

Location, location, location

Austen Clark
Department of Philosophy
University of Connecticut
Storrs, CT 06269-2054

Elaboration of a talk given at ZenCon (Zenon Pylyshyn Conference)
University of Guelph, 1 May 2005. Draft of September 2006, under
submission for publication in the volume of conference proceedings.

Abstract

Thanks to Pylyshyn's work in the imagery debate, theorists are now apt to be at least a bit more cautious when they launch into descriptions of depictive representation. Even if the full force of his critique is not everywhere acknowledged, more care and vigilance is exercised when broaching talk of scanning mental images, focusing on such entities, or rotating them. The very places found in mental images--places that once rested so serenely under the gaze of the mind's eye--now tend to arouse feelings of disquietude, if not alarm. In this paper I will argue that there is an admirable theoretical continuity between Pylyshyn's critique of pictorial representation in mental imagery and his critique of "location based" models in visual perception and visual selective attention. If it is not appropriate these days for decent minds to look at inner pictures, is it any more appropriate for them to move the spotlight of attention across the master map? If our talk of places in mental images is bankrupt, then what are we to make of our talk of feature maps in the brain? Might not much of that real estate get foreclosed as well? Pylyshyn suggests the answer is yes, and proposes we reallocate now, into object-based alternatives. I shall argue the situation is not quite so dire. Following the analytical lead of P. J. O'Rourke, I shall propose a four way classification. There are indeed Bad Locations, but there are also Good Locations. Similarly, there are Good Objects, but there are also Bad Objects. I hope to clarify some of the distinctions between these.

1. Imagery, Round One

To understand Pylyshyn on perception, it is useful, and perhaps essential, first to understand his contributions on what might seem to be a distinct topic: mental imagery. The 1980's imagery debate was a portentous one for mental pictures, and Pylyshyn played a decisive role in it. Many of his recent (2001, 2003) arguments about the architecture of visual perception, and against "location-based" models, show a

striking and admirable continuity with those earlier arguments about the forms of representation implicated in mental imagery. As he puts it near the beginning of his recent book:

we must dispense with the "picture in the head" ... we must also revise our ideas concerning the nature of the mechanisms involved in vision and concerning the nature of the internal informational states corresponding to percepts or images. (Pylyshyn 2003, 3)

In the imagery debate we had bad inferences from experimental data to claims for a distinct, pictorial form of representation. Some of those same patterns of inference are found as well in the "objects *v.* locations" debate in visual perception.

What is the bad pattern of inference? The fundamental issue is: Do any available experimental results entitle us to believe that subjects in imagery tasks use a form of representation that is distinct in kind from the forms used in linguistic tasks? Do they provide *any reason at all* to think this? Pylyshyn says, forthrightly and firmly, "no". The question is whether results establish use of a distinct *form* of representation: of a "pictorial" or "depictive" form, as opposed to a "propositional" variety. To do this results must be traceable to a feature of the cognitive architecture, not simply to implicit knowledge, task demands, strategies, or some other labile cause.

What would it be to manifest a depictive form?

Let us try to be clear on what we take to be the central issue: does visual mental imagery rely (in part) on a distinct type of representation, namely, one that depicts rather than describes? By "depict" we mean that each portion of the the representation is a representation of a portion of the object such that the distances among portions of the representation correspond to the distances among the corresponding portions of the object (as seen from a specific point of view; see Kosslyn 1994...) (Kosslyn, Thompson & Ganis 2002, 198)

A depictive representation is a type of picture, which specifies the locations and values of configurations of points in a space. ... In a depictive representation, each part of an object is represented by a pattern of points, and the spatial relation among these patterns in the functional space correspond to the spatial relations among the parts themselves. Depictive representations convey meaning via their resemblance to an object, with parts of the representation corresponding to parts of the object... (Kosslyn 1994, 5)

Pylyshyn's position:

what I shall argue is not true is that the information in the visual store is pictorial in any sense; i.e., the stored information does not act as though it is a stable and reconstructed extension of the retina. (Pylyshyn 2003, 15)

In the opinion of this spectator, the first round of the imagery debate ended roughly as follows. Two widespread, deep, and stubborn sets of reasons for holding to the pictorial form were by Pylyshyn isolated, illuminated, targeted, terminated, dissected, sliced, stained, and mounted. What was left was taken out back and buried. Unfortunately, those scraps seem to reanimate; they don't stay buried for long. The two, seemingly immortal, irrepressible reasons for mental pictures were (and are), first, that introspection reveals the pictorial form directly. The experience of having a mental image is like the experience of seeing something spread out in front of you. How can you deny that you seem to be looking at a picture? A good lawyer could make any witness who denies such a thing seem (at the very least) disingenuous; more likely a scoundrel and a liar, deserving to be convicted. Second, the intentionalist fallacy. When we talk about "the image" it can become almost impossible to tell whether we are talking about the thing imagined or the thing that does the imagining. Mental pictures suffer from the same queasy ambiguity. But in straightforward contexts, at least, it is straightforward: places in the things one represents need not be represented by places in one's representings. If we carefully avoid these two mistakes, what is left of the argument for the claim that mental imagery must employ a distinct pictorial form? Not much. Pylyshyn also provided many arguments in detail about the inadequacies of "depictive" models. The most potent: that the content of the image depends on the subject's beliefs about the objects in the domain in question.

2. Imagery, Round Two

Round two of the imagery debate opened with the publication in 1994 of Stephen Kosslyn's *Image and Brain*, optimistically subtitled *The Resolution of the Imagery Debate*. (The analogy that springs to mind is a philosopher proposing a final resting place for *zombies*.) Accounts of depictive representation are amended, and the arguments acquire a neuroscience garnish. The key amendment is that the spatial properties and relations of the image are now construed as properties and relations in a "functional space". The basic idea: talk of spatial properties and relations ascribed to the image should not be taken literally. Instead all those attributions are a kind of "as if" talk, where what we're really talking about are the values returned by the procedures that read, write, and manipulate information in the image. Those procedures function in a way that is analogous to operations applied to a literal two dimensional display. If the image is an array in a computer, we have

procedures that access and manipulate distances between points. Those distances (the values returned by these procedures) would be true of a literal two dimensional surface. But this doesn't require that values of adjacent cells in the array be physically next to one another. Basically this is a move to Roger Shepard's idea of second order isomorphism: the image models spatial relations, but it need not itself employ spatial relations to do so.

Second, and more importantly for my purposes, neuroscience is claimed to provide evidence for some key features of depictions. First, that visual mental imagery uses some of the same brain mechanisms as does visual perception (in particular V1), and second, that neuroscience shows that those mechanisms use depictive representation. Kosslyn says:

Without question, topographically organized cortical areas support depictive representations that are used in visual perception. These areas are not simply physically topographically organized, they function to depict information. For example, scotomas--blind spots--arise following damage to topographically organized visual cortex; damage to nearby regions of cortex results in blind spots that are nearby in the visual field. Moreover, transcranial magnetic stimulation of nearby occipital cortical sites produces phosphenes or scotomas localized at nearby locations in the visual field. These facts testify that topographically organized areas do play a key role in vision, and that they functionally depict information. (Kosslyn, Thompson & Ganis 2002, 200)

the actual physical wiring is designed to "read" the depictive aspects of the representation in early visual cortex. In so doing, the interpretive function is not arbitrary; it is tailor made for the representation, which is depictive. (Kosslyn, Thompson & Ganis 2002, 199)

What defines round two as qualitatively distinct from round one is this appeal to neuroscience: the reference to topographically organized "feature maps", conjoined to the claim that some of the same mechanisms could support visual imagery.

Now the appeal to neuroscience adds yet another kind of image to the already confusing mix (fMRI images of the brain), and yet another kind of map ("feature maps"). If we can avoid being distracted by these pictures, however, the critical premise is easy to spot: that "topographically organized cortical areas support depictive representations". What are we to make of this premise? Pylyshyn gives a characteristically forthright response:

Even if we found real colored stereo pictures displayed on the visual cortex, the problems raised thus far in this and the previous chapter would remain and would continue to stand as evidence that these

cortical pictures were not serving the function attributed to them. (Pylyshyn 2003, 388)

The scraps have reanimated and reorganized; the debate is up and running, once again. And with that I can state the point of this paper. Theoretical objections to "depictive" representation, if they are cogent, would apply not just to imagery, but to everything, including visual perception. So, in particular, they would seem to rule out certain accounts of "location based" effects in selective attention. If places in a mental picture are problematic, what are we to make (for example) of the notion of a "spotlight of attention" moving across the "master map", traversing intermediary locations as it moves, in its own inscrutable fashion, from *A* to *B*? For this to make sense we need places which the spotlight traverses, or across which the "window of attention" moves. Such places have alarming similarities to those found in mental images. How, if at all, can we make sense of the locations posited in location-based models? Perhaps the very notion of a "feature map" is at risk. Does any and every account of feature maps endorse some sort of "inner picture" model? In what sense, if any, are "feature maps" *maps*?

My goal in what remains is to sort some theoretical commitments on these topics into two bins: good and bad. The task is necessary and unpleasant. Theorists must sort out which aspects of an analogical model apply to the real system, and which do not. Here our analogical model for a visual state is a picture or a road map. When we talk of feature maps as "maps", which of the properties of maps must be taken literally? Which are meant only as metaphors?

The task can be unpleasant, but I hope here to render it less so by following the analytical lead of P. J. O'Rourke in his masterpiece of economic analysis, *Eat the Rich*. O'Rourke (1998, 1) says: "I had one fundamental question about economics: Why do some places prosper and thrive while others just suck?" Why indeed? The question applies to visual places too. O'Rourke follows this question with four chapters, entitled Good Capitalism, Bad Capitalism, Good Socialism, Bad Socialism. Here I shall try to distinguish Good Objects from Bad Objects, and Good Locations from Bad Locations. Because Pylyshyn's critique focuses on the badness of Bad Locations, I shall start there.

3. Bad Locations

Economically speaking, Bad Locations correspond to Bad Socialism: Cuba. O'Rourke visited Havana in 1997 and said it "looked like 1960 Cleveland after a thirty-seven year strike by painters and cleaning ladies" (1998, 80). A compelling candidate for a Bad Place! Visually

speaking, Bad Locations are any of those found in models of visual perception that succumb to the same errors as models of pictorial or depictive representation. How can one succumb to the same errors? Let us count the ways.

3.1. *The ones in a mental image or in an inner picture*

These are bad if they are stipulated to be not just places where the representation is located, or places that it represents, but places *in* it that represent places that the organism perceives. So these are stipulated to be places in the image or picture that "map" onto places in the world. The mapping is semantically significant. They are allegedly homomorphic to, and thereby depictive of, places in the world.

...images are experienced as distributed in space...Because they are experienced as distributed in space, we find it natural to believe that there are "places" on the image--indeed it seems nearly inconceivable that an image should fail to have distinct places on it. This leads naturally to the belief that there must be a medium where "places" have a real existence. (Pylyshyn 2003, 371)

But, as he argued mightily in the imagery debates, round one, this conclusion is not mandatory. No available evidence requires us to postulate representations of this form. Pylyshyn puts his conclusion these days even more firmly. "We will have to jettison the phenomenal image", he says (Pylyshyn 2003, 47). What is tossed overboard is strictly the depictive form, not the phenomenology of imagery. That is, it is still true that to some people it seems as if they sometimes look at inner pictures. That's what they report. The claim is that this "phenomenon" (or appearance) of imagery is consistent with representations that are everywhere propositional.

3.2 *Places in your percept not within your current field of view*

A very similar point can be made about the phenomenology of visual perception. While it might seem to common sense, and to some introspectors, that seeing things is a matter of apprehending an inner picture, Pylyshyn rightly insists that such appearances can be explained in ways other than by postulating internal pictorial representation.

We cannot escape the impression that what we have in our heads is a detailed, stable, extended, and veridical display that corresponds to the scene before us. ... We find not only that we must dispense with the "picture in the head," but that we must also revise our ideas concerning the nature of the mechanisms involved in vision and concerning the nature of the internal informational states corresponding to percepts or

images. (Pylyshyn 2003, 3)

One way to diagnose whether you suffer from an objectionable form of the "inner picture" model is to ask: Does that inner display extend, spatially or temporally, beyond the limits of what can, in a given moment, be seen? If the answer is "yes", your theoretical commitments clearly include some Bad Places. If the answer is "no", you might or might not be infected. As will be seen, further tests are necessary.

Pylyshyn does not deny the existence of retinotopically organized feature maps, as found in V1-V4. But each of these is confined to registering information derived from the array of retinal receptors. They neither can nor need to register information about regions that currently cannot activate any receptors: all those regions in the ambient optic array whose light fails to intersect any part of the retinal array. Nevertheless, it might seem as if visual perception involves a comprehensive or panoramic inner picture, one that includes many of those momentarily unseen portions of the scene.

It has been suggested that what we "see" extends beyond the boundaries of both time and space provided by sensors in the fovea. So we assume that there is a place where the spatially extended information resides and where visual information is held for a period of time... (Pylyshyn 2003, 28)

This last assumption is one that Pylyshyn is most eager to deny. While there might be retinotopic maps, there is, says Pylyshyn, no panoramic inner picture: no extension of the retinotopic maps so as to include, in the same map, portions of the distal scene that are currently unseen. So places in a retinotopic map are (tentatively) OK (more on this below); places represented by retinotopic maps are OK; but there the map talk stops. There is no further (much less final) comprehensive map, into which all the retinotopic versions--all the gleanings from each glimpse--can be arrayed. Gaze control and saccadic integration are not managed by larger and more comprehensive versions of the retinotopic maps found in V1 to V4.

3.3 The ones identified using a particular "reference frame" or using particular "coordinates"

Talk of "reference frames" is often just a way of specifying a category of bodily motion invariance: which motions (of stimulus or of body parts relative to one another) will, and which will not, alter the proposed state (whether it be neural or representational). To say that a sensory state "employs a eye-centered reference frame" means that the state won't change as long as spatial relations between the stimulus and

the eyeball are unchanged. To say that it employs a "head centered reference frame" means that changes in that state are correlated instead with changes in the spatial relations between the stimulus and the head. Since the eyes can move in the head, these are distinct; a stimulus can have a fixed location in an eye-centered reference frame even while it moves in terms of the head, and vice versa. Such terminology is a useful and unobjectionable shorthand.

But talk of reference frames can have a more fulsome interpretation, where we assume there is an origin and some fixed points (axes) relative to which locations and other spatial properties and relations are determined. Often theorists can slide into this talk without even noticing that it says rather more than mere motion invariance. For example, Cohen & Anderson (2004) say "A reference frame can be defined as a set of axes that describes the location of an object" (104). Note that this description does not require the animal to use those axes! They then proceed to say

Sensory targets are often coded in different reference frames. For example, the location of a visual stimulus is initially coded based on the pattern of light that falls on the retinas, and is thus in retinal coordinates. ...The location of a tactile stimulus is coded by the pattern of activation in the array of receptors that lie under the skin's surface and, consequently, it is coded in a body-centered reference frame. (Cohen & Anderson 2004, 104)

These inferences (the "thus, in retinal coordinates"; "consequently...in a body centered reference frame") simply do not follow, unless we read "coordinates" and "reference frame" very loosely. It might seem churlish to criticize what is here probably an innocent use of an analogy, and indeed, there is nothing to criticize as long as the theorist recognizes that this is *merely* an analogy. The danger with analogies, though, is that unintended portions of them creep unbidden and unnoticed into one's theories.

Similarly, talk of "coordinates" can just be a way of describing the data (as in spaces derived from multidimensional scaling); but it is dangerous if one presumes the animal actually employs them to identify anything. If we really mean "coordinates", then this presumes that we have an origin point, axes, and metrical level measurements of distance along those axes (the real number plane, or perhaps polar coordinates). It also implies that mechanisms of spatial discrimination use those coordinates, as coordinates, to pick out the locations of things. This I think no one seriously believes, despite the occasionally fanciful diagrams.

3.4. *The place lit up by the spotlight of attention*

If we assume that this is not a place in the world, but is rather one located on the master map of locations, then it *may* go onto the list of Bad Places. It depends on how one understands the "map" talk. If we presume that the master map is literally a map, or that differences in places in the map are used to represent differences in places in the world, then such places are heir to all the theoretical difficulties associated with places in a mental image, and are, indeed, Bad. If one endorses some semantically significant relation between places in the map and locations in the world, then it is prey to all the difficulties just noted for the fulsome sense of "coordinates".

One particularly clear diagnostic indicator: if one assumes that when attention shifts from stimulus *A* (in the world) to stimulus *B*, then the spotlight of attention must traverse locations on some "master map" intermediary between those used to represent the place of *A* and those used to represent the place of *B*, then one has endorsed some Bad Locations. Those "intermediary" locations are the Bad ones. The assumption that there are, and perhaps must be, such intermediary locations in the map indicates conclusively that one thinks of the spatial relations in the map as semantically significant. That satisfies the definition of "depictive". These implications are not evaded by the expedient of turning all the talk into talk of "functional" space.

3.5. *Empty ones*

Sometimes Pylyshyn charges location-based models with the crime of representing empty space: places as such; unoccupied or unfilled places; places with nothing in them. These sound Bad indeed.

The theoretical question for us reduces to whether it is possible for visual indexes to point to locations *as such* (i.e., to unfilled places) and that question is not yet settled experimentally, although there is some evidence that the position of an object can persist after the object has disappeared (Krekelberg 2001), and that at least unitary focal attention may move through the empty space between objects, though perhaps not continuously and not at a voluntarily controlled speed... (Pylyshyn 2003, 252)

The contrast is stark: the choice is between models that direct attention at empty places, and those that direct it at familiar, fulsome, objects:

there is reason to believe that at least some forms of attention can only be directed at certain kinds of visible objects and not to *unoccupied places* in a visual scene, and that it may also be directed at several distinct objects. (Pylyshyn 2003, 160)

the evidence... suggests that the focus of attention is in general on certain primitive *objects* in the visual field rather than on unfilled *places*... (Pylyshyn 2003, 181)

Is this a fair contrast? It is true that a location-based model worthy of the name should allow that differences in the direction of attention need not always be framed in terms of (or be resolvable into) differences in the objects to which attention is directed. Instead attention can be directed as finely as spatial discriminability allows. But do such models require or imply that attention can be directed to unfilled places?

Well, they might; but only if an animal sometimes encountered such locales. "Empty" can mean various things: (a) it contains nothing at all; (b) it contains nothing that would provide physical stimuli; (c) it contains nothing sufficient to stimulate any transducer of the organism in question; (d) it contains no perceptible physical objects. (a) is a literal vacuum. Case (b) is also extra-terrestrial: it might include fields and particles that do not interact with any transducers. Not a vacuum, but filled with a soup of quarks, say. Strictly, (c) is more or less impossible to produce, unless it is the same as (b): even a silent pitch black room contains stimuli for thermal sensation, as well as vestibular ones. In practice one must think of both (c) and (d) as confined to one modality. So a pitch black room would give a *visual* example of (c). In contrast, (d) could include the *ganzfeld*, or for that matter a very foggy evening; the regions contain visual stimuli but no discriminable objects.

It would tax any animal to discriminate among places that are literally devoid of stimuli (as in (a) or (b)). An animal would have that capacity only if its forebears had routinely been challenged by the need to discriminate one empty location from another. The analogous burden to place on the other side would be to require the animal to be able to discriminate objects as such: objects that lack any properties at all. These are what philosophers call "bare" particulars: manifesting the pure objecthood of objects, isolated from all their distracting properties. I don't think it is fair to require object-based models to be able to tell two of these apart. Similarly, on this interpretation of "empty", a location-based model need not even try to satisfy the request to tell apart two empty places.

But if by "empty" one means simply that the animal has spatial discriminative capacity even if it is not confronted by any discriminably distinct objects, then I think the answer is yes, it does. The wafts of cloud in a white-out or a *ganzfeld* serve as examples. Different patches of cloud or portions of *ganzfeld* remain spatially discriminable from one another.

A better contrast might be between places that are filled with

distinct *objects* and places that are not. An object-based model implies that where there fail to be distinct objects there cannot be differences in how selective attention is directed. A location-based model allows such differences as long as the organism still has the capacity to make spatial discriminations in that region. It asserts that when we write the operating principles for the directing of selective attention, the variables employed do not need always to refer to objects; they can range over any features that can be spatially discriminated from one another.

4. Good objects

Visually speaking, good objects are all and only the ones fit to serve as values of variables in the true model of what the visual system represents. Economically, the analog for Good Objects is Good Capitalism: Wall Street. O'Rourke says of this place: "The traders spend their day in that eerie, perfect state the rest of us achieve only sometimes when we're playing sports, having sex, gambling, or driving fast. Think of traders as doing all these things at once, minus perhaps the sex. ...All free markets are mysterious in their behaviour, but the New York Stock Exchange contains a mystery I never expected--transcendent bliss" (1998, 21)

The preceding problems with Bad Locations are used by Pylyshyn to argue for the thesis that visual indices are bound, not to locations, but to objects.

In what follows, discussion will be confined to ... the view that focal attention is typically directed at *objects* rather than at *places*, and therefore that the earliest stages of vision are concerned with individuating objects and that when visual properties are encoded, they are encoded as *properties of individual objects*. (Pylyshyn 2003, 181)

Medium sized package goods are good objects. Many visual proto objects turn out to be identical to medium sized package goods. So many visual proto objects are perfectly OK.

5. Good Locations

Now the problem is just this: are all locations posited in location-based models Bad Ones? Are any of them are good? Good Locations in O'Rourke's typology correspond to Good Socialism: Sweden. "Sweden was the only country I'd ever been to with no visible crazy people. Where were the the mutterers, the twitchers, the loony importunate? Every Swede seemed reasonable, constrained, and self-possessed. I stared at the quaint narrow houses, the clean and boring shops, the well-

behaved white people. They appeared to be Disney creations..." (O'Rourke 1998, 56)

My question is whether there are any Good Locations in the intentional domain. How can we construe the talk of locations in location-based models, or the talk of maps in feature maps, so as to avoid the very real dangers of which Pylyshyn has warned us? Specifically, is any theorist who wants to pitch a tent somewhere in the location-based domain (or on a feature map) necessarily camping in a Bad Location?

To start, it helps to note that Pylyshyn does endorse some good locations--some unproblematic spatial domains. They include:

- 5.1 Locations of objects and of their parts.
- 5.2 The location of the brain.
- 5.3 Location of mental representations within the brain.
- 5.4 Locations in topographically organized areas in V1-V4.
- 5.5 Locations as represented in retinotopic maps.
- 5.6 Locations of "feature clusters".

But what then of feature maps? Must these contain, or be maps of, Bad Locations? V1 is one of many alleged "feature maps" in the cortex. What's going on in those? And is Kosslyn right to say that "without question" they support depictive representations?

The core notion of a "feature map" in neuroscience is, I think, a region of cortex organized topographically. But everything hangs on how one understands the term "topographical". The simplest interpretation is anatomical. The fibers coursing into the cortical area come from some source region or regions, also within the nervous system. In a "topographical" organization, there are local regions in the source within which neighboring cells project, more or less, to neighboring cells in the destination. There might be several such local regions, between which there can be abrupt discontinuities in the projections. A prominent example is found in the retina: the left side of each retina projects to the left side of the brain, and the right to the right. So we find a topological "tear" right down the middle of the retina. But within each region, neighborhood relations are (pretty much) retained.

In sensory areas, cells in a feature map can often be associated with receptive fields: regions in circumambient space within which stimuli of a specified kind can affect the activation level of the given cell. This yields a second way to understand the topographic organization. Cells

that are neighbors in the cortical region in question often have receptive fields that are neighbors in circumambient space. When they do, one can see a very strong reason to call the thing a "map": it is a topographically organized array within the organism that seems to represent places outside and around the organism. But as will be seen shortly, the notion that cells in feature maps preserve neighborhood relations among points in space is never strictly speaking true, and it is often very misleading.

It should be obvious that mere topographic organization is not by itself sufficient to show that the cortical region in question employs pictorial or depictive representation. That way of organizing the fiber bundles can be better ascribed to physiological economy (fewer crossovers and shorter bundles) or neural development (easy ways to grow the things) than to features of our cognitive architecture. Furthermore, the cortical region may be representing something other than location altogether. For example, an auditory feature map can be topographically organized, respecting neighborhood relations on the basilar membrane, but this makes it a *tonotopic* map, of different frequencies, not different places. Mustached bats have auditory maps across which we get systematic variation in Doppler shift (see Suga 1990). It is not mapping space, but rather relative velocities.

What then is needed for these regions of cortex to be, also, maps of space? This conclusion is not automatic! A second obvious necessary condition can be put as follows: the region must enable some spatial discriminations. It carries information about spatial properties and relations of its targets in such a way as to allow the organism to navigate. Without this it wouldn't contribute to what I think of as "feature-placing".

Is that enough? Are these regions of cortex "without question" depictive? If we consider V1, for example, the best possible case for calling it a "feature map" give us three premises. First, that we have an orderly projection of fiber bundles from its source (mostly LGN) to V1. So, second, neighbors in V1 typically have receptive fields that are neighbors. (And it functions in accord with this principle, as Kosslyn points out. Damage to V1 causes scotomata whose perimetry can help the neuropsychologist identify where the damage took place.) Third, thanks to V1, the creature can make certain spatial discriminations that it otherwise cannot make. If you doubt this, just consider what it loses in those scotomata.

These three premises, so far, do not imply that the map is a "map of space", i.e., that points and distances within V1 map homomorphically onto points and distances within the ambient optic array. For it to be

literally a map of space, it would have to sustain those spatial discriminations in just one way, via a homomorphism with spatial properties. As Kosslyn puts it, it must be such that "distances among portions of the representation correspond to the distances among the corresponding portions of the object" (Kosslyn, Thompson & Ganis 2002, 198). The pattern of inference here seems eerily familiar. In fact, thanks to Pylyshyn, we can recognize it. It is exactly the pattern used to sustain the idea that mental imagery must involve inner pictures.

That V1 is required for certain sorts of spatial discriminative capacities shows that information in V1 is used by the organism to improve its steering. It does not show that the information in V1 is organized just like a map or a picture. The structure might enable spatial discriminations (of some particular sort) without itself modeling space. If you look at its finer structure, I think it's pretty clear it does *not* model space. In fact, perhaps no feature maps are maps of space in the "depictive" sense. V1 is certainly a big array of measurements, but values in adjacent cells are not invariably measurements of adjacent places.

Details of the structure of V1 make this clear. The details in question are not subtle or contentious; most of them have been known since the work of Hubel & Wiesel. In particular, the ocular dominance pattern, and the arrangement of "orientation slabs", royally messes up the neighborhood relations. In a given orientation "slab" within (layer III of) a cortical column, all the cells will fire maximally to a edge, bar, or slit of a given orientation. Cells in the neighboring slab do not register the same orientation in neighboring receptive fields, but instead a different orientation (in different receptive fields). And we have a block of orientation slabs for the left eye immediately adjacent to a block for the right eye. These are the left eye view and the right eye view of the same location in external space.

The critical point: if you move half a millimeter in one direction, you might not change the receptive field at all, but instead move to a region receiving input from that same receptive field, but from the other eye. Move in another direction and the receptive field will shift, but so will orientation. Move in a third direction and only the optimal orientation shifts. These distances do not map uniformly onto distances in the ambient array. Ergo, homomorphism fails. V1 is not depictive.

6. Hypothesis: How a feature map represents

How then does a feature map represent? One minimal but plausible description of the content of a feature map is: it indicates the spatial

incidence of features. It might do more than this, but it does at least this. That is, it registers information about discriminable features, in such a way as to sustain varieties of spatial discrimination that can serve to guide the organism. The latter two conditions focus on downstream consumers of the information, not what causes it. Registration of information in a feature map endows the creature with some spatial discriminative capacity. If that map is used, the steering can improve. To carry on its other business, the animal relies on the constellation of features being as therein represented.

One way to get at the spatial content of a feature map, guaranteed to work for every feature map, is to ask: what sorts of spatial discrimination does this particular feature map enable? That is, which spatial discriminations are possible using this map that were not or would not be possible without it? For some cortical regions dubbed "feature maps" by neuroscientists, the answer could well be "none"--in which case the map is not a representation of the spatial incidence of features at all. (Such a map will not employ the representation form I identify below as "feature placing".) The idea: if feature map *M* is representing the spatial incidence of features, then it is being used as a representation of the spatial incidence. The information in it about spatial properties and relations is exploited. One way to show that it is exploited is to show that certain kinds of spatial discriminations could not be made without it; without map *M* working normally, the guidance system and steering--the navigational and spatial competence of the organism--suffers some decrements.

The focus on downstream consumers is a way of showing that the registration of information is used as a representation; that it has a content that is used. To tie representations to the world, show that they improve the capacity to get around. But feature maps can do this without necessarily being pictorial or depictive; they can satisfy the condition without being, literally, maps or inner pictures.

Psychological theory right now lacks any deductive proofs, or even compelling arguments, that establish how information *must* be organized to endow creatures with some new spatial discriminative capacity. It's too early to invoke a priori principles in this domain. (It follows that there's never a good time to be a priori--but that's another question.) So, in particular, there is no compelling reason to think that information *must* be organized depictively in a feature map if that feature map enables a creature to make spatial discriminations which it otherwise could not. Here again we should thank Pylyshyn: his work on mental imagery showed how, in principle, a set of propositions could do the job.

What then does V1 represent? To answer this question, analyze what use downstream consumers make of the information registered in it. A first stab: these cells in layer III of V1 represent "(edginess of orientation theta) (thereabouts)". Edginess is the feature; "thereabouts" indicates its incidence. Those cells in layer III of V1 have the job of registering differences in orientations, in such a way as to allow spatial discrimination of them. If they do that job, the animal can rely upon those indicators, and thereby steer a bit more successfully than if it lacked them.

More generally, I have proposed that we call this form of representation "feature-placing". It "indicates the incidence of features" in the space surrounding the organism. The name is partly in honor of Sir Peter Strawson's work (1954, 1974) on "feature-placing languages", which contain just a few demonstratives ("here" and "there") and non-sortal universals (feature terms, like "muddy" or "slippery".) A paradigm feature-placing sentence is "Here it is muddy, there it is slippery." Such sentences indicate regions and attribute features to them. Strawson argued that these languages could proceed without the individuation of objects. The same seems true of the representations employed in feature maps. It seems a bit much to claim that V1 "refers" to places, "identifies" regions, or "demonstrates" locales. All the latter locutions arguably invoke some portion of the apparatus of individuation. Feature-placing is prior to, and can provide the basis for, the introduction of that rather heavy machinery.

Another way to put it is that feature maps in V1-V4 transact their business in a location-based way. A particular feature map can endow a creature with new spatial discriminative capacities without also endowing it with an ontology of objects. It can get the spatial discriminative job done without investing in that sort of machinery. A skimpy basis can suffice; the business can be run on an ontological shoe-string. It is also important to insist that the regions visually discriminated are not inner, or mental, ones. They are not inside the organism or inside the mind. If the job is to guide spatial discriminations then representing those places will not help. Visual "thereabouts" are always, resolutely, in the ambient array, not in the retina. The cortical feature map might be retinocentric (it uses an "eye centered" reference frame) but it is *not* retinotopic. It is not *about* the states of the retina, but instead about features in the world.

If V1 were representing places on the retina, then it should represent the blind spot as empty. But patterns are completed "across" the blind spot, as shown by Gatass and Ramachandran's experiments on scotoma and "filling in" (see Churchland & Ramachandran 1994). The filling in

across the optic disk can give a veridical "perception" of the distal place, even though it would be a non-veridical representation of what is going on at the retina. V1 cells in the "Gatass condition" fire just as they would if there were a stimulus stimulating the non-existent receptors in the optic disk. If we were representing places on the retina, this would be a non-veridical representation (Churchland & Ramachandran 1994, 82)

So I think there is good reason to say that what these parts (of layer III) in V1 are representing is something akin to "(edginess of orientation theta)(thereabouts)." "Thereabouts" indicates a region of circumambient space--a region of visual perimetry, in the ambient optic array. "Edginess of orientation theta" indicates a feature discriminable in some portion of that space. The orientation is of an edge in external space, not across the eyeball. It is feature-placing, and both the features and the places are distal.

7. Bad Objects.

That concludes my plea for the possibility that not all Locations in the intentional domain are Bad. Symmetry demands that we also consider the possibility that not all Objects are Good. This is our last quadrant: Bad Objects. In O'Rourke's typology, it corresponds to Bad Capitalism: Albania. Albania, he says, "has the distinction of being the only country ever destroyed by a chain letter--a nation devastated by a Ponzi racket" (O'Rourke 1998, 36). Chain letters and Ponzi rackets in completely unregulated markets can be tough on widows and orphans. Likewise, visually speaking, Bad Objects are kinds of objects to which a pure object-based model is at least somewhat vulnerable.

7.1 *Merely virtual ones*

By "merely virtual" I mean an object that seems to exist, or appears to exist, but does not. The ogres, wizards and dragons displayed on computer screens in some computer games are paradigm examples. The experience of looking at such a screen can be very much like seeing a dragon, but there is no dragon there to be seen.

It is a bad idea ever to allow merely virtual objects to serve as the referents of visual indices. Such an index is supposed to be entirely non-descriptive, gaining all its representational capacities from direct access to the referent itself. So if in fact there is no referent, there is nothing to which the index can be attached. An index attached to such a thing is attached to no thing.

Now in many of the experiments in multiple object tracking,

subjects are not in fact tracking objects, in any ordinary sense of the word. Instead they are looking at a computer display and tracking figures on the screen. What exactly is the object to which a visual index is attached? Pylyshyn says "the observer may be indexing clusters on the screen or, more likely, a virtual distal object, where only the part of the chain from the scene to the observer is real" (217). I think the latter alternative invites indoors some Bad Objects. Suppose one can index a merely virtual object. Then in one episode of a computer game an index might be attached to a dragon, and in another, to an ogre. But indices are supposed to be non-descriptive, and neither dragons nor ogres exist. So what is the difference between indexing a dragon and indexing an ogre?

An index gets its content entirely from what it points at. It does not encode any properties, contains no description, etc. So if it is pointing at nothing, it should have no content. So if it is pointing at an object that is a merely virtual object, there should be nothing that differentiates one such pointer from another. So there can't be difference between pointing at nonexistent object *A* vs. nonexistent object *B*.

For this reason it seems preferable to keep the door shut, and adopt the other alternative: what is indexed must be something literally seen, on the screen. Similarly, for the same reason, it is hard to see how an index could ever get attached to a non-visible object. Pylyshyn wonders "What exactly the index points to when the object is temporarily out of view" (2003, 268 n 20). Nothing comes to mind!

The problem in both cases is that reference failure is catastrophic for an index. In such a case there is no thing to which it points, and reference does not succeed by description. So in what sense is it "referring" or pointing at all? This should be a case of an indexical without a referent. How could it have any content at all? If we style these pointers on those found in programs, this one should give an "out of bounds" memory error, cause the blue screen of death to appear before the mind's eye, and make the mind itself lock up. Abort, retry, fail?

7.2 *Non-individual ones*

If vision is to be object-based through and through, from the get-go, then the values of variables in all of its representations, everywhere, are always, and only, objects. Even at the earliest stages, the representanda are objects. The worry here is simply that some of those earliest stages do not have the wherewithal to represent their objects *as* objects. In particular, they lack the wherewithal to represent that which makes one of them one, and not two.

To use the technical terminology: these "objects" lack criteria of individuation. And if they lack individuation, it will seem feckless, at least to some philosophers, to call them "objects" at all. If "this" and "that" are bound to objects, then one can distinguish the possibility of encountering first this one and then that one from the possibility of encountering this same one twice. Otherwise the application of the apparatus of individuation--count nouns, identity, sortals, indefinite pronouns, articles, etc.--is not required.

Consider the early stages of visual representation, in V1 to V4. You, the neuroscientist, laboriously describe how one of them works. Someone in the audience rises to ask, "but does this particular state, at this stage, represent exactly one x , or does it represent both one x and a y such that y is not identical to x ?" Even though the question is probably from a philosopher, and I am a philosopher, I would sympathize with your plight. Such a question seems somehow maladroit, ill-informed, out of place. In these stages there is nothing available yet that would be, or could be, sufficient to answer the question of what makes one thing one, or distinct things distinct. These stages operate in a regime that is prior to, and free from, such worries.

If this sympathy is not entirely misplaced--if the notion of such regimes is at all plausible--then these stages are representing the "things" they represent without representing them as falling under criteria of individuation. If we insist that even these stages are representing objects, these will be "non-individual" objects. They lack individuation. Nothing is such as to make one of them one, and not two. Common sense would cavil at calling such things (such values of variables) "objects". If we cannot count them, what justifies the distinction between singular and plural? Quine (1974) and Geach (1980) and Wiggins (2001) have argued, at length, that the acquisition of the apparatus of individuation is no mean feat. Unless we think that V1 (for example) can acquire such a thing, the variables therein range over features or regions, but not objects.

7.3 The ones numbered more than six

This is the most variegated kind of Bad Object, because it is not a kind at all. Like vulnerability to Ponzi schemes, the problem here seems to be a structural limitation of visual indices. Indices are limited to five or six. What happens when we run out?

In particular: Can we account for the spatial discriminative capacities that become possible when a creature acquires a feature map by supposing instead that all the reference of its representations proceeds through five or six visual indices? To be object-based all the

way down, you must think that all such information can be registered in a system of object-files (or, more broadly, a system in which all the variables are bound to objects). Consider, for example, V1. In order to explain how this map (V1) endows the creature with (say) the ability to discriminate horizontal lines from slightly off-horizontal lines, we have to think of feature detection and registration across a vast swath of space, sensitive throughout to minute differences in orientation. It has somehow to register that there is an edge or bar or stripe extending from x to y ; and then register orientation of that edge from point to point.

How would a FINST system represent a pattern of (say) nine parallel lines, tilted slightly? We have more lines than we have FINSTs, yet even registration of the features of one line seems to require lots and lots of terms and relations (edginess, connectedness, continuity, straightness, orientation, is parallel with, etc.).

Location-based theorists surmise that at least some of the information must be registered in data structures that contain variables that range over something other than objects. The books can be organized differently; the business might be transacted in an ontologically skimpy, location-based way.

8. Conclusion

To sum up. Some clearly Bad Locations are: the ones in a mental image or in the inner picture. Places in your percept that are not within your current field of view. And, finally, the ones identified using coordinate systems or reference frames.

In contrast, the presumption is that almost any Object is Good, particularly if it is one that can be bought or sold in a capitalist economy. Things you can track, and, when the funds become available, purchase. Medium sized package goods are, therefore, the paradigm Good Objects.

There are also some Bad Objects, however. Merely virtual ones qualify: the ones that do not exist, even though they have an index attached to them. Sadly, these too are sometimes bought and sold in capitalist economies. Other Bad ones include objects that lack individuation. If you buy one of these you don't know what you bought. Finally, those numbered more than six. These are bad because they can't be indexed.

Close examination of Pylyshyn's theory shows that it allows for the existence of at least some Good Locations. These include: locations of objects and of their parts. The location of the brain. Locations of mental representations within the brain. Locations in topographically organized

areas in V1-V4. Locations as represented in topographic maps. Locations of "feature clusters".

In terms of this typology, are "feature maps" Good, or Bad? I have argued that they can be Good, though to stay that way they must eschew any claim to be depictive.

The upshot? Let us leave the last word to P. J. O'Rourke: "Money turns out to be strange, insubstantial, and practically impossible to define ... economic theory was really about value. But value is something that's personal and relative, and changes all the time. Money can't be valued. And value can't be priced...I should never have worried that I didn't know what I was talking about. Economics is an entire scientific discipline of not knowing what you're talking about." (O'Rourke 1998, 122-23)

References

- Churchland, Patricia S. & Ramachandran, Vilayanur S. (1994). Filling in: Why Dennett is wrong. in Antti Revonsuo & Matti Kamppinen, (eds) *Consciousness in Philosophy and Cognitive Neuroscience*. Hillsdale New Jersey: Lawrence Erlbaum Publishers, 65-91.
- Cohen, Yale E. and Andersen, Richard A. (2004). Multimodal spatial representations in the primate parietal lobe. In Charles Spence and Jon Driver (eds.) *Crossmodal Space and Crossmodal Attention*. Oxford: Oxford University Press, 99-121.
- Geach, P. T. (1980). *Reference and Generality*. 3rd. ed. Ithaca: Cornell University Press.
- Graziana, Michael S. A., Gross, Charles G., Taylor, Charlotte S. R., and Moore, Tirin (2004). A system of multimodal areas in the primate brain. In Charles Spence and Jon Driver (eds.) *Crossmodal Space and Crossmodal Attention*. Oxford: Oxford University Press, 51-67. (ch 3 of that book)
- Konishi, Masakazu (1992). The neural algorithm for sound localization in the owl. *The Harvey Lectures*, Series 86, 47-64.
- Kosslyn, S. M. (1994). *Image and Brain: The resolution of the imagery debate*. Cambridge, MA: MIT Press.
- Kosslyn, Stephen M., Thompson, William L., and Ganiz, Giorgio (2002). Mental imagery doesn't work like that. (Reply to Pylyshyn 2002). *Behavioral and Brain Sciences* 25(2), 198-200.
- O'Rourke, P. J. (1998). *Eat the Rich*. New York: Atlantic Monthly Press.
- Pylyshyn, Zenon (2001). Visual indexes, preconceptual objects, and situated vision. *Cognition* 80: 127-158.
- Pylyshyn, Zenon (2002). Mental imagery? In search of a theory. *Behavioral and Brain Sciences* 25(2), 157-237.
- Pylyshyn, Zenon (2003). *Seeing and Visualizing: It's Not What you Think*. Cambridge, MA: MIT Press.
- Quine, W. V. O. (1974). *The Roots of Reference*. La Salle, Ill.: Open Court.
- Stein, Barry E., Stanford, Terrence R., Wallace, Mark T., Vauhan, J. William, and Jiang, Wan (2004). Crossmodal spatial interactions in subcortical and cortical circuits. In Charles Spence and Jon Driver (eds.) *Crossmodal Space and Crossmodal Attention*. Oxford: Oxford University Press, 25-50. (ch 2 of that book)
- Strawson, P. F. (1954). Particular and general. *Proceedings of the Aristotelian Society* 54: 233-260.
- Strawson, P. F. (1974). *Subject and Predicate in Logic and Grammar*. London: Methuen & Co. Ltd.
- Suga N. (1990). Cortical computation maps for auditory imaging. *Neural Networks* 3:3-21.
- Suga N, Olsen JF, and Butman JA (1990). Specialized subsystems for processing biologically important complex sounds: cross correlation analysis for ranging in the bat's brain. In *The Brain: Cold Spring Harbor Symposia on Quantitative Biology* 55:585-97.
- Wiggins David (2001). *Sameness and Substance Renewed*. Cambridge: Cambridge University Press.