The shadow knows: a primer on the informational structure of cast shadows

Roberto Casati

► To cite this version:

Roberto Casati. The shadow knows: a primer on the informational structure of cast shadows. Perception, 2004, 33 (1), pp.1385-1396. ijn_00000668

HAL Id: ijn_00000668

https://hal.science/ijn_00000668

Submitted on 8 Feb 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Paper accepted for publication on <i>Perception</i> .	Submitted version, revised 12.03.2004. Please do not copy or circulate.
--	--

Roberto Casati
Institut Nicod
CNRS-EHESS-ENS

The Shadow Knows:
A primer on the informational structure of cast shadows

Introduction

Perception *from* shadow is a large if gerrymandered topic in visual cognition. Shape and distance from shadow are the subject of countless articles in natural and artificial vision science.¹ In this vast literature, the visual system has been documented to be able to take into account a number of informational structures associated with cast, attached and self-shadows. Cast shadows have been the last character of the lot to be extensively studied, but in recent years their contribution to the spatial organization of the visual scene has been increasingly acknowledged (Mamassian, Knill, & Kersten, 1998).

In this article I shall primarily review not what the visual system *actually* exploits in cast shadows, but what it is there that it *could* exploit (although it may not be well equipped to exploit). The study concerns the various types of information that *can* be extracted from shadows by systems that are significantly like ours (including embodied artificial systems) in an environment that is significantly like ours. If the study is hence an exercise in ecological optics,² it is only partly such an exercise, for two reasons. First, not only geometrical information will be dealt with. A series of

¹ See the bibliography compiled by Pascal Mamassian, updated in 1999, at http://www.psy.gla.ac.uk/~pascal/shad_biblio.html.

² Modulo, that is, considerations à la Fodor and Pylyshyn (Fodor & Pylyshyn, 1981) about the sort of properties that it is reasonable to expect that the system extracts, given the particular constraints imposed by the make-up of the visual system itself and by features of the representations it uses. The study is ecological also insofar as some features of objects can be read off their shadows just because our environment is such and such (for instance, because the angular size of the Sun is about ½ deg). See also (Pylyshyn, 1999).

important works (see (Knill, Mamassian, & Kersten, 1997) and the bibliography therein) documents the wealth of geometrical approaches to shadows. Here geometry will be put in a broader perspective, which includes non-geometrical facts. Second, the possibility will be left open that there may exist robust contingencies in the environment to provide low-level cues that trump whatever sophisticated rules for extracting invariants one can think of.

Methodological caveats

A few methodological points are in order. First, a distinction between automatic and controlled operation of the visual system will be presupposed. This may be not particularly momentous for the study of information, but it will be convenient to make use of the distinction at some crucial points. Second, we should be aware that shadows are themselves projections and that at the same time their image is projected onto the retina. Some of the complications due to this double projection can be left out. Finally, in the course of the presentation some hypotheses will be offered that may be worth studying experimentally.

Terminology

Standard terminological items, such as ‘cast shadow’ or ‘attached shadow’ will be taken for granted. However, unless otherwise specified, ‘shadow’ in what follows will be generally coextensive with ‘cast shadow’. To further fix the terminology, we call the object that casts a shadow a ‘shadow-caster’, shortened as ‘caster’, and the surface upon which the shadow is cast the ‘screen’. It usually takes three to produce a shadow;³ in addition to the caster and the screen, a source of light is required (here for short called ‘source’). To avoid ambiguities, we shall say that the caster *casts* a shadow, and that the source of light *projects* a shadow. In the text to follow, italicized phrases will in the norm used for describing some aspects of the informational interest of shadows.

³ The concept of a shadow is not very clearly demarcated. (van Fraassen, 1989), ch. 9, indicates that it may be uncertain, in some cases, whether a certain shadowed area is a shadow of any object. (Casati, 2000) points out that we may even to make arbitrary decisions about whether to call some areas ‘shadows’ or ‘spots of light’. However, for the present purposes these conceptual indeterminacies are of little import. The basic idea that it takes three to produce a shadow can be challenged also on geometrical bases; local, shadow-like deficiencies of light that look like a shadow can be produced by cunningly arranged reflections, without a caster.

What shadows are about

A shadow can inform about any of the three main actors – source, caster and screen. The information conveyed is of a different sort in each case. A cast shadow first and foremost informs about the *presence* of the three actors. Secondly, about their *location*. Yet more specifically, it can inform about some properties of each of the three actors: the *width* or the *intensity* of the light source, the *shape* and *slant* of the caster, the *texture* of the projection surface, or other elements yet. In all these cases, it is not claimed that shadows can inform about the relevant properties without further assistance from other elements of the scene. Reverse optics, as well as interpretation of shadows at large, is a complex exercise and requires a delicate equilibrium between available online information and background constraints.

Shadows are projections, *more or less so*

Mathematically, cast shadows are projections [fig. 1]. Consider a square casting its shadow onto a screen underneath. If the source of light is at infinity and rays are parallel, we have a metric projection whenever the square lies in a plane parallel to the plane of the screen (all relevant⁴ geometric properties of the square are preserved by its shadow, fig. 1a), as a limit case of an affine projection whenever the square does not lie in a plane parallel to the plane of the caster (parallelism of sides preserved, metrics not preserved fig. 1b and 1c). If the (pointlike) light source is moved close to the square – which in turn is parallel to the screen – we obtain a similarity projection (the shadow is just a larger square than the caster, [fig. 1d]). If keeping the source at a finite distance we tilt either the square caster or the screen by keeping one side of the square parallel to the screen we have a more generic central projection ([fig. 1e]); and if we tilt either the square caster or the screen by keeping no side of the square parallel to the screen we obtain an even more generic central projection ([fig 1f]: metrics and parallelism of the sides are not preserved, but cross-ratio is preserved, see [fig. 2]).

Most of these properties were known since Renaissance painters' theories and practices about shadows were systematized in the 16th century by mathematicians interested in perspective (Da Costa Kauffmann, 1993), and then widely and routinely

⁴ Non relevant properties here include the thickness of the square.

applied, especially to architectural drawings (Deutsche Architektur Museum, 2002). At bottom, perspective and shadows have the same underlying mathematics: in some conditions, cast shadows are just shading put in perspective.⁵

However, well behaved as it may be, the mathematical structure of shadows is only a partial guide to their informational structure. Sure enough, some intricacies can be formalized mathematically with (relative) ease. To give but two examples: (1) penumbra effects do not introduce new complexity, as the penumbra can be reduced to the superposition of many sharp shadows from distinct pointlike sources; (2) projection of shadows by oddly shaped objects or on oddly curved surfaces can be described mathematically in many circumstances, provided one can decompose the objects and the surfaces into sums of mathematically well behaved items (such as elemental bodies, as is customary in pictorial practice, revived in theories of *geons* (Biederman, 1987)). The point is rather that *not all* geometrical features of projections are interesting or useful from an informational point of view, even though they may be displayed by shadows, and that some non-geometric features are interesting as well. Accordingly, we should not confine ourselves to the mathematical properties of shadows, and among those properties we should take care of isolating what *can* be used by a visual system like ours – mainly a situated, movable system, which sees from a viewpoint – in an environment like ours – furnished with some adaptively crucial illuminants (Sun and Moon, the cloudy sky, some openings), and with a fairly recent, hence adaptively less relevant, proliferation of artificial light sources.

Basic shadow information: existence and number

Vision scientists are generally intrigued about geometric information conveyed by shadows. But shadows can inform about a much larger set of properties of the environment. Some of these appear trivial, just because they are so commonplace. To give a few examples:

Shadows can inform about *the presence of casters*. (In particular, they inform about the presence of objects *located outside the visual field* or about the presence of

⁵ So much so that it remains one of the largest oddities in the history of art why painters who were able to compute and depict perspective with high mathematical accuracy were not able to paint cast shadows correctly (Casati, 2004). This could have happened for various cultural reason, but also because the task is compounded: shadows are themselves perspective images, which are further depicted in perspective (foreshortened).

occluded objects.) At the same time, shadows can inform about the *presence of a source of light* (again, in particular, they can inform about the presence of a source located outside the visual field or about an occluded source). Finally, shadows can inform about the *presence of a screen*. (In the 1998 Peter Weir movie *The Truman Show* the shadow of Jim Carrey's hand revealed that what Truman thought to be the sky was only a decoy.)

Existence and number are of course related – insofar as you know about the existence of an object, you know that there is *one* object of a kind around. However, it may be harder to ascertain *how many* objects exactly are there, and this holds, in our case, for the three main actors (source, caster and screen). From the information that there are objects around you cannot infer without further ado *how many* objects there are.

Information about the source of light

Cast shadows can inform about some aspects of the position of a source of light.

Assume you can see at the same time two distinct shadows, cast by two poles, and projected by the same source. Within a certain range, that is, whenever the distal convergence of shadows is not under discrimination threshold, *the position – hence the distance – of a pointlike source* can be retrieved by the intersection of the directions of shadows. When the shadows appear (distally) parallel, the source is at a distance that is larger than a contextually dependent value. In such a case, shadows at least inform about *the direction in which the source can be found*.

Contrast at the shadow profile (modulo the uniformity of the screen's color) is an indication *of the intensity of the light source*. Difference in the luminance (“darkness”) of any two shadows from different sources is an indication of *the relative intensity of the sources*.

The presence of a penumbra gives an indication of the fact that *the light source is extended*, and to some extent – usually in a very crude way – *of the shape of the light source*.

Penumbra information

The penumbra may appear at a first glance a *sui generis* “fuzzy” phenomenon, but conceptually it is the result of taking the union of all the shadows of a caster, projected from each of the light points of an extended luminous body, and then subtracting their geometric intersection. *The presence of a penumbra informs about*

the fact that the light source is extended. Together with information about the distance of the shadow caster from the screen, the *width of the penumbra band informs about the angular size of the light source*. (Hering, 1964) suggested that the penumbra is a phenomenological hallmark of the shadow, distinguishing it from other types of visual items. However this is typical of natural illumination. Artificial sources may be pointlike and induce no penumbra effect. On the other hand, given the adaptive importance of the natural illuminants, it may be hypothesized that all too sharp shadows will not be recognized as shadows.

Information about the caster

As it has long been recognized by painters, a cast shadow is an *image* of the caster. Tales about the origins of painting involve a character's tracing the outline of a shadow cast by the sitter (Stoichita, 1997). The main element in this story gripping one's imagination is information about the caster, possibly because of the ecological significance of material objects. However, here lurks an intriguing "shadow paradox": in order for them to be used as images and to provide information about anything else from them, the visual system needs to take shadows *as shadows*. There are other options available: cast shadows could be taken as just dark patches, or, in particular conditions, in which recognitional abilities as of an object of type O are elicited, shadows can be mistaken for the object O itself. (The latter may be termed "*trompe-l'oeil* shadows".) In either case, the contribution of shadows⁶ to the reconstruction of the spatial structure of the scene would be different from that which is in place when the visual systems treats them as shadows.⁷

But how much of the caster can be captured in its shadow?

Shadows are profiles: Terminator information

Attached shadows provide one of the most essential cues to the local surface features of an object, the curvature of the object at the *terminator* line dividing light from

⁶ Both cast and attached, in fact.

⁷ The differences between the informational capabilities of shadows to represent their caster objects are independent of the issue whether shadows themselves are treated by the visual system as objects. One possible way of testing this hypothesis is to use experimental paradigms that distribute the attentional load across multiple objects. For instance, some of the targets in the Multiple Object Tracking paradigm (Pylyshyn, 1994) could be shadows of other visually displayed objects; shadows and casters would be visually disconnected from each other. Would the total number of trackable items increase?

shade. The normal to the surface at the terminator coincides with the normal to the direction of light, which is tangent to the surface at the terminator (Palmer, 1999) based on (Horn, 1974). Now, the terminator can in principle get projected in its entirety in a cast shadow. *The outline of the cast shadow is an image of the terminator.* To that extent, the shadow outline's informational content is at least hostage to the informational content of the terminator, which varies as a function of the position and orientation of the caster relative to the source (the terminator is not fixed: there is no "dark side" of the Moon, contrary to what is erroneously implied by many a toy in which the Moon is painted half black). As the notion of a terminator is a source-dependent notion, some terminators may be more informative than others, depending upon contingencies about the location of the source and the orientation of the caster, in the same sense in which some views of an object happen to be more informative than others (Tarr & Buelthoff, 1998). It follows that *some shadows are more "telling" than others.* Shadows of standing people projected on the ground by the noon Sun near the Zenith are less readable than the shadows of the same people projected at sunset on a wall, and among the latter some shadows will be more readable than others – silhouettes of a face's profile, where nose and chin are prominent, contain more information than silhouettes of a face's outline, where the ears and hair are prominent. In sum, terminators and profiles have varying informational goodness, and this variety is reflected in their shadow.

Other-perspectival information

Consider, now, what appears to be a basic geometrical-ecological axiom about shadows

(GA): *the position of a source cannot coincide with the subject's viewpoint.*

This axiom holds in a world like ours, where eyes do not typically emit light (think of the awkward flatness and near shadowlessness of pictures taken with a flash). From the fact that a shadow can inform about the presence of an occluded caster, that it informs about the caster's terminator, and from the basic geometrical-ecological axiom about shadow (GA), it follows that, whenever the caster is a 3d object,⁸

(1) The profile of a cast shadow never coincides with the *visually* accessible profile of the shadow-caster.

⁸ 2d objects such as sheets hanging from a rope may cast a shadow that matches their profile.

Consequence (1) further entails that the profile of a shadow is informationally rich in yet another sense: it can inform the viewer *about some aspect of the shadow-caster's shape that may not be accessible to the viewer*. This can be summarized by saying that to some extent the shadow *gives the viewer access to the viewpoint of the light source*. It makes it possible to guess what an observer could see of the caster if he or she was placed in the position of the light source. An indirect proof of the difficulty in disentangling the viewpoint of the source from that of the observer comes from a large number of paintings in which both the caster and the shadow are visible, and the visible profile of the caster perfectly matches the profile of the shadow.⁹ Most likely the painter first represented the caster, and later on computed online the shape of the shadow by a quick and dirty rule – proximal profile matching. Not only this tells us that painters do a lot of online computing when they draw; it is also an indication of the relative tolerance the visual system has for suchlike mistakes, provided a rough association of an object with a shadow can be established.¹⁰

Sub-terminator information: silence about the internal structure of an object

Informationally rich as they may be relative to some aspects of the caster, shadows appear to be silent about its internal structure or about any surface structure which does not lie at the terminator, if the caster is an opaque object. A shadow profile is the projection of the terminator of the shadow caster; everything optically within the terminator gets projected as a structureless dark patch. (Different is the case for semi-transparent objects, whose internal parts may let some light pass, thus unveiling some sub-terminator information.)

Sub-terminator information: 3d unfolding of objects and the fundamental viewpoint ambiguity

Although shadows are portions of surfaces, under conditions they can easily be read three-dimensionally (Wallach & O'Connell, 1953). In particular, moving shadows of moving casters can make it possible to retrieve the 3d unfolding of the caster. It looks as if the visual system relies on a “rigid body assumption” in performing this task.

⁹ As an instance, consider a very early depiction of a shadow in an illumination by Belbello da Pavia in the Visconti Hours, *Spies of Jericho Escape* (Florence, Biblioteca Nazionale) (Kirsch, 1972) where the profile of the shadow matches the profile of the figures as we see them, and not the profile of the terminator.

¹⁰ A similar indiscriminating tolerance has been documented for the depiction of mirrors and mirrored objects (Bertamini M., 2003).

The temporal unfolding of the shadow appearance is compatible with its being the appearance of a rotating, rigid body, and this interpretation dominates possible competitors, such as that as of¹¹ a constantly deformed rubber band.¹² In these cases the shadow itself is seen as having a 3d structure – that is, it isn't seen as just being the shadow of a 3d structure.

Consider now a cubic frame, casting a Necker-cube-like shadow. The shadow can correspond to two different positions of the cubic frame relative to an observer. The Necker-cube as seen in the shadow will appear “as from above” or “as from below” [Fig. 3]. If the cubic frame caster of the Necker-cube-like shadow rotates, the shadow itself will rotate in a way that leaves the Necker ambiguity open (it will be seen as rotating either clockwise or counterclockwise, with sudden turns). Interestingly, only one of the two readings of the shadow is correct, but this fact appears to escape automatic recognition by the visual system. Only one of the two interpretations (as from above, or as from below) is correct because the shadow witnesses the viewpoint of the light source – if the light source is above the cubic frame, the correct reading is “as from above”. The “as from above”/“as from below” ambiguity generalizes to all cast shadows in which it is possible to read a 3d structure. It should be interesting to study subjects' preferences, if any, for either horn of the ambiguity, as well as the reasons for the preference. Be it as it may, the example illustrates the informational overdeterminacy of cast shadows. They not only hint at what the source “sees”, but also at what would be visible from the viewpoint of a source situated antipodically relative to the object and the screen.

Information about caster-screen distance

Shadows are cast across a distance. A *ceteris paribus* rule of thumb for reading distance of the caster from the screen consists in visually determining whether the caster is or is not in contact with its shadow. This is first and foremost a binary reading, as it tells us *whether the caster is or is not in contact with the screen*. Noise in the reading can be traced back to the particular viewpoint of the observer. Some parts of the caster may *appear* to be in touch with the shadow when they aren't; hence

¹¹ “As of” and “as” clauses are generally used to indicate the content of a representation – what the system takes it to be the case, independently of whether it is indeed the case.

¹² This holds for shadows of rigid bodies, which is possible to tell from shadows of non-rigid bodies, as shown by (von Fieandt & Gibson, 1959)(I am grateful to an anonymous referee for pointing out this reference).

only the judgment that the object is *not* in contact with the screen (because the shadow is not in contact with the caster) would be foolproof.

A finer reading, disposing of part of the noise, and paving the way to an analog (graded) judgment about the distance of the caster from the screen, would require finding out a salient point on the shadow and the corresponding shadow casting point on the object, such that the two points are in touch – in that case the caster touches the screen. If the two points are not in touch, the distance between them *provides a cue to the distance between the caster and the screen*. The cue will be in many cases indirect, as the line between the two points will normally be foreshortened on the image. It is nevertheless a powerful cue, that can in some conditions overrule competing information (Mamassian et al., 1998).

Penumbra information about caster-screen distance

In the adaptive environment, the dominant type of illuminant is either the Sun or the Moon.¹³ Both are so distant from the Earth that their rays can for many a purpose be considered as parallel to each other, and the projection of shadows affine.¹⁴ However, this simplification would only concern a pointlike faraway illuminant, whereas in the sky both Sun and Moon¹⁵ subtend an arc of about ½ deg. Hence part of the light emitted or reflected from any of two diametrically opposite points on any of the celestial body's visible circumference is bound to intersect at an angle with light emitted from the other point. As a consequence, a light cone opens up anytime a shadow caster is sunlit or moonlit, and the projection on the screen of the space within the cone is blurred, thereby creating a 'softening' effect that increases as the shadow casting point is more and more remote from the surface. That is, *the width of the penumbra informs about the distance of the caster from the surface*. Interestingly, this is a form of absolute – non-relative – knowledge, anchored to the absolute environmental datum of the angular size of the Sun in the sky.

¹³ The luminosity of the Sun at Earth is about 450.000 times the luminosity of the Moon; shadows projected by the Moon when Sun and Moon are both present in the sky are canceled out by shadows projected by the Sun.

¹⁴ It would be interesting to test whether this strong environmental feature translates into a default processing of shadows as if they were affine projections.

¹⁵ Venus, the third brightest ordinary celestial body, casts very feeble shadows, sometimes visible in a moonless night. Venus can be considered pointlike for vision, thus adding interesting variance in the population of astronomical illuminants, but the paucity of its shadows is a sufficient reason not to study them.

A limit case of direct access to the caster-screen distance

There is yet another, more direct way of inferring distance, in a limit case. Again, as the Sun's angular size is constant, a pinhole in a caster intercepting light will reconstitute an image of the Sun (cut within the shadow of the caster) whose size is in direct proportion to the distance of the pinhole from the projection surface. This is apparent in images of the Sun projected through trees' foliage, where pinholes frequently occur at the superposition of different leaves. (If r is the radius of the image of the Sun, the distance is $(r \tan (30''/2))$ [$\tan 30'' = 0.00872$]) It would be interesting to see if any sensitivity to such a strong environmental feature can be documented in the visual system.

Information about the screen

At present, this is still an ill-studied area (Knill et al., 1997). Obviously, the final shape of shadows owes much to the configuration of the screen. A shadow cast upon a flat surface will look pretty different if cast over a flight of stairs. The difference can be used to *infer the geometry of the screen*. All the more so if the shadow is moving. Assuming by default that caster and source are kept constant,¹⁶ *changes in the moving shadow inform about the local geometry of the screen*. One can picture the situation as that of shadow that in moving over a surface appears to bend, shrink or expand, signaling edges, hollows or bumps, respectively. However, anecdotal observations suggest that perception may not be especially tuned to this type of information and may simply stop at the consideration that the shadow behaves in a strange way.

Making screen features visible

When it comes to the screen, it turns out that depending on lighting conditions, shadows can hide as much as they can reveal.¹⁷ Some marks on a surface can be hardly visible because of strong illumination, which reduces the contrast between visual features. A shadow can *reduce the contrast and make the features visible*. Even perception of the color of a surface can be enhanced by the presence of a shadow cast across it.

¹⁶ Except, that is, for the movement of the caster.

¹⁷ J.M.Kennedy, personal communication.

Shadow constancy

Expanses of color can be seen as constant under a variety of local differences in luminance: for instance a white case on a checkerboard maintains its whiteness in spite of being partially shadowed. Does the dual phenomenon occur? Shadows are – at some level – expanses of color, and they may cross borders, such as the border between a black and a white case. Is there any shadow constancy? To my knowledge, there is no discussion in the literature of the conditions that underlie the phenomenon of shadow constancy – the phenomenon itself does not appear to have been acknowledged. Adelson's checkerboard (Adelson, 2001) is usually presented as an instance of color surface constancy, which is not impaired by the presence of shadows, that modifies local luminance. However, it is as well an instance of shadow constancy, which is not impaired by the change in surface, that modifies local luminance as well. (There is possibly some relationship with the conditions under which a transparent film is seen as of constant color in spite of local difference in luminance (Metelli, 1974).)

The informationally relevant point here is that *shadow constancy* can provide information about *the opaque structure of the caster* (constancy being here an indication of evenness).

Digression on stretching the limits of ecological perception: astronomical geometries

The visual system exploits shadow properties in the automatic mode; however, visual information can be accessed in a controlled mode and made an explicit premise of inferences whose conclusion is an assessment of some properties of an object. (As the example of the pinhole shows.)

Indeed, pre-telescopic astronomy undertook the study of shadows quite seriously. Celestial bodies are beyond the limits of ecological perception and but their most basic visual properties (luminance, orientation) can be processed automatically. Reasoning in a controlled way upon celestial bodies appears to be biased by the construction of elemental mental models under incomplete information (Vosniadou, 1990). Shadows are been instrumental in controlled processing of information about celestial bodies (Van Helden, 1985).

I shall mention only three instances. (1) from the study of cast shadows one can infer that The Moon is smaller than the Earth. To begin with, from the fact that the Sun and

the Moon are very close to being on two opposite sides of the Earth at full Moon, *and* that there is an eclipse of the Moon only rarely, one can infer that the Earth is much smaller than one might have thought relative to the Moon (by comparing the two visually), and quite distant from both Sun and Moon. A very large Earth would eclipse the Moon more often than it does. Now, the Earth is small, but not so small. Drawing on the fact that during an eclipse of the Moon the terminator of the Earth's shadow is less curve than the profile of the Moon, Aristotle inferred that the Moon is smaller than the Earth (relying on the further assumption that the Sun is much farther away than the Moon than the Earth)(Casati, 2003). And the Earth is large indeed: the flat horizon of the Earth at sunset during an eclipse is a tiny portion of a terminator that gets projected onto the Moon as a curve (Kennedy, 2002). (2) Another example is Galileo's 1610 method to find out about the height of mountains on the Moon (Galilei, 1957). Galileo ingeniously inferred the height of some lunar mountains by measuring the distance of their tops from the terminator, when the tops were visible as points of light in the shaded part of the Moon (at their sunrise or sunset). (3) Yet another example is the confirmation of Huygen's hypothesis that Saturn was surrounded by rings by Campani in 1664, who observed that rings cast a shadow on the planet, which in turn casts a shadow on the rings.

Noise in shadow information: the case of number

Given the number and the complexity of the factors that make it possible for a cast shadow to convey information about source, caster and screen, it is to be expected that noise can enter the picture at many stages and significantly degrade the information. We have seen some examples of such degraded information: uninformative terminators, spurious contacts between an object and its shadow. To these one may add the unevenness of surfaces, especially as these may get partially occluded at the shadow's position – so that the shadow can be visually discontinuous, split across distant portions of the visual field.

Consider, as an example of the latter problem, the question of how shadows can inform about the *number* of sources, casters and screens. Conditions have to be favorable. If they are indeed favorable, a rule of thumb for counting casters may sound: each self-connected, maximal shaded area corresponds to a (self-connected, unitary) caster, so that if the shadow is unitary, the caster is unitary, it is one. But the

rule is easily violated in unfavorable conditions. A unitary caster can spread its shadow over two split surfaces, and conversely two casters' shadows can merge into a self-connected, maximal shaded area. Or again: two light sources can be so arranged that a single object casts two shadows. However, one can to some extent triangulate: knowing how many screens there are, and how many sources, *the number of shadows informs about the number of casters*. This reminds us of the highly contextual information provided by shadows – something else must be known in order for shadows to work reliably as informants. Reverse optics and shadow interpretation is hardly successful in highly unconstrained cases.

Conclusions

In this paper I have argued that cast shadows contain information that can be processed either automatically or in a controlled way by the visual system. The type of information conveyed by shadows can be binary (presence or absence of casters, contact or lack of contact, unity or lack of unity)¹⁸ or else continuous or analog (shape, distance). In all these cases, it was not claimed that the visual system is sensitive to this type of information, although in many cases it surely is, and although it may turn out that at some yet unexplored levels of automatic information processing the system will be discovered to take into account the informational riches of cast shadows. I suggested some open questions and hypotheses about shadow perception:

- Is penumbra is a mandatory phenomenological mark of shadow?
- Does the strong environmental feature of the sun at (virtual) infinity translate into a default processing of shadows as if they were affine projection?
- Are some shadows actually perceived as more telling than others?
- Are there conditions under which shadows are mistaken for the real thing?
- Are shadows treated by the visual system as full-blown objects?
- Is there a preference for any of the as-from-above/as-from-below aspect of shadows?
- Is it difficult to disentangle the viewpoint of the source from that of the observer?

¹⁸ Research for this paper was made possible partly by a Cognitique Grant of the French Ministry of Education, ACI Cognitique 2000 2P3470, "The role of recognition in pictorial representation". Thanks are due to Nicolas Bullot, Pascal Mamassian, and to two anonymous referees for helpful comments on an earlier version of the paper. Thanks to the late Paolo Bozzi for many inspiring conversations.

- Can a sensitivity to strong environmental constraints such as the constant angular size of the Sun be documented?
- Can the phenomenon of shadow constancy be documented and does it obey idiosyncratic principles?

These questions may be worth an experimental try. The present approach will be incomplete until the hypotheses will be made operational.

References

- Adelson, E. H. (2001). On Seeing Stuff: the Perception of Materials by Humans and Machines. In B. E. Rogowitz & T. N. Pappas (Eds.), *Proceedings of the SPIE: Human Vision and Electronic Imaging VI* (Vol. Vol. 4299, pp. 1-12).
- Bertamini M., H., H., Spooner, A. (2003). Naive optics: Predicting and perceiving reflections in mirrors. *Journal of Experimental Psychology: Human Perception and Performance*, 29(5), 982-1002.
- Biederman, I. (1987). Recognition-by-components. A theory of human image understanding. *Psychological Review*, 94, 115-147.
- Casati, R. (2000). *The Shadow Club*. New York: Knopf.
- Casati, R. (2003). The availability of large size from shadow: looking for hidden assumptions. *Perception*, 32, 1021-1023.
- Casati, R. (2004). Methodological issues in the study of the depiction of cast shadows. *Journal of Aesthetics and Art Criticism*(Forthcoming).
- Da Costa Kauffmann, T. (1993). The perspective of shadows, *The Mastery of Nature*. Princeton: Princeton University Press.
- Deutsche Architektur Museum, D. (Ed.). (2002). *The secret of the shadow: light and shadow in architecture*. Tübingen: Wasmuth.
- Fodor, J., & Pylyshyn, Z. (1981). How direct is visual perception? Some reflections on Gibson's 'ecological approach'. *Cognition*, 9, 139-196.
- Galilei, G. (1957). *Discoveries and Opinions of Galileo* (S. Drake, Trans.). New York: Anchor.
- Hering, E. (1964). *Outlines of a theory of light sense*. Cambridge, Mass.: Harvard University Press.
- Horn, B. K. P. (1974). Obtaining Shape from Shading Information. In P. H. Winston (Ed.), *The Psychology of Computer Vision* (pp. 115-156). New York: McGraw Hill.
- Kennedy, J. M. (2002). Eclipse, shape from shadow, and perceiving astronomic size - Earth's'. *Perception*, 31, 1027-1029.
- Kirsch, E. W. (1972). *The Visconti Hours, facsimile edition*. New York: George Braziller.
- Knill, D. C., Mamassian, P., & Kersten, D. (1997). Geometry of shadows. *Journal of the Optical society of America*, 14(12), 3216-3232.
- Mamassian, P., Knill, D. C., & Kersten, D. (1998). The perception of cast shadows. *TICS*, 2, 288-295.
- Metelli, F. (1974). The perception of transparency. *Scientific American*, 230(4), 91-98.
- Palmer, S. (1999). *Vision Science*. Cambridge, Mass.: MIT Press.

- Pylyshyn, Z. (1994). Some primitive mechanisms of spatial attention. *Cognition*, 50, 363-384.
- Pylyshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22, 341-423.
- Stoichita, V. (1997). *A short History of the Shadow*. London: Reaktion Books.
- Tarr, M. J., & Buelhoff, H. H. (1998). Image-based object recognition in man, monkey and machine. *COGNITION*, 67((1-2)), 1-20.
- van Fraassen, B. (1989). *Laws and Symmetry*. Oxford: Clarendon Press.
- Van Helden, A. (1985). *Measuring the universe*. Chicago: University of Chicago Press.
- von Fieandt, K., & Gibson, J. J. (1959). The sensitivity of the eye to two kinds of continuous transformation of a shadow pattern. *Journal of Experimental Psychology*, 57, 344-347.
- Vosniadou, S. (1990). Conceptual development in astronomy. In S. M. Glynn, Yeany, R.H., Britton, B.K. (Ed.), *The psychology of learning science* (pp. 144-175).
- Wallach, H., & O'Connell, C. (1953). The Kinetic Depth Effect. *Journal of Experimental Psychology*, 45(205-217).

A square caster casts its shadow onto a screen. Left column, plane of the caster parallel to the plane of the screen; middle column, square tilted with two sides of the square parallel to plane of screen; right column, square tilted with no sides of the square parallel to plane of screen. Top row, source at infinity; bottom row, source close. Projection of the shadow is metric (a), affine (b-c), similar (d), projective (d-e).

Fig. 3

Necker-cube shadow (top), and two possible interpretations (bottom).

Left, “as from above”. Right, “as from below”. The portion of the frame closest to the observer is highlighted.