

(Forthcoming in *Perceptual Ephemera* eds. Crowther & Mac Cumhaill, Oxford University Press. This is an uncorrected draft – please do not quote from this!)

ON SILHOUETTES, SURFACES AND SORENSEN

THOMAS RALEIGH

In his book “Seeing Dark Things” (2008), Roy Sorensen provides many wonderfully ingenious arguments for many surprising, counter-intuitive claims. One such claim in particular is that when we see a silhouetted object – i.e. an opaque object lit entirely from behind – we literally see its back-side – i.e. we see the full expanse of the surface facing away from us that is blocking the incoming light. Sorensen himself admits that this seems a tough pill to swallow, later characterising it as “the most controversial thesis of the book” (2011, p199).

I will argue against Sorensen’s controversial thesis and in favour of what seems to me to be a much more natural and commonsensical alternative: when we see a silhouetted object, what we see is its edge and only its edge – so we do not see its entire back-side. More precisely, we can see an outline of the object, defined by where the limits of the light-blocking object meet an area of surrounding space which is allowing light through towards our location. (I will say something more about what is meant by ‘edge’, ‘surface’ and ‘outline’ in section 1, below.)

My conclusion then is not meant to be very surprising. Nor can I hope to emulate the elegance and originality of Sorensen’s argumentation. However, I do hope to show that my position, unlike Sorensen’s, has the virtue of being true. And in the course of defending my somewhat uninteresting conclusion, I hope to touch upon one or two interesting points and principles about causation and vision.

(1) PRELIMINARIES

To begin, I need to say something about how I will be understanding the key terms ‘edge’, ‘outline’ and ‘surface’.

Both surfaces and edges can be thought of as varieties of boundaries or limits of objects. Mathematically speaking, a surface is typically a 2-dimensional boundary of a 3-dimensional solid¹. An edge is often defined as a 1-dimensional line between two vertices, which forms some or all of the boundary of a polygon, or forms a boundary where two 2-dimensional faces of a higher-dimensional solid meet. On this standard usage then, a 2-D square has 4 edges (sides), whilst a 3-D cube has 12². However, it will be convenient to use ‘edge’ in a more liberal, inclusive sense to speak of *any* line on the

¹ There are also more exotic 2-D surfaces that cannot be embedded in 3-D space – such as the surface of a ‘Klein Bottle’.

² This definition doesn’t apply quite so neatly to circles, which could be said either to have no edges or to have infinitely many!

surface of a 3-d object. For example, whilst a knife's blade may have one especially salient edge – which we might naturally speak of as '*the* edge' – for present purposes any arbitrary line on the surface of the blade would also count as an edge. This liberal usage is, I take it, in line with the talk of 'edge-detection' that occurs in vision science and artificial vision – an edge of an object in this sense need not be any sort of sharp discontinuity in the shape of the surface; it can be *any* line on its surface where the object ends and the surrounding space begins *relative to the perspective of the viewing subject*. E.g. my visual system can detect the 'edge' of a perfectly smooth sphere I am seeing. This is just a line running along the surface of the sphere, which is determined by my viewpoint on it and which marks the visually salient contrast between the object and its background.

The 'outline' of an object, as I am using the term, is just a kind of edge, or some edges, in this liberal sense. It is the part/parts of the object where the object ends, immediately next to the adjacent space/medium, defined relative to some specific viewpoint on the object. So if I am looking at the sphere from one viewpoint and you are looking at it from a viewpoint precisely opposite on the other side of the object, we can see different portions of the sphere's surface, but we are seeing the same edge or outline of the object.

Of course, how exactly the notions of surfaces and edges as boundaries/limits of 3-D solids apply to, or are related to, actual physical objects, composed of matter, of scattered molecules and atoms, is an interesting and tricky question. Stroll (1988) distinguishes between two broad conceptions of surfaces: on one they are treated as abstractions ('A-surface'), on the other they are treated as physical entities which are parts of other physical entities ('P-surface'). Stroll writes:

'...as I am using these terms, it is a necessary condition of something's being a physical object that it have some mass or bulk, and thus that it occupy some space; whereas it is a necessary condition of something's being an abstract object or entity that it lack bulk and therefore cannot displace space.' (Stroll, 1988, 40)

Stroll takes it to be evident that it is only on the conception of a surface (or an edge³) as a portion of bulky matter that it could be a candidate object of perception:

'On the assumption that there are abstract objects, it is an essential feature of such objects that they cannot be seen, or physically manipulated in any way. For our inquiry here, the important inference to draw is that if we wish to describe what is seen when one is said to be seeing a surface, or if we wish to describe the conditions that must be satisfied before one can be said to be seeing a surface, it is clear that we must be talking about a P-surface rather than an A-surface... That is, we have to be talking about something that in principle can be scratched, waxed, polished, or dusted, be rough, pitted, damp, and so forth.' (ibid. 80)

³ Stroll draws the same distinction between an abstract conception and a substantial physical conception for edges in the final chapter of the book – see pp196-199.

Following Casati & Varzi (1999), I will assume that a boundary of an object (whether a surface or an edge) is a *part* of the object, rather than being a property of the object⁴.

‘The choice between a boundary-based and a boundary-free mereotopology can hardly be eluded... Boundaries seem needed in order to account for certain basic intuitions – to do justice to the ordinary concept of a surface, for instance. But they also give rise to deep puzzles that a boundary-free theory would escape altogether. Eventually we shall opt for the first attitude, treating boundaries as *bona fide* spatial entities along with ordinary objects.’ (Casati & Varzi, 1999, 71)

Sorensen likewise takes surfaces to be genuine parts of objects. And, echoing Stroll’s distinction between an abstract and a physical conception of surfaces, he is very clear that he is treating the back-surface as a *substantial material part* of a material object:

‘A surveyor is free to concentrate on the abstract surfaces of objects: he is interested in geometrical relations. A vision scientist must concentrate on concrete surfaces because only they satisfy the causal requirements for seeing...

The scientist will go to whatever depth is necessary for the action. But he will go no further. For opaque objects, this means he stops early.’ (Sorensen, 2008, 45-6)

‘In frontlit conditions, we see an object by virtue of the light transmitted by its front layer. In backlit conditions, we see the object by virtue of the light blocked by its back layer. A layer blocks by a different mechanism than it reflects. Consequently, the surface of a backlit object consists of the portion of the object that is just enough to block the light. This absorptive surface must be distinguished from the reflective surface.’ (ibid, 48)

For Sorensen then, the part of the object that we see in back-lit conditions is some substantial portion of its matter – in the case of the silhouetted moon, it will be a massive concave bowl-shaped portion of its constituent rock.

In contrast, Casati & Varzi treat boundary parts (or ‘tangential parts’ as they sometimes call them) as *different* from extended or interior parts of the object, in that they ontologically depend upon the object for their existence and they do not take up space or have divisible bulk.

‘Our position is that boundaries are ontologically on a par with (albeit parasitic upon) extended parts. But unlike extended parts, spatial boundaries have a peculiar relation to space (just as temporal boundaries have a peculiar relation to time): they are located in space, yet do not take up any space.’ (Casati & Varzi, 1999, 5)

⁴ An example of a philosopher who took the opposite view, that a surface is property of an object rather than a genuine part, is H. H. Price: “‘Surface’ it is true, is a substantive in grammar; but it is not the name of a particular existent, but of an attribute.’ (Price, 1932, 106)

Moreover, they wish to remain sanguinely agnostic as to how boundaries, thus conceived, are related to the (micro-physical) material constituents of objects described:

“Our primary concern is with the ordinary conception of middle-sized reality... We can assume that it is an obvious and uncontroversial fact that ordinary physical objects (i.e. objects interpreted physically as aggregate of molecules) are not strictly speaking dense and do not have boundaries of any kind (at least, not boundaries of the smooth, continuous sort countenanced by our unreflective view of the world)... But we are not interested in the problem of how middle-size boundaries ‘emerge’ out of microscopic discontinuities... the physical framework drops out as largely irrelevant in the folk justifications we provide for the content of our perception and action. Mary decided to leave the highway because she thought that the road was bumpy, and she thought so because she saw it (and felt it) as uneven. This is the kind of description we use in our everyday justifications, and it is also the description of the world we need to rely on in order to explain the behaviour of a cat or to teach a machine... to move around and interact autonomously with the environment.

It is this intermediate ‘world’ that we are concerned with. We are uncommitted with regard to its status (whether it should be reduced to, supervene on, or added on top of an underlying physical reality).” (ibid, 73)

Casati & Varzi’s conception of surfaces (boundaries) then does not seem to neatly fit the concrete vs. abstract distinction that Stroll and Sorensen both invoke. The pertinent distinction here is rather between different ‘levels of description’, or perhaps between the manifest and the scientific image of the object. Speaking at the familiar everyday level of medium sized dry goods, Casati & Varzi wish to treat the boundaries of familiar objects as genuine (though ontologically dependent) physical parts of those objects located in the physical space of the environment, which can be perceived and which can figure in many kinds of causal explanations pitched at this everyday, medium-sized scale, but which do *not* take up (displace) any space or have any bulk. It is then a further question how these manifest features of manifest objects relate to the micro-physical constituents of objects qua swarms of scattered particles.

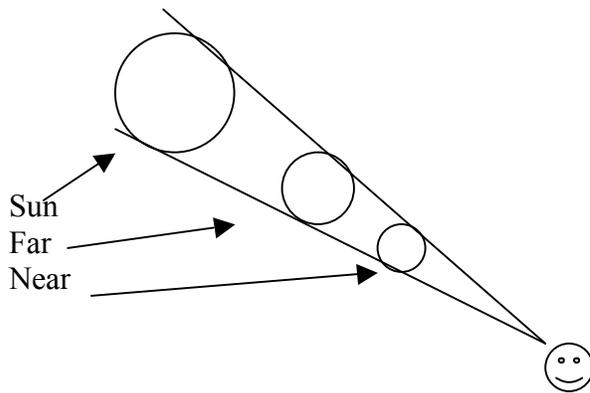
It is not quite so obvious as Sorensen implies then that in order to figure in the sorts of causal explanations that vision scientists are interested in giving, a boundary (surfaces or edges) *must* be understood to be identical with some substantial portion of the constituent matter making up the object. It might be that, at least sometimes, a perfectly good scientific explanation of what/how we see can be given which treats the surface or edge of the perceived object, a la Casati & Varzi, as a smoothly continuous boundary located out there in the environment’s 3-dimensional space, whilst remaining agnostic as to how this boundary is related to the object’s constituent particles and forces.

As far as the argument of this paper goes, it will not matter whether the surfaces and edges that we see are thought of, a la Sorensen and Stroll, as substantial portions of matter, or, a la Casati & Varzi, as dependent tangential parts – which are genuine physical parts of objects, but which do not have divisible bulk. So long as we are clear that they are not abstracta, but are real visible parts of things, located out there in the

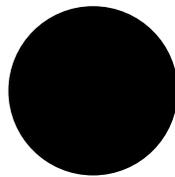
physical environment. However, I will briefly return to these differing conceptions in section 5, below.

(2) BACKS VS. BOUNDARIES

So why does Sorensen think we see the (entire) back-surface of back-lit things? His position arises out of his discussion of his “intersecting eclipses” puzzle (a variant of the “Yale shadow puzzle”). Suppose there is a double eclipse of the sun, as in figure 1:



From the observer's position, near and far have exactly the same relative size and both perfectly block out the view of the sun. The observer sees something like figure 2:



Sorensen asks: “Do I see Near or Far? Common sense answers that I see Near rather than Far.” (2008, 20) But Sorensen answers, against common sense, but on the basis of the widely endorsed ‘causal theory of perception’⁵, that we must see Far rather than Near. Very briefly, the argument is that Near is causally idle (with respect to my eyes) – it is neither reflecting light, nor emitting light but most importantly it is not blocking light either. Near is entirely contained within the shadow cast by far, so it cannot itself be casting a shadow. Whereas Far *is* blocking light and so is causally active. Add the plausible principle that “S sees O only if O is a cause of what S sees”, and we conclude

⁵ The classic statement of this theory is Grice (1961). Other influential endorsements of the theory include: Strawson (1974), Pears (1976), Peacocke (1979).

that we see Far not Near. (Sorensen also points out that seeing silhouetted objects is a form of ‘contrastive seeing’, which is a different and more ‘ancient’ species of perception than our more familiar, everyday kind of seeing. Thus we have some further, independent reason to think that our common-sense intuitions as to which object, or which *part* of an object is seen may not be reliable in the case of silhouetted objects.)

I am perfectly happy to allow that we see Far rather than Near, and also to allow the general claim that when we see the silhouette of O we are seeing O itself – i.e. we’re seeing O directly, rather than, say, seeing O’s shadow and so only seeing O ‘indirectly’. Where I want to draw the line is the further claim that when we see O silhouetted, *the part* of O that we are seeing is the back surface. How does Sorensen reach this further claim?

“When I see the silhouette of an object, I see the object. I see an object only if I see part of it. That part must cause my perception in an appropriate way. In the case of naked, opaque silhouetted objects, the only part that can play this causal role is the object’s absorption layer. Therefore, when I see a silhouetted object, I see its back surface.” (Sorensen, 2008, 50-51, underlining added)

Lets grant the assumption that to see an object one must see some part of it⁶. I deny that the only possible candidate part is the object’s entire back surface – a much better candidate is the object’s outline or edge (relative to the the viewer’s perspective). Of course, I do not deny that the back surface of the silhouetted object plays *some kind* of causal role in our seeing the object. But I will deny that it plays the “*appropriate*” sort of causal role that would amount to its being seen (to its being the part of the silhouetted object that is seen).

I think the initial steps of Sorensen’s reasoning in favour of the back surface might be made slightly more explicit as follows:

- (1) In back-lit conditions, we see O because it is blocking the incoming light (Sorensen’s phrase is “in virtue of O’s blocking the light”).
- (2) In back-lit conditions, the part of O that is blocking the light is the back surface of O.
- (3) So, in back-lit conditions, it is because of the back surface of O that we see O.

So far, so plausible. But we have not yet got to the claim that the part of O that we see is the back surface. Sorensen still needs a further premise roughly along the following lines:

- If it is because of some part of O that I see O, then I see that part of O.

However, as it stands this principle is too inclusive for it does not discriminate between different kinds of causal roles – any kind of causal “because of” would suffice for the part to be seen.

⁶ For a well-known argument against this assumption, see Clarke (1965). For a response to Clarke see Stroll (1988), chapters 5-6.

There are familiar problems concerning how to restrict the kind of causal chain that leads from an object to the subject's visual system so as to capture only the "appropriate kind" of chains, which amount to the subject perceiving the object, and to rule out the various kinds of deviant chains – e.g. such as cause veridical hallucinations⁷. And even once we have a kosher, perceptual causal chain, we still need a basis for ruling that it is *this* object that is seen rather than some other causally active item upstream or downstream in the same causal chain. It is important to notice that even on the assumption that we are indeed seeing O, very similar issues, concerning how to distinguish 'the right kind' of causal role, will crop up all over again when we consider *which parts* of O we see. I take it that Sorensen is alluding to these sorts of issues with his phrase "cause my perception *in an appropriate way*" (italics added), though he says nothing further as to what appropriateness consists in.

E.g. Suppose I am looking at an opaque sphere. And suppose further that some mechanism hidden in the centre of the sphere sends out radio waves which turn the lights on, or which cause the scales to fall from my eyes etc. Then, in one way, it is "because of" this inner part of the sphere that I see the sphere. But this would be a quite different kind of causal role to the role that the facing surface plays in reflecting light into my eyes – whereas the latter seems to be an 'appropriate' kind of role for something to be visible, the former presumably is not.

And it is not just a matter of ruling out "weird" or "deviant" cases. Even in "normal" cases we do not see everything that plays some sort of causally explanatory role in our seeing as we do. E.g. When looking at an opaque sphere in normal, front-lit conditions, the matter forming the core of the sphere may be causing the matter on the surface to have the shape it does (and so to reflect light in the way it does). But I do not see the core of O, though I'm seeing O looking as it does *because of* the core of O playing *some kind* of causal role. Again: this kind of role is not the appropriate kind. Thus the key premise needs to be amended somehow:

- If it is *appropriately* because of some part of O that I see O, then I see that part of O.

But what independent grip do we have on whether a part is playing the appropriate kind of role, apart from/prior to settling whether the part in question is visible?

Over the years, going back at least to Grice's 1961 'The causal theory of perception'⁸, a number of different prominent philosophers have attempted to provide Necessary and Sufficient causal conditions for seeing an object, which purport to weed out the deviant causal chains from the kosher chains⁹. But to my knowledge nobody has been concerned

⁷ Again, the locus classicus for this sort of discussion is Grice (1961). See also Lewis (1980).

⁸ Earlier works which discuss causal conditions on perceiving an object include: Chisholm's *Perceiving* (1957), and H.H. Price's *Perception* (1932).

⁹ E.g. Frank Jackson (1977), David Lewis (1980), Michael Tye (1982), Alva Noe (2003).

with the causal conditions that hold when we count as seeing, or not seeing, some *part* of O, given that we are indeed seeing O in the first place. So here is my stab at providing such conditions:

PPP (Parts Perception Principle)

Assumption: we assume that S is *consciously* seeing O¹⁰ – so O is causing S’s visual experience ‘in the right way’, whatever that way is and whether or not we can specify it in any neat, finite, non-circular manner. The following conditions then apply concerning whether some part p of O is seen by S.

- When S does *not* see p:
Either: Part p is only causally influencing O’s overall appearance to S by causally influencing the intrinsic appearance of some OTHER distinct part of O to S.
Or: p is simply not causally influencing O’s overall appearance to S at all.
- When S *does* see p:
Part p is causally influencing O’s overall appearance to S where such appearing does not necessarily require p to be causally influencing the intrinsic appearance of any other part of O to S.

Let me now unpack and illustrate this principle so as to make it seem a bit more intuitive. The basic idea that PPP is meant to capture is that the parts of an object which are currently visible to the perceiving subject are those that her visual experience is *directly* responsive to, whereas the unseen parts are those that are, at most, only *indirectly* affecting how the object looks to her. I assume, in line with our ordinary talk, that we can (at least sometimes) count as seeing an object, even though we do not see every part of it. Intuitively, I count as seeing the cricket ball even though I cannot see its inner cork core – I can only see the facing portion of the ball’s leather surface. And PPP delivers the intuitively right verdict as to which parts of the ball I see when I am looking at it in normal, front-lit conditions: the inner core of the ball can only causally contribute to the overall appearance of the ball if it somehow contributes to the appearance of some *other* part of the ball – viz: part of its facing surface. E.g. if the inner core is constantly changing shape slightly, this change will only be visible to me in virtue of causing the facing surface to change shape and so causally contributing to *the facing surface’s appearance* (to me). Whereas, there can be changes to part of the facing leather surface – e.g. a change of colour – which affect the ball’s overall appearance to me without requiring any other part of the ball to change its appearance (to me).

More generally, when S *does* see part p of O, even if there happen to be lots of complex causal dependencies and interactions between p and various other parts of O that are also seen, the idea is that we *could* have ‘blocked off’ all the other regions of S’s visual field except for the region that reveals part p of O – in which case *no* other parts of O would have been visually appearing to S at all, yet p still would have been causally contributing

¹⁰ So PPP cannot be used for cases of unconscious visual perception such as ‘Blindsight’.

to O's appearing to S via just the same causal channel. So p is causally contributing to O's appearing to S where just this same contribution *could have been* occurring without p's causally influencing how any other part of O appears to S. (Even if in fact p is causally influencing how other parts of O appear to S.)

Note: PPP speaks of the intrinsic appearance of the other parts, as the intrinsic appearance of one part of O can trivially influence the *relative* appearances of the other parts. Compare: if Tweedle-dee suddenly doubles in size, this (usually) leaves Tweedle-dum's appearance intrinsically unchanged, but of course it means that Tweedle-dum gains the new relative appearance property of 'looking smaller than Tweedle-dee' and loses the relative appearance property of 'looking the same size as Tweedle-dee'.

Now clearly some physical behaviours by a part of O would mean that the part *changes* from being unseen to being seen (or vice-versa). In normal front-lit conditions, part of an object that is not initially part of the facing surface could change shape or location in such a way as to become a new part of the facing surface. E.g. If the cork material in the core of the cricket ball expands so that it pokes out of the back of the ball and curves around to the side, this matter can go from being an invisible to a visible part, without any change of (intrinsic) appearance to the other parts of the ball. The part that pokes out now forms a *new* facing surface, a surface that can reflect light towards the subject's retina. This part then has gone from playing an indirect/dependent role to a direct/independent role in how the ball appears to S. In other words: just at the point where the cork part begins to stick out from the side and becomes visible, it *changes* from only contributing to ball's overall appearance by contributing to the appearance of some other part, to simply contributing to the ball's overall appearance by *itself* appearing some way.

Let me now mention a limitation of PPP – in telling us what is the case when a part p is or is not seen, the conditions make appeal to how *other* parts of O are appearing to S. So if we were completely in the dark as to which parts of O are or are not seen in the 1st place, we would not be able to use PPP as our guide for finding out. For in order to tell which condition p satisfies – seen or unseen – we'd have to be able to tell whether p was influencing the appearance of other parts of O, so we'd already have to be able to tell which other parts of O are seen. Perhaps this is a pretty serious limitation on PPP's usefulness for certain philosophical purposes. Nevertheless, even if this limitation means that PPP cannot serve as a useful *guide* when we are totally ignorant of which parts are seen, this in no way casts doubt on the *truth* of the conditions that PPP specifies.

As well as capturing our intuitions in normal everyday cases, PPP also delivers the intuitively correct verdict in the sort of cases, mentioned above, featuring more bizarre causal chains. E.g. suppose the inner core of the cricket ball is somehow transmitting a signal that casually affects the colour of the ambient lighting or the state of my eyes. So here the inner core is causally influencing the overall appearance of the ball by causing the surrounding lighting to be white light rather than red light, or by causing my visual system to go red-green colour-blind. Nevertheless, these influences by the inner core on the ball's overall appearance to me both still essentially involve causally influencing how the front surface of the ball *appears* to me, even though there may be no causal influence

by the inner core on the intrinsic physical nature of the front surface itself. So the influence of the inner core on the ball's overall appearance is still indirect, via influencing the *appearance* of some other part – so we do not see the inner core.

We are now ready, at last, to ask: how does PPP apply to cases of seeing silhouetted objects? Well, I think it tells in favour of the edge-theory against the entire-back-surface theory. For consider some part of the surface of a silhouetted object that does *not form* any part of the object's outline/edge (recall: this outline is determined by the viewpoint of the perceiving subject). For this non-outline-forming portion to causally contribute to the object's overall appearance to S it *must* make some causal contribution to the appearance of the object's outline (to S). So it must either causally contribute to how some *other*, distinct part of the object appears to S – viz, its outline, – or it must change or move in such a way that it itself comes to form a part of the object's new, changed outline or edge. Whereas, the outline/edge of the silhouetted object can causally contribute to the object's overall appearance to S without causally contributing to how any other distinct part of the object appears (to S).

When we see a silhouetted object, there is a crucial asymmetry between the outline and the (rest of the) back surface. Changes to the outline can be directly visible, whereas changes to any of non-outline-including portion of the surface are only visible insofar as they result in visible changes to the outline. We should conclude that the object's outline/edge is the only part of O that we see when we see O silhouetted. We do not see any part of the back surface (or front surface) that is distinct from the outline.

Whilst it is true that when O is silhouetted we see it because it is blocking light, it is not, of course, *simply* in virtue of its blocking light that we see O. After all, when locked in a totally dark room with blazing sunshine outside, you do not see the wall in front of you despite the fact that the wall is blocking light, and so causally contributing to your visual experience, or lack thereof¹¹. There needs, of course, to be a *contrast* – a surrounding area in which light is not blocked. Indeed, more generally an object needs to *contrast* with its surroundings in order to be seen. This is not a *totally* general requirement – e.g. when you are looking at a uniformly coloured expanse of the surface of a wall that totally fills your visual field, so there *just are no* surroundings to contrast with, it seems that you still count as seeing the wall¹². So the requirement seems to be that if an object has visible surroundings in the subject's visual field, then the object must form a contrast with those surroundings in order to be seen. Sorensen, of course, is well aware of the importance of contrast:

¹¹ Another of Sorensen's (2008) potentially surprising claims is that in such a pitch-black scenario we do *not* count as seeing *nothing*, for we can see the darkness.

¹² On the other hand, empirical results from psycho-physics suggest that if there is *absolutely* no change at all in the pattern of visual stimulus at the retina, after 2 or 3 seconds the subject experiences something like loss of any externally directed visual experience – see Hubel (1988) *Eye, Brain and Vision*, p.81. So perhaps you could only count as seeing an object that provides a *totally* uniform pattern of stimulus for a few seconds before lapsing into a lack visual perception.

“The universality of this reliance on contrast explains why any object can be camouflaged by muting its contrast with its surroundings. Seen objects normally make a *holistic* contribution to the scene. They are visible by virtue of the differences that they make with their surroundings.” (2008, 55)

So I think we should prefer the following line of reasoning to that which I attributed (above) to Sorensen:

(1*) In back-lit conditions, S sees O because of O’s blocking the incoming light *so as to form a contrast* with an adjacent area from which light reaches S’s viewpoint.

(2*) In back-lit conditions, the part of O that forms a contrast with an adjacent area of light is the edge or outline of O.

(3*) So, in back-lit conditions, it is because of the edge or outline of O that we see O.

PPP then provides the general theoretical rationale for taking the causal-role played by O’s edge to be the *appropriate* kind of role for that part of O to be seen. The causal role that the edge of O plays in my seeing O, unlike the role that the rest of the back surface plays, is the right kind of role to be a *visible* part, for the edge of O, unlike the rest of the back surface of O, makes a direct (i.e. potentially-independent) causal contribution to O’s overall appearance to S.

We have now considered Sorensen’s core argument for the ‘back-side theory’ and the general principles I have advanced against the back-side theory and in favour of the edge or outline theory. However, Sorensen provides a number of further considerations in favour of his view as well as providing some arguments against the edge-theory. And so I will also provide some further considerations in favour of my position and against Sorensen’s arguments. But first, I will briefly digress and say something about camouflage . . .

(3) CAMOUFLAGE

*‘You can blend-in in the country, you can stand out in the fashion world
Be invisible to a white tail, irresistible to a redneck girl
Camouflage, camouflage, oh you’re my favourite colour camouflage’
(‘Camouflage’, Brad Paisley)*

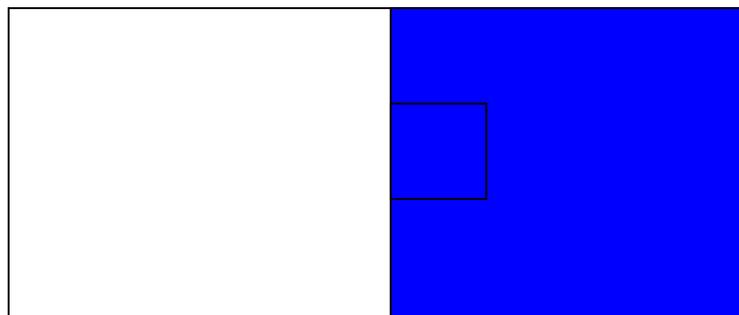
In a case of perfect camouflage, where S is *totally* unable to discriminate O from its surroundings, it is very natural to say that S simply does not see O¹³. Now whilst I can imagine ways of denying this, it seems clear at least that Sorensen is happy to accept this natural thought (see the quotation at the end of the previous section). But, assuming now

¹³ The classic statement in the recent philosophical literature of this sort of differentiation condition on seeing an object is Dretske (1969). Dretske makes the point already mentioned in section 2 (above), that this requirement is not *totally* general as one can sometimes count as seeing an object that entirely fills one’s visual field, for which there are thus no currently visible surroundings for it to be visually differentiated from.

for the sake of argument that in such cases of perfect camouflage we cannot see the object, O, can we nevertheless see O's facing surface? I think the answer must be that we *can* see at least part of O's facing surface, despite the fact that we cannot (lets assume) see O. Suppose O's surface has polka dots and zig-zags on it, but that O is perfectly camouflaged on a polka-dot zig-zag background. We can see particular dots and lines on O's surface – we can discriminate these parts from their surroundings – so we can see at least part of O's facing surface, though we apparently cannot see O. Perhaps it will be protested that this is not a case of perfect or total camouflage. But even when O's surface is a single uniform colour, perfectly blending with a uniform background colour, I think we still have to allow that we can see at least some of the facing surface of O although we cannot see O. For we might happen to attend to an area of (what is in fact) O's surface and think truly on the basis of what we see "That (arbitrary) region of the scene is blue just like the rest of the scene". Or we could be instructed to attend to the most central 1 cm² of the scene from our viewpoint, and the area we then happen to attend to just happens to be an arbitrary 1cm² portion of the perfectly camouflaged O's surface. If we are then asked "What colour is that region you are attending to?" – we could correctly answer 'blue'. Which strongly suggests that we must count as seeing that region of O's surface. Indeed it seems that in these sorts of perfect camouflage cases, we can see *all* of O's facing surface *except for* its edge or outline (the limit/boundary of the object where it meets with its surroundings, relative to the subject's viewpoint). The possibly surprising moral to draw then from these camouflage cases: whilst seeing O entails seeing some part of O, the reverse direction does not hold – seeing a part of O does *not* entail seeing O.

However, this raises a potentially interesting/puzzling point: whilst it seems plausible that to count as seeing an *object* one must be able to visually discriminate it from its surroundings (assuming that it does not fill one's visual field entirely), if I am right about the preceding camouflage cases, it seems that this requirement of discrimination does *not* apply to mere arbitrary portions of objects. For, I have suggested, I can count as seeing an arbitrary or indefinite region of the facing surface of O, even though I *cannot* visually discriminate such an arbitrary portion from its surroundings. So perhaps then some kind of implicit mereological theory is embedded in our concept of seeing. For it seems that what is required to see a "genuine" whole object – whose boundaries are, in some sense, real, mind-independent features – *differs* from what is required to see mere arbitrary regions or portions – whose limits are, in some sense, imposed or projected or undefined.

Actually, there can be cases of perfect camouflage where we can see *some* of the camouflaged object's edge or outline – as in fig 3, below.



Suppose that the small blue square here is perfectly camouflaged within the larger blue region, but is positioned such that one of its edges is perfectly aligned with the underlying boundary between the background blue and white regions. Here, one of the small blue square's edges *does* form a contrast with its surroundings, and so we can visually discriminate this portion of, what is in fact, the square's outline. However, despite being able to discriminate a part of it from its surroundings, we still (plausibly) don't count as seeing the small blue square, as we cannot discriminate *it* from its surroundings. (Even though S can visually discriminate what is, in fact, part of one edge of the small blue square, S is still completely unable to visually locate the small blue square anywhere within the larger background blue area.)

Cases of camouflage and cases of silhouettes are then opposites in the following respect:

- In cases of (front-lit) camouflage we can see *all* of O's facing surface *except* for its boundary/limit/edge. So we do not count as seeing O itself.
- In cases of silhouettes, we can see *none* of O's surface (front or back) *except* for its boundary/limit/edge. So we do count as seeing O itself.

We will briefly return to camouflage cases in section 5, below.

(4) SORENSEN'S FURTHER CONSIDERATIONS

I turn now to discussing two other considerations that Sorensen adduces in favour of his back-surface theory.

(i) "Seeing a surface does not require seeing much detail. Just before a rocket passes out of view, we see the rocket as a glint or speck in the sky." (Sorensen, 2008, 61)

The implication here is that when we see the glint/speck in the sky we count as seeing the front surface of the rocket. And so, in normal cases of seeing via light reflected from the object's facing surface, we can sometimes count as seeing this surface even though we are visually highly insensitive to the details of this surface – its colour, its bulges and declivities, etc. So likewise, Sorensen suggests, we can count as seeing the back surface of a silhouetted object even though we are visually highly insensitive to details of this back surface.

Firstly, it is not obvious that we have to accept this description of the case of the distant, front-lit rocket as one in which we can see its surface. When the rocket is a tiny glint I

think it is pretty natural to say the following sorts of things: “I can still see the rocket, just about.” “Can you see its facing surface?” “Well no, it is just a speck in the sky.” Or: “How much of its facing surface can you see?” “Well NONE, it is just a speck in the sky.” Again, it is *in some sense* because of the causal role of the front surface reflecting light into our eyes that we see the glinting rocket, but it is not at all obvious that it follows that we see this part of the rocket – the front surface – in this instance.

But whatever we want to say about the distant rocket case, the question is whether in the silhouette case it is better to say we see the entire back surface (though perhaps in poor detail or resolution), or that we see only the outline or edge? Well, if we think of a sliding scale of greater resolution, or visual acuity, as we get closer to the silhouetted object, or as the intensity of the light source behind O increases, *what* is it that we see in greater detail? With the front-lit rocket, as we get closer and closer, or as the lighting gets better and better, we can make out more and more detail of the *entire* front surface. But with a back-lit object, it is only the *outline* that we would see more and more clearly, the rest of the back surface would continue to provide just as little information as when the silhouette was a mere speck in the distance. (The back surface could be moving and changing shape in crazy radical ways and we would be totally oblivious so long as the outline is unaffected.)

Sorensen concedes much of this (speaking of the moon eclipsing the sun): “True, little of the far side [i.e. back-side] can be differentiated by purely contrastive seeing. Only the outer rim offers positive detail. The middle supplies only the negative information that there are no huge tunnels running straight through the moon.” (ibid, 60)

I think Sorensen here fail to properly distinguish genuinely *visual* information with what might be inferred on the basis of vision plus background knowledge. When I see a silhouette, the strictly *visual* “information” I get is that an object, or objects, make a certain outline (form a boundary with the surrounding light) that is shaped thus and so (relative to my position). Likewise I get the information that the object, or objects, do not form a tunnel in the middle *that is currently allowing light through to my position*. But seeing the moon silhouetted is quite compatible with the moon having a huge tunnel through it, so long as there is something else, farther back or in front, that blocks any light from coming through such a tunnel. Or it could just be that whilst there are sources of light that create a contrast with the edge of the moon, there is no source of light that illuminates the middle of the back of the moon, and so even though there is a tunnel right through it, no light comes through to the eye, and so this region simply looks like more black darkness.

Compare: When I am locked in a pitch-black room, do I *visually* gain the “negative information” that there is no hole in the wall I am facing? No. That is an inference you can only make, *given* background knowledge that the sun is indeed shining on the other side and so *would* come through any holes. But so far as *visual* information goes, you are just getting *nothing* about the wall. The wall may have all sorts of holes if it is pitch dark outside or if the room is contained within another larger room that blocks the sunlight etc. And likewise with the silhouetted moon: so far as any of the non-outline-forming middle

region of blackness is concerned, one is *visually* getting no information whatsoever, positive or negative. (Though of course a lack of visual information can itself be informative given the right background knowledge.)

The basic point here: given that the edge/outline is an alternative candidate to the entire back-surface for the part of O that is seen in back-lit conditions, Sorensen's concession that only the edge offers positive detail should reasonably be thought to provide strong evidence that the edge is a *better* candidate for the part of O that we see than the entire back side.

It is also notable that in a different chapter of the book, when elaborating a different ingenious puzzle-case involving a brick that perfectly fits into a cast shadow (which he calls "the disappearing act") Sorensen makes the following argument:

"While parked [in the shadow], the brick could grow yellow spots and spin. That would not affect what you see. The brick could contract into the shape of a bowling pin, wobble, and then further shrink into non-existence. Once again, what you see would not change. This low degree of functional dependence is due to the fact that the brick is causally idle. To be seen, an object must be a cause of what is seen. Therefore, the brick cannot be seen." (ibid, 66)

But what goes for the perfectly shadowed brick seems to go for the back surface in the silhouette case. Parts of the back surface can change shape in massive, radical ways, but *so long as* the outline (the boundary formed with the light) remains the same, these changes will be totally invisible as they will make no causal difference to the subject's experience. The way things look to you is functionally dependent only on the outline formed by the blocking surface, not on any of the rest of the back surface.

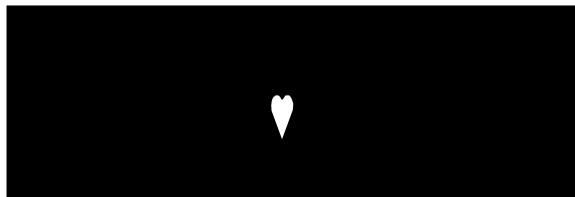
Of course, to repeat, the back surface of the silhouetted O is *in some sense* causally active in that it is blocking light and this blocking role is part of the total causal story of why the subject's visual experience is as it is. But, to repeat, *merely* blocking light that would otherwise be seen does not a silhouette make. (Recall the pitch-black room whose walls are blocking the bright sunlight behind.) And a part's playing *some* kind of important casual role in the object's getting seen, does not guarantee that the part is itself seen.

(ii) "Suppose a fluorescent balloon meanders in front of a solar eclipse. First, the glowing balloon blocks our view of the bottom edge of the moon, and then the middle of the moon. But if the balloon *blocks* our view of the middle of the moon, then we must previously have seen the middle of the moon. Therefore we are not just seeing the edge of the moon during the solar eclipse." (ibid, 50)

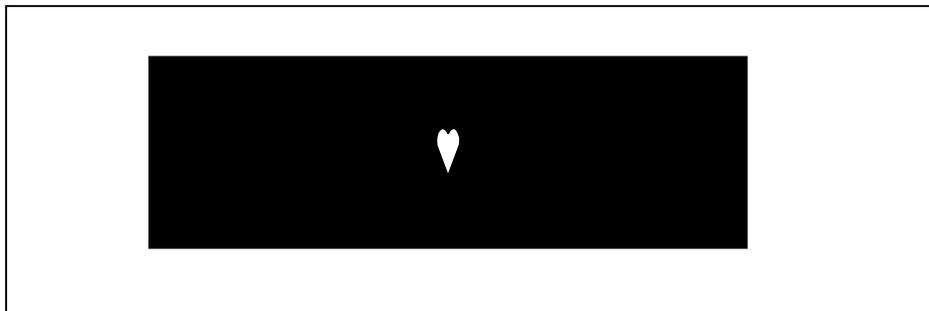
This passage strikes me as having some kind of initial, intuitively persuasive force. But I think much of this initial plausibility is due to the fact that Sorensen here no longer emphasises that it is the middle of the *back surface* of the moon that the balloon is supposed to block our view of. If we change the key sentence to: "First, the glowing balloon blocks our view of the bottom edge of the moon, and then it blocks our view of

the middle of *the back side of the moon*”, there is much less intuitive force; it becomes clear that Sorensen is soft-peddalling the controversial part of his theory. Moreover, this passage is not really much of an *argument*: It simply assumes that we do indeed have a view of the middle of the back surface of the moon for the balloon to block.

I think it is perfectly reasonable to allow that the balloon does block our view of the edge of the moon, but to deny that when it floats a little further it is then blocking our view of the middle (of the back) of the moon – for we simply have no such view. Here is an example to persuade you of this. Suppose the balloon is floating and glowing against complete darkness (figure 5):



One can see the balloon but there is no other surface one can see. Now suppose that in fact one is in a large pitch-black hall and outside there is bright sunlight. So you are in fact looking in the direction of a wall that is blocking outside light. Clearly we simply do not see the wall at all (neither its front nor its back surface). So, does the balloon here block your view of the middle of (the front/back surface) the wall? No. You have no view of the middle of the (front/back surface) of the wall. If you happen to know that in fact you are looking in the direction of the wall, then you can know that the balloon *would* be blocking your view of the wall. But the balloon is not currently blocking your view of the wall. Now why should it be different when one can see the silhouette of the wall (figure 6)?



Here one can see the edge of the wall, its outline, forming a contrast with the surrounding area of light. And so when you're looking at the glowing balloon, you realise you are

looking *in the direction of* the middle of the (front/back surface of the) wall. (You can see *that* the balloon is positioned in front of the middle of the wall.) But the balloon is still not blocking your view of the middle of the wall, for you still do not have any view of the middle of the front or back surfaces of the wall. You are still totally visually insensitive to this part of the wall. Once more, the front or back surfaces could be morphing and changing shape & colour in crazy radical ways – bits detaching and flying around etc – and one will be visually entirely oblivious *unless* these changes somehow affect the outline of the wall. You can see the wall by seeing the boundary it makes with the surrounding light relative to your position. You cannot see any other non-outline-forming part of the wall. So the balloon is not blocking your view of the middle of the wall's back surface nor the middle of its front surface. It is only when the balloon crosses the outline of the wall, the part of the wall that you *can* see, that the balloon partially blocks what you can see of the wall.

(5) SORENSEN'S EDGINESS ABOUT EDGES

Sorensen also provides a couple of arguments specifically targeted against my preferred position – that we see a silhouetted object's edge or outline.

His main argument against the edge theory is that edges considered as abstract limits are not suitable entities to do the appropriate causal work, but if they are treated as substantial portions of matter then the edge theory becomes just an unmotivated variant of the back-side theory:

“To lessen the appearance of a miracle, some people suggest that I am only seeing the edge of the moon during a solar eclipse. However, the edge can only block light if it is a physical object rather than an abstract mathematical limit. To be visible from Earth, that edge would have to be many meters wide. If I can see the backside of this giant ring of matter, then I would still be seeing a surface that is behind many kilometers of solid rock! Reducing the quantitative scale of the miracle secures no advantage . . .” (Sorensen, 2008, 50)

And in his later discussion of the perfectly shadowed brick “disappearing act” Sorensen repeats the contention that boundaries are too abstract to perform causal roles of visibilia:

“The only part of this thing that would not be blocked by the brick would be its boundary with the environment. This abstract surface cannot sustain the causal relations needed for visibility.” (ibid, 69)

I have two points to make in response here.

Firstly: even granting his assumption that the edge theorist must treat the edge/outline of the moon as a substantial ring of matter, Sorensen makes the further assumption that the *part* of this ring of matter that is seen is the backside of it – or putting it slightly differently, that the ring of matter the subject supposedly sees would be a smaller part of

the whole ‘absorptive’ back surface that is located *behind* a layer of rock which the subject cannot see. But why should an edge theorist be committed to this? According to the edge or outline theory what is seen is precisely the edge or outline, the part of the object where it forms a contrast with the surrounding area of light (relative to the subject’s viewpoint). If we need to identify a substantial portion of matter as the contrast-forming portion, this would *not* be a portion located behind some other distinct intervening portion, it would simply be a ring of matter located *right at the edge* (i.e. with no other distinct portion of matter either in front of or behind it, relative to the subject’s viewing position).

After all, Sorensen himself, the avowed back-surface theorist, is *not* committed to the idea that the part of the back-surface that is seen is *its* back-surface¹⁴ – rather, he holds that we simply see whatever portion of matter is ‘necessary for the action’, where the action which he takes to be crucial is blocking the light. The edge-theorist can accept this general methodological precept to seek whatever is necessary for the action, but the crucial action now is forming an outline with the surrounding light (relative to the subject’s viewpoint). So, to repeat, the edge that is seen, on this physical-lump-of-matter construal of ‘the edge’, would be the portion of the object’s matter that forms a visible outline contrasting with adjacent light.

Once we remove the extraneous additional commitment that one sees the back-side of the portion of matter, this version of the edge-theory is *not* equally counter-intuitive or apparently miraculous as the entire-back-side theory. For it is not committed to the idea that what is seen lies *behind* some other distinct portion of the object that is not seen – rather the outline-forming portion/portions of matter lies right on the edge of the object, neither in front of nor behind any other distinct portion/portions of the object (relative to the given viewpoint).

Secondly, one might question Sorensen’s assumption that an edge-theorist must understand the edge to be a lump of matter in order for it to figure in the relevant causal explanations of the subject’s visual experience. Notice that this seems to ignore the sort of position that Casati & Varzi wished to take (see section 1, above), on which boundaries such as surfaces and edges are features of the middle-sized manifest environment, located out there in space but which take up no space and are not lumps of

¹⁴ Sorensen rejects the following argument (given by Leonardo Da Vinci (2002) and discussed in chapter 2 of Stroll (1988)): a surface *must* be abstract, for whatever candidate portion of matter is proposed as being the surface one could always maintain that only the outermost part of this part really forms the surface. Thus we are led to conclude that the surface must be of infinitesimal thickness – i.e. abstract. He writes:

‘Leonardo’s whittling argument illicitly privileges immaterial surfaces. He manipulates the domain of discourse for the quantifier ‘outermost’. What counts as the parts of an object varies with purposes. Sometimes the parts include a coat of paint, the enamel, the peel, and so on; sometimes not. The domain is nondissective: part of parts need not be in the domain of discourse. If I say that the surface of my couch is plaid, I am not refuted by the observation that none of the small parts of the surface is plaid.’ (45)

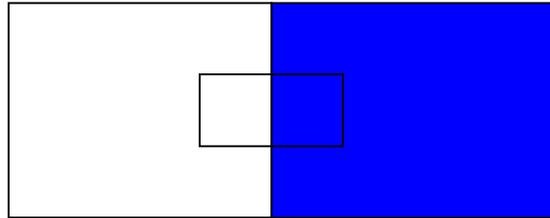
matter. These are held to be perfectly perceptible features in the environment and fit to figure in causal explanations pitched at the level of everyday medium-sized objects – though how exactly they relate to the microphysical world described by physics is left open. But then how ordinary middle-sized objects themselves relate to the leptons and quarks of fundamental physics is, *prima facie*, no less difficult and interesting a question. So far then, there is no obviously more pressing problem with the claim that we see the boundaries/limits of familiar objects – their edges and surfaces in Casati & Varzi’s sense – than with the claim that we see the familiar object itself in the first place.

I find the claim that only a substantial portion of the object’s physical matter is fit to play the relevant causal role especially surprising coming from Sorensen of all philosophers, who elsewhere in the book eloquently argues for a liberal view of what can be a cause, arguing that absences and privations – such as holes, cracks, shadows, silences – can all be causes. When we see O silhouetted we see a boundary between light and dark. Such a boundary is a perfectly mind-independent feature out in the world – though its position on the object is determined relative to the viewer’s location. There really is a line of difference out there in the environment where the boundary of the object that is blocking light from reaching the subject’s viewpoint meets the adjacent space/medium that is allowing light to continue on to that viewpoint. Now this boundary is not a standard 3-dimensional material object, but it is nevertheless a fairly non-mysterious environmental feature, located in 3-dimensional space. Such a boundary is not obviously any more “abstract” than an absence or a hole is. It does not itself absorb light, nor does it reflect or emit light. For it is formed/determined by a difference between 2 regions, one of which is occupied by light-blocking matter and one of which is allowing light through. But I cannot see any obvious reason to deny that such a boundary, such a mind-independent, locatable difference out there in the environment, can be a suitable cause of a difference in my visual experience. The boundary out there in the world, between where light can flow freely in my direction and where it gets blocked from reaching me, causes there to be a boundary in the proximal stimulus-pattern at my retina, which causes there to be an experienced boundary in the my visual field. True, part of the full causal story of why/how there is such boundary will be that the entire back surface of O is engaged in blocking light. But that provides no obvious reason why the boundary itself should not be counted as a genuine causal factor, influencing the subject’s visual experience. (Compare: part of the full causal story of why/how the whole back surface of O is blocking light, might be that the unseen matter in the inner core of O is holding the back surface in place. But this in no way undermines the back surface from playing a genuine causal role. Likewise then, the part played by O’s entire back surface in blocking light should not undermine the genuine causal credentials of O’s edge.)

Do we really want to accept that the boundaries or limits of things, where they start and where they stop, are *invisible* features because they are too “abstract” to do causal work? Such a position certainly does not square with ordinary talk about seeing the edges or ends of things, apparently quotidian features in our environments that we ordinarily take

ourselves to see. Nor does it seem to square with vision-science's robustly causal talk of "edge-detection" by areas in the visual system¹⁵.

Lets turn now to considering another case of camouflage (as in figure 7):



Suppose that the smaller, inner white and blue squares (X and Y) are both perfectly camouflaged against their respectively white and blue backgrounds. So I can see neither smaller square. However, the boundary between the two small squares *is* seen. I.e. I can see (what is in fact) the boundary, the difference between X and Y, without necessarily being able to see either X or Y. This boundary may not reflect light (or emit it or block it etc), but it is still a perfectly real mind-independent feature with a location in space, and it seems perfectly natural to say that the presence of this boundary out in the environment causes me to experience a boundary between white and blue.

Recall Sorensen's (entirely correct) emphasis on the holistic nature of vision (previously quoted in section 2, above):

"Seen objects normally make a *holistic* contribution to the scene. They are visible by virtue of the differences that they make with their surroundings." (p55)

We might add that seen features are sometimes visible in virtue of *being* differences (being constituted by differences) between adjacent parts of the scene. In the case of seeing O in silhouette, the part of O that is seen is the boundary/edge that is defined/constituted by O forming a difference with the surrounding space as to whether incoming light reaches the given viewpoint.

Of course in the above camouflage case, one *can* see what are in fact X and Y's facing surfaces (though one cannot see X or Y). Whereas in the case of seeing a silhouette, so I claim, one does not see O's back (or front) surface (except insofar as one sees the

¹⁵ Perhaps the seminal work on edge-detection within the primary visual cortex is Hubel & Wiesel's (1959, 1962) experiments on cats. Another important early paper is Shapley and Tolhurst's (1973) 'Edge detectors in human vision'. For a more recent overview of the psychophysical and physiological literature see Elder & Sachs (2004). And ever since the original 'Marr-Hildreth algorithm', see Marr & Hildreth (1980), there has been a huge amount of work in computational vision science seeking to provide algorithms that the visual system (or artificial visual systems) might use to perform the task of edge-detection. For one fairly recent comparison of some of the best-known of these algorithms see: Nadernejad et al (2008).

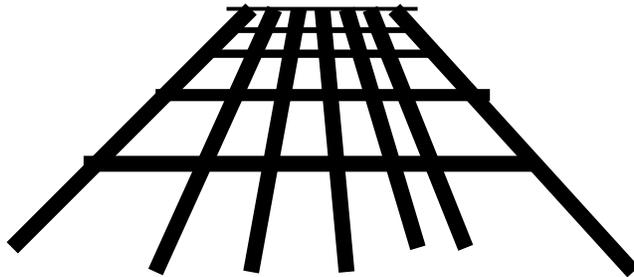
edge/outline that is the part of O where the front-facing and back-facing surfaces meet) – though one can indeed see O. One might think this is an unintuitive cost of my position: how can you see an object without seeing any of its surface?! Quick Answer: well you can see the very edge of the surface, where it forms a contrast with the light. (On the Casati & Varzi view, this is a dependent part of the surface just as it is a dependent part of the object as a whole.) Moreover, Sorensen himself is quite happy to allow that there are cases where one can see X without seeing *any* surface of X.

“A balloonist approaching a cloud first sees it by virtue of the light its outer portion reflects. As he enters the cloud, he sees the cloud from the inside – but not because any neatly delimited surface is available to him...

As the balloonist exits the cloud, he can see the sky, but not because he see the surface of the sky... You cannot make out which is the last part of the sky you can see. You are akin to a diver who can see the ocean all around him even when not looking at the surface or the bottom.” (p47-48)

So we can see clouds and oceans and the sky without seeing any surface. To this list I propose we add any silhouetted object – we see such an object by seeing only a boundary or edge of its surfaces, as defined by the surrounding light and one’s relative position.

In a later paper, replying to his critics, Sorensen (2011) gives one other consideration against the edge theory. His objection comes down to the fact that silhouettes “convey depth” – or as Sorensen also puts it: “silhouettes are 3-dimensional”. Sorensen’s example is looking at a back-lit wire-mesh grill that is extending away from one’s viewing position: (figure 8).



We see this silhouetted object receding into the distance.

As a quick clarification: Sorensen allows that a silhouette will only convey depth or 3-D shape given some contextual richness, such as the pattern of change over time or the position of the silhouette within a familiar 3-D environment etc “The three-dimensionality of silhouettes is easier to appreciate in a dynamic setting...most of the depth clues are relations between many objects.” (2011, 202-203) In this respect the depth perceived in a silhouette is no different to depth perceived in front-lit conditions – the more rich contextual clues there are, the more easily depth can be seen. In severely attenuated/simple/artificial cases, depth becomes hard to perceive or can be misperceived

(e.g. think of the Ames room illusion). Likewise, when a silhouetted object is not moving and is not situated in the midst of a familiar 3-d environment, then it does not convey depth.

I confess that I cannot see the force of this objection to the edge theory at all. As I have emphasised, the boundary, or boundaries, formed by the object where it meets the surrounding light, relative to the subject's viewpoint, are located out there in environmental 3-D space. This is the case either on the Stroll/Sorensen bulky matter conception, or on the Casati & Varzi bulk-less limit conception – though on this latter understanding, whilst an edge is located in 3-D space, it is not something that *takes up* any 3-D space. Either way, an edge is something located in environmental space and so can be located at different depths from the perceiver. In the case of the wire-mesh grill, one sees the intricate grid-shaped boundary (or boundaries) formed with the surrounding light. This environmental feature (or set of features) *really is receding* away from the observer. A boundary, just as much as a surface, can recede away from some location.

To be fair to Sorensen, he is replying to O'Callaghan (2011) who claims that a silhouette is “just an object's occlusion shape from a given perspective”. An occlusion shape is a shape that exactly blocks your view of the object. The idea then is perhaps that in seeing a silhouette one's experience does not represent any depth properties of the object, it represents only the relational property of the occlusion shape that is formed by the object in one's visual field. I think this totally mishandles the commonplace/intuitive idea that we see the silhouetted object's edge. The boundary of an object, where it meets a surrounding area of light-filled space, relative to a viewing location, is *not* a 2-dimensional feature of the subject's visual field. It is a perfectly mind-independent feature, located out in 3-d space¹⁶.

(6) THREE MORE PROBLEMS FOR THE BACK-SURFACE THEORY

In this final section, I will raise three further issues/difficulties for Sorensen's Back-surface theory.

(1) As already mentioned when we considered the possibility of a ‘huge tunnel’ running through the moon, to see an object silhouetted it is *not* necessary that its *entire* back surface be blocking light. Suppose that no light at all is falling on the middle of the back of O, but light is falling on all of the periphery of the back of O. We see O silhouetted, as we can see its outline. But *not all* of the back surface is blocking light. In general, to see a silhouetted object, the back surface only needs to block light so that an outline is created – the rest of the back surface can be causally dormant.

Sorensen's position depends upon the claim that: “the silhouette is the surface of the object that makes the object visible by virtue of the light it blocks”. So in the case where

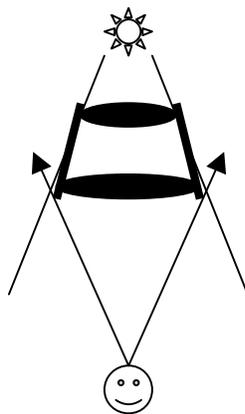
¹⁶ And to be fair in turn to O'Callaghan, *he* is mostly concerned with a quite different strand in Sorensen's book – hearing silences – and only mentions the silhouette case in passing.

only some of the back surface is actively blocking light, Sorensen must presumably hold that we only see part of the back surface, the part that *is* actively blocking light. Suppose then that in the above case, where only the periphery of the back of O is blocking light, we start to switch a light source on and off that illuminates the central portion of the back surface of O. This makes no difference to how O looks to S. But presumably Sorensen must say that as the light goes on and off, S alternates from seeing *part* of the back surface to seeing *all* of the back surface, despite there being no phenomenological change in S's experience. Or: suppose the periphery of the back surface is blocking light and so is the centre of the back surface, but there is a ring of the back surface in between these two regions that is not blocking light – Sorensen presumably has to say that S sees two disconnected regions of the back surface.

(2) Suppose that a perfectly black sphere is lit from both the front *and* the back, so it forms a contrast with its surroundings *both* by not reflecting the front-lighting *and* by blocking the back-lighting. Sorensen seems to be committed to the idea that we can simultaneously see *both* the front and back surfaces. For both front and back surfaces are making equal contributions to O contrasting with the surrounding light – so we should see the *entire* surface of the sphere here!

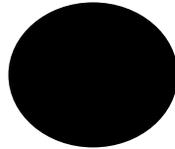
To be clear: I am not suggesting that seeing both the front and back surfaces of an object at the same time is *always* absurd. Think here of Sellars's (1982) pink ice-cubes, whose front and back surfaces we can see at the same time¹⁷. What I am suggesting is that it would be absurd *in this case* to hold that we can see both the front and the back of the opaque sphere at the same time.

(3) It is not only absorption *surfaces* that create silhouettes. An absorption *volume* can create a silhouette too. Consider a material that for every 1cm of thickness absorbs 1% of the incoming light. At 1 meter thickness, all the light is absorbed. Now suppose that a 1m deep truncated conical volume of this material is backlit/silhouetted (figure 9)



¹⁷ See sections §25-26 of his “Sensa or Sensings: Reflections on the Ontology of Perception” (1982). G. E. Moore (1918) had earlier provided the similar example of a soap-bubble, all of (the surface of) which we can see at a single instant.

Once more, the silhouetted volume will simply look as in figure 2:



Now there is no obvious candidate *surface* that is blocking the light. The entire volume is equally causally responsible for blocking the light. The 1cm layer nearest to the sun has hardly blocked any of the incoming light at all, so we presumably do not see this back surface. And it looks pretty arbitrary to claim that the absorption layer here is the lowest 1cm, farthest from the sun. This bottom-most part of the volume has also only blocked a tiny amount of light, and has done no more absorbing than any other 1cm layer. So do we see *all* of this 3-dimensional volume? That sounds pretty implausible to me. Surely we can see *none* of the internal volume.

It is much more plausible – indeed it is correct! – to say that what we see is the *outline* of this volume – the border between light and non-light that is formed by the volume’s circular edge relative to the viewer. This is where the visible boundary between incoming light and blocked light, which is created by the volume, is located relative to the given viewing position.

In view of these three unattractive consequences of the back-side theory, I think we should conclude that the part of a silhouetted object that we can see is its outline. And that is, both for the object and for the argument of the present paper, the end of the matter!¹⁸

¹⁸ Earlier versions of this paper were presented at the University of Geneva, at the National University of Colombia in Bogota and at King’s College London. I am very grateful to the audiences on all those occasions for their questions and comments. I am extremely grateful to both the editors of the present volume for their extensive, perceptive and wonderfully helpful written comments, which I only regret I have not been able to do better justice to in the final version of this paper. And lastly, I owe a huge debt of gratitude to Roy Sorensen – not only for writing such a brilliantly thought-provoking book, but also for generously and patiently discussing my paper with me over a very enjoyable conference dinner in Geneva.

References

- Casati, R. & Varzi, A. (1999) *Parts and Places: The Structures of Spatial Representation*, Cambridge: M.I.T. Press.
- Chisholm, R. (1957) *Perceiving: A Philosophical Study*, Ithaca: Cornell University Press.
- Clarke, T. (1965) 'Seeing Surfaces and Physical Objects', in *Philosophy in America*, (ed.) M. Black, Ithaca: Cornell University Press.
- Da Vinci, L. (2002) *The Notebooks of Leonardo Da Vinci*, (ed.) MacCurdy, E., Konecky & Konecky.
- Dretske, F. (1969) *Seeing and Knowing*, Chicago: University of Chicago Press.
- Elder, J. & Sachs, A. (2004) 'Psychophysical receptive fields of edge detection mechanisms' *Vision Research* 44 (8): 795-813.
- Grice, H. P. (1961) 'The Causal Theory of Perception' *Proceedings of the Aristotelian Society* (Supplementary Volume) 35:121 - 168.
- Jackson, F. (1977) *Perception*. Cambridge: Cambridge University Press
- Lewis, D. (1980) 'Veridical hallucination and prosthetic vision' *Australasian Journal of Philosophy* 58(3), 239-249.
- Hubel D. & Wiesel T. (1959) 'Receptive fields of single neurones in the cat's striate cortex' *Journal of Physiology* 148: 574–591
- Hubel D. & Wiesel T. (1962) 'Receptive fields, binocular interaction and functional architecture in the cat's visual cortex' *Journal of Physiology* 160: 106–154.
- Marr, D. & Hildreth, E. (1980) 'Theory of Edge Detection' *Proceedings of the Royal Society of London. Series B, Biological Sciences* 207 (1167): 187-217.
- Moore, G. E. (1918) 'The Presidential Address: Some Judgements of Perception', *Proceedings of the Aristotelian Society*, Vol. 19 (1918 - 1919): 1-29
- Nadernejad, E., Sharifzadeh, S. & Hassanpour, H. (2008) 'Edge Detection Techniques: Evaluations and Comparison' *Applied Mathematical Sciences*, 2 (31): 1507-1520.
- Noe, A. (2003) 'Causation and perception: The puzzle unravelled' *Analysis*, 63(2): 93-100.
- O'Callaghan, C. (2011) 'On Privations and Their Perception' *Acta Analytica* 26 (2): 175-186.
- Pears, D. (1976) 'The Causal Conditions of Perception', *Synthese* 33: 25-40.
- Peacocke, C. (1979) *Holistic Explanation*, Oxford: Clarendon Press.
- Price, H. H. (1932) *Perception*. London: Methuen & Co. Ltd.
- Sellars, W. (1982) "Sensa or Sensings: Reflections on the Ontology of Perception", *Philosophical Studies* 41: 83–111
- Shapley, R. & Tolhurst, D. (1973) 'Edge Detectors in Human Vision' *Journal of Physiology* 229(1): 165-83.
- Sorensen, R. (2008) *Seeing Dark Things: The Philosophy of Shadows*. Oxford: Oxford University Press
- Sorensen, R. (2011) 'Silhouettes: A Reply from the Dark Side', *Acta Analytica* 26 (2):199-211.

- Strawson, P. F. (1974) 'Causation in Perception', in *Freedom and Resentment and Other Essays*, London: Methuen.
- Stoll, A. (1988) *Surfaces*, Minneapolis: University of Minnesota Press
- Tye, M. (1982) 'A causal analysis of seeing' *Philosophy and Phenomenological Research*, 42: 311-325.