object which is moving in depth. In other words, cast shadows are non-commensurate with other depth cues because they provide global rather than local information about spatial layout. Nevertheless, an interaction between cast shadows and other depth cues has been found experimentally, in particular the cues of size change, binocular disparity and motion parallax[18,35]. As such, cast shadows pose a challenge for models of depth cue interaction where depth in each visual direction is averaged across all depth cues[36].

Conclusions

In this review, we have outlined the role of cast shadows for the perception of surface shape and spatial layout. Cast shadows are those shadows that are projected on a remote surface. We found that even though they were potentially informative about surface shape, cast shadows served very weakly, if at all, in this function. On the other hand, cast shadows clearly provide very salient cues for the relative dispositions of objects in space, particularly when an object and its cast shadow are moving. This raises some unique and difficult conceptual issues for perception. The issues revolve around three problems: segmenting and labeling cast shadows in scenes, linking cast shadows with the objects which cast them and interpreting spatial relations from the changing displacement between an object and its shadow in an image.

Segmentation and labeling

Before they can be used as cues to spatial layout, cast shadows must be segmented from the background and labelled as cast shadows. In static images, shadows can easily be confused with other events on surfaces: black paint (surface markings), pieces of black fabric (like in Dali's 'Disappearing Bust of Voltaire') or even stain on clothes (like in Rembrandt's 'The Night Watch' cited by Arnheim[37]). The fact that they are just like transparent surfaces[38] which are darker than their immediate vicinity[39] is barely sufficient to detect them. However, the presence of penumbra may help distinguish shadows from sharp surface markings[40].

In dynamic scenes, a feature appears which can very reliably determine proper shadow labeling: the relative motion of an object and its cast shadow is constrained to follow a line connecting the object to the light source. Assuming the light source is fixed, the virtual line connecting an object and its cast shadow will thus revolve around a common point (the projection of the light source in the image). For light sources distant from a scene, the motion is constrained

Outstanding questions

Psychological

- Shadows in a scene sometimes fail to appear like shadows, and, conversely, dark materials can be misperceived as shadows. What are the criteria used by human observers to categorize a patch in the image as a shadow? How good are we at discriminating attached from cast shadows?
- What are the mechanisms which led us to select the particular constraints used by our visual system? What is the contribution of exposition to natural scenes to the final selection of constraints?

Neural

- It is now believed that there exists two streams of visual processing in the cortex: one occipito-temporal stream dealing with the shape of objects and one occipito-parietal stream dealing with their location in space. If this view is correct, are attached shadows processed in the former and cast shadows in the latter stream?
- Do special neural mechanisms exist in the visual system for measuring the particular type of correlated motion which is found between objects and their cast shadow? For instance, is the firing rate of cortical cells (in particular in areas V1 and MT of macaque monkeys) affected by correlated motion in another part of the visual field?

Computational

- How can a visual system compute spatial layout from cast shadow information? Can it be done with a simple heuristic which embodies fixed assumptions about light source direction, background surface shape and viewing position, or must it resolve the information cooperatively with estimates of these other scene properties?
- Cast shadows provide some information about the location of remote objects. How can this 'delocalized' information be integrated with more local forms of depth information such as stereopsis?

to always be in a common direction in the image. The presence of such a constrained motion in the image is a strong indicator that two moving patches are related as an object to its shadow.

The importance of correlated motion as a cue to shadow labeling may explain the divergent phenomenal effects of violating luminance constraints in the square-above-the-checkerboard movie and the ball-in-a-box movie (see descriptions above). In the former, the object (square) is static, so that the motion of the cast shadow is trivially correlated with that of the object (by virtue of being linear motion). In this special case, object-shadow correlated motion does not provide a reliable cue to shadow labeling, and violations of shadow luminance constraints (such as making the shadow lighter than the background) do reduce the efficacy of the shadow cue. In the ball-in-a-box movie, the object itself is moving, hence the object-shadow correlated motion which appears in the animation would be accidental if caused by anything other than an object casting a shadow. It therefore provides a strong cue to shadow labeling, and the phenomenal percept of motion is more robust to violations of shadow luminance constraints.

Linking objects with their shadows

In order to use cast shadow information to surface layout, shadows must be properly linked to the objects that cast them. This is yet another variant of the ubiquitous matching problem in visual perception. In static scenes, determining the correspondence between an object and its cast shadow is a very difficult problem[7]. Few simple cues are available, besides the relative distances between a shadow and candidate objects. Theoretically, one could match the shape of a shadow to that of objects in a scene. However, this matching procedure would appear to be computationally prohibitive for even the simplest objects, and this is probably the reason why having a trapezoidal rather than elliptic cast shadow in the ball-in-a-box movie did not reduce the effectiveness of the percept[18].

In dynamic scenes, in which both objects and their cast shadows are moving, the correlated motion between an object and its shadow provides a reliable and simple source of information for linking cast shadows and objects. This further emphasizes the potential importance of object-shadow correlated motion for shadow interpretation. The relative simplicity of the correlated motion feature suggests the intriguing possibility that the visual system contains low-level mechanisms for detecting such correlations. One way to experimentally test this possibility would be to decorrelate progressively the object and shadow motions and check the induced effect on the perceived object motion.

Inferring spatial relations

Finally, we come to the basic problem of actually inferring spatial relations from the relative dispositions of objects and their cast shadows. Assuming they have been properly segmented, labeled and linked, cast shadows provide information for the spatial layout of the scene (i.e. the relative positions of surfaces and objects in the world). In dynamic scenes, the changing distance between an object and its shadow provides information about the relative motions of objects in the world. Our analysis of cast shadow motion information in the previous section highlighted the different factors which determine the pattern of object-cast shadow motion appearing in an image. This analysis revealed that the interpretation of an object moving in depth was chosen against other equally valid interpretations involving a motion of the light source, the background surface, or the viewpoint. Taking into account these biases (or perceptual constraints), cast shadow motion provides crucial information about the relative motion of objects in a scene. The way perceptual constraints interact with the information in the image can be formally described within a Bayesian framework[41,42].

In summary, we found that cast shadows have a strong influence on the perception of spatial layout, but only a weak impact on the shape of the objects casting or receiving the shadows. One reason for this dichotomy might be that straightforward assumptions allow the visual system to make reliable inferences about the relative position of objects while equally powerful assumptions seem hard to come by for the interpretation of the objects' shape. However, the link between perceptual constraints and statistical properties of natural scenes remains to be ascertained.

Acknowledgements

Some of the ideas presented in this paper were presented at the annual meeting of the Association for Research in Vision and Ophthalmology in 1991 and 1992, and at the European Conference on Visual Perception in 1992. We thank Jim Enns and the other reviewer for their acute comments. Funding was provided by a grant from the National Science Foundation (NSF) SBR-9631682.

References

1 Stoichita, V.I. (1997) *A Short History of the Shadow*, Reaktion Books
2 Baxandall, M. (1995) *Shadows and Enlightenment*, Yale University Press
3 Knill, D.C., Mamassian, P. and Kersten, D. (1997) Geometry of shadows *J. Opt. Soc. Am. A* 14, 3216–3232
4 da Vinci, L. (1970) *The Notebooks of Leonardo da Vinci* (Vol. 1), Dover
5 Horn, B.K. and Brooks, M.J. (1989) *Shape from Shading*, MIT Press
6 Waltz, D.L. (1975) Understanding line drawings of scenes with shadows, in *The Psychology of Computer Vision* (Winston, P.H., ed.), pp. 19–91, McGraw-Hill
7 Shafer, S.A. (1985) *Shadows and Silhouettes in Computer Vision*, Kluwer
8 Ullman, S. (1979) *The Interpretation of Visual Motion*, MIT Press
9 Wallach, H. and O'Connell, D.N. (1953) The kinetic depth effect *J. Exp. Psychol.* 45, 205–217
10 Giblin, P.J. and Weiss, R. (1987) Reconstruction of surfaces from profiles *Proc. First Int. Conf. Comput. Vis.*, pp. 136–144, IEEE Computer Society Press
11 Cipolla, R. and Blake, A. (1992) Surface shape from the deformation of apparent contours *Int. J. Comput. Vis.* 9, 83–112
12 Todd, J.T. (1985) Perception of structure from motion: is projective correspondence of moving elements a necessary condition? *J. Exp. Psychol. Hum. Percept. Perform.*11, 689–710
13 Pollick, F.E. (1994) Perceiving shape from profiles *Percept. Psychophys.* 55, 152–161
14 Gombrich, E.H. (1995) *Shadows: The Depiction of Cast Shadows in Western Art*, National Gallery Publications, London
15 Yonas, A., Goldsmith, L.T. and Hallstrom, J.L. (1978) Development of sensitivity to information provided by cast shadows in pictures *Perception* 7, 333–341
16 Yonas, A. (1979) Attached and cast shadows, in *Perception and Pictorial Representation* (Nodine, C.F and Fisher D.F., eds), pp. 100–109, Praeger
17 Kersten, D. *et al*. (1996) Illusory motion from shadows *Nature* 379, 31
18 Kersten, D., Mamassian, P. and Knill, D.C. (1997) Moving cast shadows induce apparent motion in depth *Perception* 26, 171–192
19 Rittenhouse, D. (1786) Explanation of an optical deception *Trans. Am. Philos. Soc.* 2, 37–42
20 Brewster, D. (1826) On the optical illusion of the conversion of cameos into intaglios and of intaglios into cameos, with an account of other analogous phenomena *Edinburgh J. Sci.* 4, 99–108
21 Ramachandran, V.S. (1988) Perceiving shape from shading *Sci. Am.* 255, 76–83
22 Ramachandran, V.S. (1988) Perception of shape from shading *Nature* 331, 163–166
23 Hagen, M.A. (1976) The development of sensitivity to cast and attached shadows in pictures as information for the direction of the source of illumination *Percept. Psychophys.* 20, 25–28
24 Pentland, A.P. (1982) Finding the illuminant direction *J. Opt. Soc. Am.* 72, 448–455
25 Duncker, K. (1929) Über induzierte Bewegung: ein Beitrag zur Theorie optisch wahrgenommener Bewegung *Psychologische Forschung* 12, 180–259
26 Bridgeman, B., Kirch, M. and Sperling, A. (1981) Segregation of cognitive and motor aspects of visual function using induced motion *Percept. Psychophys.* 29, 336–342
27 Masson, G., Proteau, L. and Mestre, D. (1995) Effects of stationary and moving textured backgrounds on the visuo-oculo-manual tracking in humans *Vis. Res.* 35, 837–852
28 Smeets, J.B.J. and Brenner, E. (1995) Perception and action are based on the same visual information: distinction between position and velocity *J. Exp. Psychol. Hum. Percept. Perform.* 21, 19–31
29 von Helmholtz, H.L.F. (1910) *Treatise on Physiological Optics* (Translated Southall, J.P. 1925), Dover Publications
30 Gibson, J.J. (1950) *The Perception of the Visual World*, Houghton Mifflin
31 Warren, W.H., Morris, M.W. and Kalish, M. (1988) Perception of translational heading from optical flow *J. Exp. Psychol. Hum. Percept. Perform.* 14, 646–660
32 Nakayama, K. and Shimojo, S. (1992) Experiencing and perceiving visual surfaces *Science* 257, 1357–1363
33 Freeman, W.T. (1994) The generic viewpoint assumption in a framework for visual perception *Nature* 368, 542–545
34 Kaufman, L. (1974) *Sight and Mind: An Introduction to Visual Perception*, Oxford University Press
35 Ujike, H. (1998) Depth perception with shadow and disparity cue combinations, in *Selection and Integration of Visual Information: Proc. Int. Workshop Adv. Res. Vis. Cognit.*, pp. 137–150, STA-NIBH, Tsukuba, Japan
36 Landy, M. *et al*. (1995) Measurement and modeling of depth cue combination: in defense of weak fusion *Vis. Res.* 35, 389–412
37 Arnheim, R. (1954) *Art and Visual Perception: A Psychology of the Creative Eye* University of California Press
38 Metelli, F. (1975) Shadows without penumbra, in *Gestalttheorie in der modernen Psychologie* (Ertel, S., Kemmler, L. and Stadler, L., eds), pp. 200–209, Dietrich Steinkopff
39 Cavanagh, P. and Leclerc, Y.G. (1989) Shape from shadows *J. Exp. Psychol. Hum. Percept. Perform.* 15, 3–27
40 Marr, D. (1982) *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*, W.H. Freeman
41 Knill, D.C., Kersten, D. and Mamassian, P. (1996) Implications of a Bayesian formulation of visual information processing for psychophysics, in *Perception as Bayesian Inference* (Knill, D.C. and Richards, W., eds), pp. 239–286, Cambridge University Press
42 Mamassian, P. and Landy, M.S. Observer biases in the 3D interpretation of line drawings *Vis. Res.* (in press)