

Perception Preattentive and Phenomenal

Austen Clark
Department of Philosophy
103 Manchester Hall U-2054
University of Connecticut
Storrs, CT 06269-2054

(Penultimate draft, 6 April 2006)

In Paul Thagard (ed.), *Handbook of Philosophy of Psychology and Cognitive Science*, Dordrecht: Elsevier, 2006, pp 105-139. (This is volume 12 of the multi-volume *Handbook of Philosophy of Science* (eds. Dov Gabbay, Paul Thagard, and John Woods), Elsevier, 2006.

Abstract

Recent work in experimental psychology and neuroscience has revealed a rather surprising architecture for early (or preattentive) perceptual processes. This paper will describe some of the surprising features of that architecture, and how they bear on recent philosophical debates about the notion of phenomenal consciousness. I will argue that the common sense idea that states of phenomenal consciousness are states of a unitary kind cannot survive confrontation with the details of how our early perceptual processing works. In particular, that architecture forces us to contemplate the prospect of phenomenal consciousness being sundered in two, with states that have phenomenal character making an appearance far before the arrival of anything one could call consciousness or awareness.

Contents

- I. The architecture of early vision
- II. Competition, selection, attention
- III. Features and phenomenal properties
- IV. Preattentive phenomenal properties
- V. Phenomenal consciousness, but hold the consciousness
- VI. Philosophical implications of divorce
- VII. Perceptual psychology after divorce
- VIII. Conclusions

References

Perception Preattentive and Phenomenal

The conundrums of phenomenal character and consciousness have often motivated philosophers to study perception; and a particular area of study that will prove worthy of their attention is research into what are called "early" or "preattentive" perceptual processes. These are, roughly, processes that start at the transducers and end where selective attention has access to the results, and can select some favored few for further processing. These early processes are the ones most likely to be called "sensory"; they are at any rate simpler, and they make their appearance earlier than, the more sophisticated states that underlie perceptual judgments. If non-conceptual representation is employed anywhere in the system, it would be employed here. Animals that cannot muster the words for a perceptual judgement can nevertheless sense things. These simple sensory states are dear to the heart of those interested in phenomenal character, and I hope they find their interests piqued by preattentive phenomena.

To fit constraints of time and space this paper must set aside consideration of relations between perception and knowledge, and between perception and action, even though there is enormously interesting work underway on both fronts. We will mine single-mindedly the vein that leads into phenomenology. To narrow the topic even further, I will confine the discussion to early vision. It is still an enormous field, and it provides more than enough material with which to examine some recent discussions of phenomenal consciousness.

The reader should be forewarned: the architecture of early vision is surprising, bizarre, *weird*. I will describe some of the surprising and weird features of that architecture, and then consider how folk concepts of appearance and awareness might be applied to it. The results, like the architecture itself, are somewhat bizarre. In particular, the preattentive architecture challenges the idea--it puts enormous stress on the common sense notion--that "phenomenal character" is and must be coeval with awareness. Instead, the two split apart, with "phenomenal character" showing up first, in places as yet unoccupied by awareness. After reviewing some of the evidence, I shall argue that the most reasonable conclusion is that some states of preattentive sensing are states in which one is being appeared-to, even though one is entirely unaware of what appears, of any aspect of its appearance, or of being in the state of being appeared-to. So "phenomenal character" and consciousness fly apart; their association in "phenomenal consciousness" is a merely contingent conjunction of two distinct things. This puts some stress on our ordinary notions; one burden of the argument is to consider the least costly methods for stress relief.

I. The architecture of early vision

Models of early visual processing produced in the last forty years have grown increasingly sophisticated and complex, and in detail they differ markedly from

one another. But certain global features of the landscape are found reflected in all of them (see Neisser 1967). They are:

1. That distinct visual attributes of stimuli are initially processed largely independently of one another, in parallel, for the entire visual field;
2. That these processing streams, "channels", or "feature maps" extract information of varying orders of complexity, have complicated linkages with one another, may reside in distinct portions of the visual nervous system, and in cases of brain damage may be selectively dissociated from one another;
3. That at some point there is a competitive selection process, whereby some results of some of these parallel processes are selected for additional, capacity-limited, processing;
4. That a result of winning that selection is that discriminations in the selected domain can be made more quickly and more accurately; that the additional processing allows for the resolution of details and structure that might otherwise be unresolved, and that there may be features that are resolvable only as a result of such selective attention; and
5. That sensory representations which are not selected do not gain all these benefits, though they may gain some.

The initial stage is high-bandwidth analysis (e.g., of many megabits per second) over the entire visual field, extracting "local" or low level visual features, with different "channels" carrying on the analysis of different visual attributes. For example, motion, hue, luminance, and shape are likely candidates for being distinct attributes, analyzed in distinct channels. Each such channel carries on a high bandwidth analysis of its favorite feature for the entire visual field.

The surprise here is that the separation of channels is done by attribute (or by distinct dimensions of variation of attributes), and not something that would seem, a priori, or to a computer engineer, to be more reasonable. So for example one might think that if there are parallel channels, it would make sense to separate them by the regions they scrutinize, by their ambit. One covers the upper right quadrant, one the lower, etc; and each reports on everything visible occurring within their sector. Or perhaps the jobs of the different modules could be divvied up more intelligently and more dynamically: the visual field might be "segmented" into likely objects, and a channel is given the job of analysing everything visible about that object. The end result of such a module would be a visual representation of that object complete with color, shape, motion, texture, size, and so on. In fact a visual representation of the object would look suspiciously like a picture of the object. We add up all these visual representations, and we get a complete picture of everything one can see.

But it is evident that early vision is not organized in that way at all. Finding a *picture* at the terminus should arouse suspicions; right around the corner from it you will fall under the disturbing, lidless gaze of the mind's eye. Visual perception does not produce anything like a fully detailed picture of everything one can see, and visual representation is nothing like a picture. Or at least if it

is, it is a very odd picture (or better, set of pictures) indeed: we would need a picture of nothing but the object's color, leaving out all details of its shape; another picture of just shape, but no color; one of motion, but not of contour; and so on.

There are many different strands of evidence that support each of the five assertions, and many different points of disagreement between different models on each of the five. So for example the evidence for the first point--that distinct attributes are processed independently, and in parallel--comes from various sources. The first is neurophysiological: the careful tracing of visual pathways and different cortical areas connected by those pathways. The organization is quite intricate, with numerous independent input channels, roughly three dozen distinct cortical areas devoted to visual processing, and staggering interconnections between them. Different cortical areas seem specialized for the processing of distinct attributes, with distinct areas tentatively identified for motion, color, and form.

A second strand of evidence is neuropsychological: particular types of brain damage can "knock out" a particular dimension of variations in visual appearance, while leaving others intact. Particularly compelling is evidence of double dissociation: some types of brain damage knock out A, while leaving B; while others knock out B, and leave A. Achromatopsia and akinesia show this dissociation: one type of central nervous system damage can leave a patient colour blind, but still able to perceive motion. Another impairs discrimination of motion, while leaving color perception intact.

A third type of evidence is purely psychological, and is based on various types of experiments, including, most prominently, reaction times in visual search tasks (Treisman 1998, 1996, 1993, 1988; Wolfe 1994, 1996a, 1996b) . In a visual search task a "target" is to be located as quickly as possible within an array of varying numbers of "distractors". Particular combinations of attributes of targets and distractors can make an enormous difference to the speed of such searches, and these differential times give evidence that some classes of features are processed independently of others. For example it is found that in some classes of features, a target with a unique value within that class can be located in a fixed (and brief) interval, no matter how many distractors are present. If the target is the red letter, and we have a single red H among green H's, blue T's, and green X's, the red H will "pop out", and reaction times for its identification will be low and basically constant, no matter how many distractors are present. Similarly if we have a single letter "X" among H's and T's, of whatever colors, the X will pop out. But if the target can be identified only by a conjunction of features--a single red H among green H's, red and blue T's and red and green X's--then search becomes much more arduous, and reaction times increase linearly with the number of distractors. It is as if each item in the array must be examined in turn, and such examination takes some finite time. The latter is sometimes called "serial search", or simply "less efficient" search.

The relatively straightforward inference from this result is that feature

families which allow "pop out" are processed independently of one another, and such processing is high-bandwidth, covering the entire visual field. From the results above, for example, color and orientation (which allow the diagonals of an X to pop out from the horizontals and verticals of H's and T's) would seem to be features processed independently of one another. Each allows pop-out. Only if a target is defined by a conjunction does search become less efficient.

Pop out indicates a special kind of "preattentive" processing of the feature in question. Attention is effortlessly drawn to the singleton, no matter how many distractors are present. The processing to do this must happen before attention arrives on the scene, and it must encompass the entire visual field, since pop out occurs no matter how many distractors are present. So features in that group seem to be favored with independent, high-bandwidth, full field processing, completed preattentively.

There are other tests for the independent processing of distinct families of visible attributes (Treisman 1988; Wolfe 1996b). One is "effortless texture segmentation", or the ability to discriminate regions defined by contours across which the feature in question changes. There is some boundary across which feature F changes; and the segments of that boundary in turn form a continuous closed curve. In this case the entire region might "pop out", or be immediately noticeable. The feature difference defining the texture can be minuscule: minute differences in orientation, closure (whether we have full squares or just three sides), and so on. Whereas other differences one might think would suffice do not: N's cannot be preattentively segmented from regions of backward N's, E's and F's are texture-wise indistinguishable, and so on (see Julesz 1984). The inference is that certain feature values are the subject of preattentive parallel processing, while others are not. Orientation and closure yes; the difference between E and F, no.

With these and other sorts of experimental probes available, a considerable amount has been learned about the different kinds and categories of features processed preattentively. Jeremy Wolfe (1996b) has provided a useful list of visual features that pass both of the two tests just mentioned. He calls them "basic features":

- color
- orientation
- curvature
- vernier offset
- size/spatial frequency/scale
- motion
- pictorial depth cues
- stereoscopic depth
- gloss
- various form primitives

So for example a curved line (with basic feature number three) can be picked out quite efficiently if it is surrounded by straight lines, however many they may be (that's "pop out"); and it allows preattentive texture segmentation: a region filled with curved lines is effortlessly seen to have a different texture than one

composed of straight ones. All of the features on the list pass both tests. Some are perhaps unfamiliar. We get "vernier offset" by breaking a straight line, and giving the second segment a slight offset, or sideways jog, relative to the first. (Human vision is extraordinarily sensitive to vernier offsets, detecting them at visual angles smaller than the angle between adjacent receptors in the retina.) "Pictorial depth cues" include cues for occlusion, perspective, shape from shading, and others that artists might use to make a two dimensional canvas represent a three dimensional scene. The last ("form primitives") is the most controversial and complicated. For reasons that will become important later (see Wolfe & Bennett 1997) it is not exactly "shape", but includes local features of contours and surfaces that might together yield "cues" or evidence for shape. Wolfe mentions, for example: line termination, intersection, and closure (e.g. is the contour a closed curve?). Others are less well attested: topological features such as holes, convergence, and perhaps containment.

Different entries in the list of basic features are, intuitively, of differing orders of complexity. It takes more visual processing work to discover whether a contour is a closed curve or not than to discover whether some segment of it is oriented vertically or horizontally. Some of the features might be attributes of two dimensional polygons, while others (gloss, depth, occlusion) require more than two dimensions. Even "size" is ambiguous between a two dimensional interpretation (visual angle) and a three dimensional one closer to common sense. Preattentive features must then have some ordered information linkages to one another. In order to isolate a distinct texture region, for example, one must find boundary segments across which there is a discontinuity in the values of some feature; such segments must be linked to one another continuously; and the defined contour must be closed, without significant gaps. The relations that the data points must bear to one another become increasingly complex as one ascends in these constructions.

An example might clarify the problem. As emphasized by Marr (1982), the analysis of boundaries and edges in the optic array is a major task of early vision. At the receptor level the raw data consists of variations in the membrane potentials of approximately two hundred million receptors--200 mega-pixels, if you like. Variations in these membrane potentials are caused by varying numbers of isomerizations of photopigment molecules, which are in turn related to the intensity and wavelength of the electromagnetic radiation absorbed by the cell. An initial task is simply finding the local edges--the regions where there are large discontinuities of intensity between neighboring cells. Solving this problem requires "local differencing" operations--a fancy way of saying we need to perform hundreds of millions of subtractions (between intensity values of neighboring receptors) and have usable results every few milliseconds. (And we are downsizing, so you have to do the job yourself, and complete it while paying attention to something else.) It's a bit of a challenge! But it seems to be accomplished within the retina itself. Then these local differences need to be linked up into coherent "edges"--connected tracks, or series of points, across which we find discontinuities of intensity.

Only such edges have "orientations" (which needless to say are useful to record) and "contour" features such as curvature, kinks, or corners.

But even if we have isolated a continuous track across which there is a robust discontinuity of intensity, it is unclear what such an edge means in distal terms. It might be a "luminance" edge or a "reflectance" edge. That is, the differences in intensity might be caused by differences in the amount or character of light falling on a uniform surface, or they might be caused by differences in the way different parts of a surface reflect light of a uniform character. Perhaps the edge is the edge of a shadow; perhaps it is a place where surface characteristics change, or perhaps it is place where one object occludes another. Sorting out what kind of edge each edge might be is the next major job of early vision, and this analysis bumps us up into higher levels of complexity. A two dimensional world would lack shadows, and no object in it could occlude another. (We might have polygons that intersect, but a true flatlander has no concept that one lies on top the other, or that one is behind the other.) So early on the analysis of edges will force us to assign depths to the neighboring points; the edge acquires an orientation that is not simply two dimensional.

We confront, then, a second source of complexity in the relations among preattentive features: what might be called their geometry, or the spatial structure of the coordinate system in which their relations are defined. This too varies across different visual features. It is clear for example that some "boundary" or "edge" features can be specified determinately in a two-dimensional retino-topic coordinate scheme. So the discontinuities of intensity which define the border might be specified to obtain between two adjacent retinal locations (and we use some spherical coordinate system for the two dimensions of the retinal surface). The retinotopic coordinates of the point on which one's eyes are focused are always (0,0). But of course the eyes move in the head, the head moves on the trunk, and the trunk moves in the world; each degree of freedom spawns another potential coordinate scheme, from head centered to body centered, to, eventually, a full "allocentric" system, in which no part of the body has any privileged position. "Being to the left of" has different interpretations in these different schemes: the left half of the retina might or might not be receiving stimuli from points to the right of one's nose, which might or might not be to the right of one's sternum. And in an allocentric scheme, the relation no longer has a fixed interpretation at all, with as many variants as there are observers on the scene.

These logical and geometrical orderings of features are usefully combined in what has been called a "feature hierarchy" (see Treisman 1988), starting at the simplest and proceeding upwards into those that require more intricate relations and more intricate coordinate schemes:

Point level features: all the data needed for ascription of the feature in question can be had given just one visible point.

Merely local differences, across some span, in some feature dimension: "discontinuities of intensity". Requires relations between at least two

points: but "edge" is still too sophisticated a term for this.

Segment or edge level features: oriented and connected discontinuities of intensity, connected in a series, and so having some orientation. Here for the first time a connected series of local differences is treated *as* a series or group, and has an orientation different from the span across which there are feature differences.

2d surface properties, with no depth: properties of a 2d surface, specifically, so these are properties ascribed to points on a plane, in varying ways; no shadows or depth yet; no points on the surface vary in distance from the observer.

2.5d surface properties, so called because each point on the surface is also assigned a unique depth coordinate. (It is not fully three dimensional, because for each 2d point there is only one depth coordinate, and not a full third axis. Such surfaces can be tilted relative to the observer, have bumps, convexities, and concavities, but no three dimensional volume.)

Full 3d surfaces and volumes.

Different "basic features" make their first appearances at different points within this hierarchy. It is perhaps problematic whether there are any point level features, but color is often mentioned as a possible example. (The problem is that color perception seems always to require color contrast, and so requires relations between a point and its surround.) But color also shows up as a property of two dimensional surfaces (the famous color patch of sense-data theorists), a property of oriented surfaces (walls of the same color meeting at a corner), a property of three dimensional surfaces of objects, and a property of three dimensional volumes (as in colored mists or colored light). We must arrive at "edge level" features before any interpretation can be provided for orientation, curvature, vernier offset, and various of the form primitives; though features that might show up as mere features of edges (such as curvature, orientation, and closure) are instantiated at higher points of the feature hierarchy as well. Size and motion require at least a two dimensional coordinate scheme. The depth suggested by pictorial depth cues requires more than two dimensions, though some may suggest only orientation or curvature of a two dimensional surface, so fit in a 2.5 dimensional space. But shadows and occlusion, stereoscopic depth, and gloss have their expected interpretation only if they are placed in a coordinate scheme that is three dimensional. The difference between a reflectance and an illumination edge can be introduced only at this point; prior to it one thing cannot occlude another. Note that "motion" can in fact implicate developments across *four* dimensions: a boxer described as "weaving" and "bobbing" has a spatio-temporal trajectory, with a prominent temporal axis (as well as the spatial ones) along which developments have a characteristic form. Anyone acquainted with the idea that rhythms are temporal "shapes" will understand the point (see Goodman 1977). Features such as "glistening" or "shimmering" also require some temporal extension.

In short, features are not all of piece; some have a higher order of logical and geometrical complexity than others. The point may seem purely pedantic,

but grasping it is essential in the clinic, since otherwise the existence and symptomatology of the various *visual agnosias* are entirely inexplicable. Several of these syndromes seem precisely to be loss of the capacity to detect higher-order features, even though lower-order ones are still intact. An agnosic patient may show no loss of visual acuity, as shown by perimetry tests; yet lose the ability to visually recognize objects by their shape. In associative visual object agnosia, a patient may have unimpaired perception of local segment features--as shown by the ability to draw objects or match pairs of pictures faithfully--yet fail to recognize the object by sight. (See Farah 1990, 57-60). Instead the patient often volunteers a list of lower order form primitives, and then hazards a guess as what the object might be:

The patient could not identify common objects present visually and did not know what was on his plate until he tasted it. He identified objects immediately on touching them. When shown a stethoscope, he described it as 'a long cord with a round thing at the end' and asked if it could be a watch. He identified a can opener as 'could be a key'. Asked to name a cigarette lighter, he said 'I don't know' but named it after the examiner lit it. He said he was 'not sure' when shown a toothbrush. Asked to identify a comb, he said 'I don't know'. When shown a large matchbook, he said 'It could be container for keys'. (Farah 1990, 58)

What seems to be lost is the ability to recognize the overall shape of an object, with this capacity replaced by arduous feature-by-feature matching and hypothesis testing. Such patients can fail to recognize common objects (tea bags, rings, pens) which they can quite accurately draw.

An even more spectacular and specific agnosia is prosopagnosia, or the inability visually to recognize faces. The deficit can be quite specific to visual identification of faces. Such a patient may have no sensory loss and can recognize other common objects visually, and furthermore can recognize people by other distinguishing characteristics (voice, clothing, gait, etc); but if the face is the only clue, even a spouse can be misidentified. Facial features involve shapes, orientations, sizes, and symmetries, and are clearly higher order. Humans seem to have a "fusiform face area" whose activity is highly correlated with face recognition, which is not too surprising given the biological importance of recognizing con-specifics. In terms of our feature hierarchy this is specialized feature analysis at a rather high level.

II. Competition, Selection, Attention

"Basic features" are preattentive: their detection is completed prior to the activation of selective attention. They are noticeable or salient "attention-grabbers"; they draw attention to themselves, and so must not themselves require attention in order to be detected. Indeed, such detection is often the basis on which the selection of selective attention is made: attention might be drawn to the letter "h" because it is the only red letter in the display. Furthermore, different sets of basic features must be detected simultaneously and in parallel, since at any given moment any one of them could grab attention in the same effortless, almost instantaneous fashion. A single red letter, or a

moving one, or one that displays a unique orientation, or a single glossy one, ... and so on, could grab attention in (roughly) fixed time no matter how many distractors. So the inference seems robust: one must simultaneously and in parallel detect basic features of color, motion, orientation, glossiness, and so on.

For various reasons offered over the years, it is clear that mental life cannot forever proceed in such parallel, independent tracks. Sadly, at some point it becomes subject to central control, and acquires at the least the appearance of unity. That is, at some point there is a transition from high-bandwidth, data-driven, independent channels to something more centralized, of lesser bandwidth, and with a greater capacity to pre-empt the underlings. The contrast between processing "before" and "after" this transition point has been made in various ways. So as not to beg any questions, I shall simply dub the two sorts of processing "preattentive" (or "early") and "post-attentive" ("central"). There have been many disagreements over the years on how to characterize the differences between these two sorts of processing: parallel v. serial; bottom-up v. top-down; data-driven v. task-driven; exogenous v. endogenous; modular v. general-purpose. Likewise there have been many controversies over where the transition point is found: whether it is "early" or "late"; and why there is a transition point at all. Some have argued that the transition is needed because of the enormous amounts of data, and the lower ("serial") capacity of central processing; others that it may simply reflect the need for response selection: to choose *the* way to move now, since one cannot manage to send oneself to two distinct destinations. Psychologists continue to differ over the benefits of central processing, and over the fate of sensory representations that do not receive it. But, however else it is conceived, there is a consensus that the point of transition is the point where selective attention selects some representations, and not others, for further processing. This point controls, or exerts some control over, the flow of information from one style of processing to another. Various homely analogies are used: it is a gateway, a filter, a movable window, through which some can pass, but most are barred.

Post-attentive visual perception differs in various ways from the preattentive variety. It is "capacity-limited" in the sense that it cannot simultaneously be applied to every portion of the ambient visual array. A subset or subregion must be selected, and even then the analysis of that portion takes some time. So if central processing must be used to search every visible item, search times increase linearly with the number of items. This is often called "serial" search, and its contrast with the fast, "parallel" search in pop-out is one of the more robust findings of experimental psychology. The serial appearance of post attentive processing may be deceptive (parallel processes with limited capacities could yield the same pattern of results), and later models have moved away from the notion that there is a dichotomy between "serial" and "parallel" processing, but the basic empirical contrast (now called "efficient" vs "inefficient" search) is unarguable.

If central processing is capacity-limited in this way, then it would be

swamped if it had to deal with every visible feature simultaneously. Some selection must be made; the system must somehow establish priorities among the important or salient visibilia, and select which ones are first to receive the benefits of fuller analysis. Since resources for that fuller analysis are limited, they must be allocated. For the psychologist, this is the key function of selective attention: to select from a pool of potential candidates some favored few that are elected to receive further processing. This selective process is widely agreed to be competitive; the candidates are racing against one another. The winners of the competition are the ones that are selected. Some models even mention primary elections, followed by a general election (see Desimone and Duncan 1995). The rules of the race--the terms of the competition--can vary, depending on the current needs and desires of the needy organism. The smell of food is more compelling if one is hungry, and (surprise!) so is its taste. Bird-watchers and flower-gardeners will notice very different things if they walk together through the woods. In these and other competitions the competitors are representations, produced by the various feature channels, of stimuli one has sensed. Winners are the ones selected, by selective attention, for further, central processing. There may not always be exactly one winner--perhaps the hungry organism wants to gobble down several things simultaneously--but the losers in the competition do not gain entry to, are not selected for, central processing.

Common sense ascribes various benefits to the practice of paying attention to what you are doing, and experimental psychology has confirmed many of them. For example, if you are paying attention to a particular visual region, you are more likely to notice novel stimuli there, you will respond more quickly, and your discriminations of features in that region will be more accurate. (See Posner 1994, Posner & Rothbart 1992). "Paying attention to the region" is thought to be equivalent to devoting more processing power to representations of features and events in that region; and that additional processing allows quicker detection, more accurate discrimination, and better resolution of details and structure. Basic features can perhaps be resolved without this additional processing, as can gross morphology and some jumble of form primitives, but discerning the precise shape of the thing--the exact structure of relations among those form elements--seems to require the sort of central processing that can only be brought to bear if the target is selected by selective attention (see Wolfe & Bennett 1997).

This idea can explain why paying attention to something visible seems to place that thing in a spotlight--a rather mysterious spotlight, it turns out, since it employs no physical light. Increasing the illumination on an object can indeed help one see it better, but so can increasing the processing power devoted to analysis of visual detail. As anyone who has used the world wide web knows, the more computer power one can throw at a picture, the quicker its details are resolved. The spotlight analogy is probably as old as the human capacity to direct artificial illumination on objects of interest, but perhaps internet usage will eventually provide it with a rival. Attending to something helps one see it

better, not by directing occult illumination at the object, but rather by allocating additional processing to the analysis of visual features.

Competition for selection imposes an organization on the transition between preattentive and post-attentive perception. In any race there are winners, runners-up, and all those who also ran, surrounded by an even larger class of spectators, hangers-on, couch potatoes, and others who did not compete at all. In the spotlight metaphor we have an object or two in the focus; others in the penumbra, to which the spotlight could be slewed with minimal effort; and a larger class of things entirely out of its range. Similarly, the stream of consciousness is reputed to have a few phenomena in the foreground at any moment, and a somewhat larger but still finite collection of phenomena in the background. The latter are all the ones to which one is not directly attending, but of which one is "aware". Finally, surrounding that is a limitless but insensible netherworld of phenomena to which one cannot shift attention. As many authors have noted, it is clear that one is "aware" of things beyond those to which one is strictly paying attention; indeed the former often provide the next targets when attention shifts. A simple hypothesis is that this terminology reflects the fact that selective attention is allocated competitively. One focuses attention on the winners (they are in the focus, the foreground, the spotlight) but the runners-up are the ones of which one is "aware". The latter could have been selected for further processing, and would have been, but for the fact that something else came in first. In the next moment the currents might shift, and something of which one is currently merely "aware" might attract one's attention. If the currents are shifting rapidly in this fashion, the waters are choppy, and we call it "being distracted".

The more important point for our purposes is that if there is a selective process, the selection itself must be based in part on current information. Efficient allocation of scant central resources requires that it be at least somewhat sensitive to current developments. An animal whose attention was directed a priori would fail to manage any contingency; its steerage would be competent only among the necessary truths. So some sensory information is used in order to select where attention will be directed next. This information is found behind the stage-works, so to speak: it is not in the spotlight, or anywhere out there on the stage, but is rather being used, by the evil little homunculus who directs the spotlight, to determine where the spotlight will go a moment hence. (A homunculus is by definition a little man, and this one is evil because he controls what *you* pay attention to, and you never gave him permission.) It is not information of which you are aware, but is rather information already received by that sub-personal agency who is right now determining what you will be paying attention to a moment from now.

Furthermore, if there is some regular process of selection, it is carried out by some sort of mechanism, which like any mechanism has normal operating parameters, limits, and vulnerabilities. Posner (1994; Posner & Petersen 1990; Posner & Rothbart 1992) has done the most to reveal some of the details of how selection itself works. There are three separable processes in shifting

visual attention: disengaging it from the current focus, moving it, and then engaging it on a new focus. Each of these three subprocesses has its own normal operating parameters and areas of vulnerability. For example, it takes a certain amount of time to disengage and to shift attention, and subjects have tremendous difficulty perceiving events that happen in that interval. Change blindness and the attentional "blink" both exploit this parameter of attentional selection. Change blindness is not literally "blindness": there is no sensory loss, and the change is perfectly perceptible once it is noticed. What is blind to the change are the mechanisms of attention: the change occurs in such a way that it is difficult or impossible for it to attract attention. And if it *cannot* attract one's attention, one is "unaware" of it.

The attentional blink is a similar phenomenon which was discovered years earlier (see Raymond, Shapiro, & Arnell 1992; Raymond, Shapiro & Arnell 1995; Luck, Vogel, & Shapiro 1996). Subjects are presented a very rapid sequence of visual stimuli, and asked to respond to particular targets. For example, they press a key each time a vowel appears in a stream of letters. If two targets are presented within a certain interval of one another, the odds of missing the second target increase. But the odds do not increase in the way one might expect: if the second target is immediately after the first one, it is more likely to be identified correctly than if there is a slight lag, of about a hundred milliseconds, before it is presented. The worst performance is found when several letters intervene between first and second targets; with lesser intervals, or greater ones, identification improves. The suggestion is that the hundred millisecond lag indicates an operating parameter of the mechanisms that disengage, shift, and re-engage attention: by that time attention is engaged with the task of responding to the first target, and it is difficult to disengage it to register the second target. Just as changes cannot be noticed during an eye blink, the target is not noticed during an "attentional blink".

Both change blindness and the attentional blink show that it is possible to exploit vulnerabilities in the mechanisms that shift attention, and find stimuli that (given everything else going on) slip between the cracks of the selective process. These are stimuli to which one *cannot* shift attention. It is interesting that change blindness is called "blindness"; ordinary language makes scant distinction between events one does not sense and events one cannot notice, often treating them as equivalent. They are both events to which one is, in different senses of the word, "insensible". But change blindness and the attentional blink are not sensory failures at all; they are demonstrations of the vulnerabilities of an attentional mechanism, specifically the one which does the job of selection. If that mechanism has a hitch, then one cannot shift attention to a stimulus; one is insensible to it, and it is not overly misleading to say one is "blind" to it. It is a distinct kind of blindness from the sensory kind, and from the agnosia kind ("mind blindness"), but it is a third way in which one can fail to have normal visual experience.

And this points us to the most deeply surprising feature of the visual architecture described so far. Visual inputs are divvied up into discrete

dimensions of independently detectable "features", and processed independently and in parallel in topographically organized "feature maps". Each such feature map covers the entire visual field, and so in a certain way it is "panoramic", but only in terms of its own, favorite feature. One map is a specialist for motion, anywhere; another for color, anywhere; and so on. After all these feature extractions one might expect to find some comprehensive representation of all of them, put together; a synoptic view. But there seems to be no such thing. Beyond the feature maps we find nothing but what seem to be scattered (but carefully directed) glimpses; episodic and focused representations, serving the needs at hand. Their objects are scattered in space and time, in a surprisingly thin distribution. Change blindness, the attentional blink, and other results suggest that after the feature maps perhaps all we have are these temporary representations of conjunctions, relations, and structured properties, marking the visual cruxes for current affairs of interest, but vanishing as soon as the interests change.

The system can afford to analyse conjunctions, relations, and more structured properties only if such analysis is necessary for some on-going action. Even then it has to pick and choose. So conjunctions and more structured features are extracted only for isolated and scattered portions of space-time. Post-attentive processing is not engaged in the construction of a panoramic view, a fully detailed picture of all that is visible, but is better thought of as expedient probing and sampling, with targets carefully selected. For example, as you sit down you might employ well directed glimpses, and rather evanescent representations, to find the edge of the table, the front edge of the chair, the center of the seat, the handle of the coffee mug. Each of these targets requires rather sophisticated analysis (edges of objects, relational properties, etc), but results can be discarded as soon they have been used.

It is only in early vision--in preattentive, sensory processes--that we find anything resembling panoramic views. Even there the resemblance is not great. Perceiving, or post-attentive processing, is a narrowly focused business of probing and sampling, constructing representations that are punctate and transient, serving the needs at hand. It yields at best a pile of glimpses. In this architecture, a comprehensive picture is nowhere to be found.

III. Features and phenomenal properties

At this point in the story the sentient organism is focusing attention on some salient visible feature, and is aware of others, which were also in the running in the competition for selection. Such awareness of a visible feature counts as an example of what is ordinarily called "visual experience". The animal is aware of something it sees. One would think therefore that the state qualifies as an example of "phenomenal consciousness". If you want to check its qualifications, here is the initial baptism for the latter, technical term:

P-consciousness is experience. P-conscious properties are experiential properties. P-conscious states are experiential states, that is, a state is P-conscious if it has experiential properties. The totality of the experiential properties of a state are "what it is like" to

have it. Moving from synonyms to examples, we have P-conscious states when we see, hear, smell, taste, and have pains. P-conscious properties include the experiential properties of sensations, feelings, and perceptions, but I would also include thoughts, wants, and emotions." (Block 1997, 380)

A state in which one is aware of something one sees is, presumably, a state of visual experience. Any state in which one sees something without being aware of what one sees would not be a state in which one could be said to experience that thing, and Block quite consistently therefore excludes "subliminal perception" from the ranks of phenomenal consciousness, on the grounds that it is "perception without awareness" (see Block 1997, 393). And if the architecture so far described has the wherewithal to yield states of visual experience, then it is thereby also yielding states of "phenomenal consciousness".

Which states in the architecture are the states of phenomenal consciousness? The question gives one pause. The issue becomes quite complex, because the notion of a "phenomenal" property is quite complex. In fact the word is multiply ambiguous, and those ambiguities contribute to the problem. It has simpler and less simple interpretations, often conflated with one another.

The simplest and perhaps oldest notion of "phenomenal property" derives from ancient thought about the problem of illusion. Sometimes things are not as they appear. The arrow looks bent when it is half in water; the mountains look blue from a distance. Although the study of illusion is not a large part of perceptual psychology, the terminology needed to describe illusions looms large in the philosophical traditions. And indeed there are well known and replicatable stimulus arrangements under which people are subject to visual illusions of various sorts. So one line in the Mueller-Lyer illusion looks longer than the other (even though are both of the same length); after adapting to a bright red light, white paper will look green. In these situations we have some visual appearance which is illusory: it fails to correspond with how things really are. The ancestral home for phenomenal properties is found in these contexts. A phenomenal property is a property of appearance--of how things look, feel, sound, or in general, seem or appear. Such a property may or may not be veridical.

A slightly more recent analysis (Chisholm 1957) first picks out a class of "verbs of appearance", with which various predicates for sensible properties can be attributed, even though the referent of the sentence fails to have that property. So "the paper looks red" can be true even though the paper is, in fact, white. "Looks" is then a verb of appearance, and "red" (or any other predicate for a sensible property), used within the context of any verb of appearance, attributes a phenomenal property. "Looks red" characterizes how the paper appears, and it can look that way even though it is in fact white.

The acid test for this older notion of phenomenal property is that such a property can characterize appearances even if those appearances diverge from reality, as they do, prototypically, in sensory illusions. It remains true that one

line in the Mueller Lyer illusion continues to look longer than the other, even though measurement shows them to be the same length, and one knows them to be the same length. But the notion gets generalized to include properties of appearance even when they are veridical--even in those situations in which things *are* as they appear. So "looks red" comes to characterize the similarity in appearance across a large class of episodes: all those episodes in which something looks red but is not, and those in which something looks red and is red (see Sellars 1963). All those episodes have a shared feature--a shared phenomenal property--characterizing how something appeared. In some of them that appearance is veridical, in others not.

This is the oldest and simplest notion of "phenomenal property", but there are many other more recent and more intricate variants. Sense-data theories complicated the issue by identifying phenomenal properties with properties of odd objects called sense data; and the "qualia" debates around functionalism further confused matters by abandoning the odd objects but retaining talk of their inscrutable properties. Many philosophers today think of qualia as properties of mental states, and some identify those properties with phenomenal properties.

One question is whether phenomenal properties are characteristics of sensible appearance, or characteristics of something else. The other alternatives these days are, in order of increasing specificity: they are properties of mental states; they are properties of conscious mental states; they are characteristics of being conscious of one's own mental states. The simplest notion, as already described, is that "looks red" is a paradigm phenomenal property. A mental state itself has no known visible appearance; no mental state literally "looks red", as far we know. There is presumably some property of a mental state (specifically a visual sensory state) in virtue of which something one sees looks red, but that property is not itself the property of looking red. Similarly, if one is conscious of something that looks red, there may be some property of the state of consciousness related to the fact that the thing one sees look red, but the property of the state of consciousness is certainly not one of looking red. If it too is a phenomenal property, it is a different sort than the one that can be instantiated by things one sees.

The most intricate notion of "phenomenal" properties ties them directly to "what it is like" to have certain mental states. Block for example identifies "experiential" properties with the totality of "what it is like" to have a state. Notice that the bearer of the properties has shifted yet again. In order for there to be something it is like to have some mental state, one must be conscious of that state. Nagel (1974, 1979) thought this to be an analytic truth, and that he was characterizing what it is for a state to *be* a conscious mental state. There is something it is like to have a state if and only if that state is a state of consciousness. He used the phrase to point to the problem of consciousness, which he thought all the accounts standard at the time ignored. If we take Nagel at his word, these properties obtain in virtue of similarities and differences in what it is like to have various mental states. They are variations

in how states of which one is conscious appear to the one conscious of them; variations in the manner of appearance of mental states to one who is conscious of those states.

Suppose there is some property red* in virtue of which when one has a mental state with red*, something out there looks red. Red* will still fail to be a phenomenal property of this, the most intricate level. Instead we need some property red++, which characterizes what it is like to have a red* mental state. These are characteristics of one's consciousness of one's sensory states, so they neither look red nor suffice to make something in the surroundings look red. Instead they are closer to an individual mode of appearance of a mental state with the red* property.

In short, phenomenal properties were kicked indoors (they became properties of mental states), and then kicked upstairs (they came to characterize what it is like to be aware of one's mental states). It is certainly true that on the latter view, a state that has a phenomenal property must be a state of consciousness--in fact, it must be a property of a state in which one is conscious of one of one's mental states. But by the same token, this sort of phenomenal property is entirely distinct from what we are saying of the white paper when we say it looks red. Something looks red. One has a state of sensing something which looks red. One is (sometimes) conscious of being in a state of sensing something which looks red. And perhaps there is something distinctive and special about what it is like when one is conscious of being in a state of sensing something which looks red. But whatever that last special and distinctive property is, it certainly cannot be "looks red"! It is a property of a different bearer, in a different context, to explain different facts.

For now I propose we focus on the simplest variety: being aware of something that looks red. It gives us problems enough. The phenomenal property is the traditional sort: a characteristic of how things appear. The prototypical test is finding some stimulus arrangement in which things appear to have properties they do not in fact have. There is no need to get fancy about the assertion that sometimes things are not as they appear; "is", "is not", and "seems to be" need only their ordinary interpretations. Our target is, for example, a state in which one is aware of something which seems to be red but is not red. This counts as a state of "phenomenal consciousness" because it involves awareness. The latter involvement is also as simple as can be: one senses something, and is thereby made aware of something one senses. One need not be aware *that* the thing in question looks red; we might have just "thing" awareness, not awareness "that". Furthermore, there is no requirement that one also be conscious of seeing something or of seeming to see something. One might be aware merely of *what* one sees, and not also of the seeing of it. Perhaps the visual experience is so compelling that one is, as we say, lost in it: perhaps one is absorbed by the colours of the sunset, and is not at the same time conscious of seeing the sunset.

So the notion I want to focus is the very simplest of the ones that might qualify as "phenomenal consciousness". Conflating these different notions has

a high cost: the different versions have different bearers and their truth conditions differ. There is no logical or conceptual connection between the characteristic of a sensory episode in which something looks red and the distinct characteristic of a mental state in which there is something it is like for me to be conscious of sensing something that looks red. The cost is that we might confuse failure to explain one of these with failure to explain all of them. In fact I will argue that the simplest variety--being aware of something that looks red--is already within our grasp. We cannot explain all the different intricacies of different varieties of phenomenal consciousness, but if the simplest variety--being aware of something that looks red--still counts, then perhaps we can explain *a* variety of phenomenal consciousness.

IV. Preattentive Phenomenal Properties

The first step to showing this is to show that the traditional notion of phenomenal properties has instances within the preattentive domain. This is something of a surprise, particularly to those who think that phenomenal character requires consciousness. Some preattentively detected features can be shown to be phenomenal properties, in the traditional sense. There are preattentive illusions. To put it another way, there are sensory states that have phenomenal character in the traditional sense (character which, were one conscious of it, would yield perceptual reports such as "it seems to be moving" or "it looks like a triangle" or "it appears to have a shadow"), even though, (a) in these cases, nothing is moving, there is no triangle, and nothing is in shadow, and (b) the sensory states are preattentive, and hence unconscious. (The "hence" takes some showing.) In fact these locutions are the natural ones to apply to the three different examples of preattentive phenomenal properties that will be described: apparent motion, subjective contours (as in Kanizsa triangles), and pictorial depth cues. All three have been shown to allow efficient search (or "pop out") and can serve to attract attention, in an effortless fashion, to a target. So they seem to be detected prior to selection by selective attention. Indeed, their detection serves to direct the selective processes of selective attention.

Apparent motion is the simplest to describe. Flashing dots on a screen can give the appearance of motion, depending on the exact parameters of angular separation, brightness, and timing. If the interval is too long, subjects will see one flash, then a second, distinct flash. If it is too short, subjects will see two distinct dots flash nearly simultaneously. But if it is just right, subjects will seem to see one dot that moves across the intervening space, landing at the point of what is in fact the second flash. Nothing in fact moves on the screen, but the display generates a compelling illusion of motion.

Description of apparent motion is a paradigm example of use of the traditional notion of phenomenal properties. "Apparent motion" is not real motion; it simply the contrary-to-fact appearance of motion. The dot that moves through the intervening space is a merely phenomenal dot; there is no luminescence moving across those portions of the screen.

So it is very interesting indeed to discover that apparent motion sustains "pop out" in multi-item displays (see Wolfe 1996b, 32-33). For example, Ivry and Cohen (1990) used displays in which "X"s could appear to move on the screen horizontally or vertically. Apparent motion was induced by alternating frames on the computer screen, with the lit-up pixels spaced and timed appropriately. They showed that a single "X" which appeared to move horizontally could be picked out quite efficiently from any number of distractors appearing to move vertically, as long as the apparent motion was over fairly short distances. Long-range apparent motion resulted in less efficient search. In another experiment the targets and distractors were "forms" constructed by filling in eight of sixteen pixels in a four by four grid. In the short-range condition these "forms" could change from frame to frame, and the pixel colors could also change, yet a target appearing to move in a unique direction among apparently moving distractors could still be detected efficiently. Detecting the target in displays using longer range apparent motion was less efficient, requiring an item-by-item search. Ivry and Cohen suggest that detection of short range motion is a process of early vision, performed in parallel across the entire visual field, while long range motion "requires an attentional process in which each object in the display is examined in a serial fashion" (1990, 329). Both are separable from the discrimination of form or color.

But apparent motion is merely phenomenal. Hence a merely phenomenal property--*appearing* to move horizontally--sustains pop-out and effortless texture segmentation. It sounds portentous, but that's because of all the extra freight that the term "phenomenal" has acquired over the years. It used to be a clean little skiff but is now encrusted with barnacles and algae.

Subjective contours provide an even more compelling example. A "Kanizsa" square is not in fact a square, but instead a stimulus array that is cleverly designed so that it creates the impression of a square; it *looks* like a square. (See figure 1.).

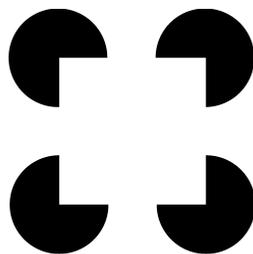


Figure 1. A Kanizsa square. A compelling illusion of something that appears to be a square, with continuous edges, lying on top of the four circles, and slightly brighter than the background.

At the vertices one places circles with cut-outs corresponding to the local details of each vertex. The circles with cut-outs are often called "pac men" because they look like the hero in the "Pacman" arcade game. The figure

creates a compelling illusion of continuous edges, running from vertex to vertex; and the illusory shape will also appear to be brighter than the background: a slightly brighter square, floating on top of the white paper. But if you examine it carefully you see that there are no line segments between the circles, and the white paper is the same shade throughout.

The edges (which appear to be there, but are not) are called "subjective contours". They reveal again the importance of edge-detection in early vision; here the mechanisms seem a bit over-enthusiastic. It does not take much to seem to see an edge; just a few cues will do. As Ramachandran (1987) and others have pointed out, in natural environments of dappled light and heterogeneous reflectance profiles, camouflage breaking would be considerably easier if mechanisms for edge detection were quick on the draw, picking out boundaries on the basis of sparse data. Kanizsa figures show just how sparse the data can be.

Our terminology for "subjective contours" is another example of good, clean, traditional talk of phenomenal properties. A subjective contour is not real; it is merely the appearance of a contour. There is no real edge at that place; the appearance is an illusion. So it is a surprise to find that subjective contours can be detected in parallel, preattentive stages of vision. The simpler sort of experiment shows "pop out" for merely subjective polygons. For example, Davis and Driver (1994) used a stimulus array in which some portions had the "pac men" circles aligned so as to create the impression of a Kanizsa square, while in other portions the circles were rotated ninety degrees. In the latter condition the cut-outs are not correctly aligned, no potential edge is given two vertices, and no impression of a subjective square arises. They found efficient search and effortless segmentation of the subjective square, regardless of the number of non-subjective (e.g., merely actual!) distractors. Enclosing the "pac man" circles with a boundary also disrupts the formation of subjective contours, and Davis and Driver showed that the same stimulus, but with pac-men encircled, also gave inefficient search. There was still one place where the alignment was correct, but without the subjective contours it had to be found by laborious, item-by-item scanning.

Davis and Driver (1998) provided further controls to show that subjective polygons can indeed be detected in parallel. That is, arrays that presented multiple subjective figures simultaneously, and in which the search target was defined as a feature in just one of those subjective figures, still allowed efficient search for the target. The target was a brown circle with cut-out, said cut-out in some trials aligned as a corner of a Kanizsa square. Non-targets were whole circles, which might or might not occupy the position of a corner for a potential square. The really clever bit was to present the notched circles by stereoscopic fusion, adjusted so that they would appear to lie either above or below the plane on which the subjective squares seemed to be located. If above the plane, then the subject would seem to see a notched circle, and through the notch of it the fourth corner of a subjective square would be visible. If below the plane, then the subject would seem to see a circle *occluded* by the corner of the

subjective square. In the latter case it would not be clear to the subject that the circle was a notched circle: it would tend to appear as a full circle, with a portion occluded. It would appear to be occluded by a figure that is itself merely "subjective".

Results confirmed that, as Davis and Driver put it, "the presence of the notched large circle target is not immediately apparent, because of amodal completion behind the abutting subjective square" (Davis and Driver 1998, 176). That is, with the notched circles apparently lying in plane behind the subjective squares, search for a notched one is difficult and inefficient. Each circle near a subjective square must be examined in turn, since it might be a full circle, partially occluded by the subjective square that seems to lie on top of it. Whereas when presented to appear above the plane of the subjective squares, search for one notched circle is quite efficient. There were other important controls, but the gist is that multiple subjective figures are indeed coded in an obligatory, parallel fashion, as otherwise there would be no apparent occlusion, and no inefficient search.

The discussion of results invokes not only squares that are not there, but also three different apparent depths: above the screen, on the screen (where the unreal squares appear) or below the screen. As mentioned, this impression was produced stereoscopically: the "brown" circles were not really brown, but instead were the product of stereoscopic fusion of red and green circles, seen through red-green 3-D spectacles. This abuts, but does not encroach upon, my third category of examples: experiments showing preattentive parallel extraction of *pictorial* depth cues. A pictorial depth cue is a cue that can be presented on a flat piece of paper, without any special spectacles required, but which suffices to give an appearance of differing depths. Perspective and occlusion cues are well-known. Early work on machine vision at MIT explored how the the exact topology of intersections (Y joints v. T joints) could be informative, and it was a target of analysis for early vision in Marr's models as well. (A T vs. a Y is a nice local feature that automatic mechanisms of early vision pick up; see Enns and Rensink 1991.) Shape from shading and surface contours are other, potent, indicators. Many of these yield pop out: occlusion, slant from texture, junction topology (T v Y junctions), shadows, and shape from shading, for example (see Wolfe 1996b, 39).

In the two dimensional realm there are no shadows, but merely darker or brighter patches of the picture. Nevertheless, careful arrangement of the lighter and darker patches can give a compelling impression of depth. Not all arrangements will do; some might have inconsistent cues for the direction of the light source, and others are inconsistent with an opaque three dimensional object which occludes some of that light. The surface which the shadow falls upon must be a possible three dimensional surface, and the way the shadow falls upon it yields "shape from shading" cues about *its* shape. There are further subtleties in variations of brightness, reflections, glare, light scatter, and diffusion, but these can be ignored for now. A picture that fails to satisfy enough of these constraints will fail to "look like" a normal object.

Given a picture that does succeed in depicting a three dimensional object, lit up from one direction, casting a shadow, it is possible to rearrange its light and dark patches to yield a stimulus that has exactly the same surface areas at the same reflectance levels, but which fails to "look like" a three dimensional object. This was one of the manipulations that Enns and Rensink (1990) used to test preattentive sensitivity to pictorial depth cues. It was found that in the 3d panel--that is, in the panel presenting what *appear to be* three dimensional cubes--the one cube that appears to be lit from the side "pops out" and can be identified efficiently, independently of the number of distractors. But the corresponding target in the 2d panel is difficult to find, and requires each element to be scrutinized in turn. Somehow the visual system can effortlessly pick out the image with valid shape-from-shading cues, but must resort to serial search if the appearance of depth is eliminated. So shape from shading, despite its geometrical complexity, is detected in automatic, parallel, preattentive processes.

But just as a subjective contour is not a real contour, so the shape one seems to see when one looks at these pictures is not real. The object depicted appears to be three dimensional, while the picture is flat. As in our other examples, "pictorial depth cues" are cues to how things appear, not necessarily to how they are. And if the cue is *pictorial*, then it is derived from arranging patches on a two dimensional plane so as to look like a three dimensional object. In short, "looks like a cube" ascribes a phenomenal property, a characteristic of appearance. That figures so characterized can pop out preattentively hence yields another example of a preattentive phenomenal property.

V. Phenomenal consciousness, but hold the consciousness

It bears emphasis that all the features so far discussed are preattentive: representations of them are completed prior to, and without, the assistance of any perceptual or cognitive capacities that might be summoned by the activation of selective attention. For if basic features are the ones allowing "pop out" and effortless texture segmentation, then these features are the ones that help determine which parts of the optic array *attract* attention. They are features to which attention is drawn. It follows that their extraction must be complete--the sensing of them must have already happened--*before* attention is drawn to the vicinity.

Motion, for example, is a terrific attention-getter; and results in the last section showed that merely apparent motion can be as well. The one area of the screen where things appear to be moving horizontally grabs attention, and can be identified in finite time, more or less independently of the number of other items that appear to be moving vertically. The task does not require deliberation or effort: the horizontal motion is noticeable and salient, and attention is drawn to it, rather than being required for it. It is salient because it is the one region where apparent motion has a horizontal direction, rather than vertical.

It follows that the detection of the direction of apparent motion and of its singularity in one area must be completed prior to the moment that attention gets slewed to the area in question. The visual system can detect where things *appear* to be moving horizontally without one needing to focus attention on the candidates. As noted in section III, the selection process itself requires some sort of data and some priorities in terms of which the selection can be made; and what these phenomena are taken to show is that basic features help provide the data in terms of which the selections of selective attention are made. To revert to the analogy of that section: this is information that is passed along offstage, to the little guy behind the curtains who controls where the spotlight is to go next. It is information he gets *before* the spotlight is moved, and here (given the imminent pop out) it is enough to determine where it will be moved.

Suppose the director yells "freeze!" at exactly that moment. Our homunculus has just received intelligence that there is an interesting development in quadrant delta. The subject of the experiment is *about to* have attention drawn to quadrant delta, because that is the one region where apparent motion has a horizontal direction, rather than a vertical one. The appearance, and the singularity of its features in delta, must have already been registered; and that is *why* attention is about to be drawn to region delta. I suggest that not only is the subject not yet attending to developments in delta, but also that at that moment the subject is not yet aware of developments in delta.

It is worth laying out the argument for this claim. The first (essential) premise is that the situation in question is one in which we are testing for "pop-out": it is designed to test for the involuntary, stimulus-driven, exogenous "grabbing" of attention by sudden presentations of novel stimulus arrays. The critical time interval occurs immediately after the presentation of the new stimulus array, and it is brief enough that the subject does not have time to shift attention deliberately. Instead attention is about to be grabbed by something, at least briefly. The question is what directs that exogenous "grabbing" of attention.

Now if attention shifts involuntarily in this way, to an abruptly presented novel stimulus, then one only becomes aware of that stimulus after attention has shifted to it. Because it is novel, it wasn't even in the "background" previously; one could not have attended to it at any moment prior to the moment of its presentation. After that moment, one attends to it as quickly as possible. So if we have "pop-out" of some feature, then one is not aware of that popped-out feature until one focuses attention on it.

The other interesting fact about pop-out is that it is not random. Attention is drawn to the unique instance of basic feature F , which stands out among the many instances of basic feature G . This implies that information about the difference between F and G , and about the uniqueness of the occurrence of F , must be registered and received by the mechanisms of selective attention, which then use it to direct attention to that one instance of F . So if we have "pop out" of some feature, then the presence of the feature must be detected by the

selection mechanism, and registered in the salience map as the next target, *before* attention can be directed to it. It is so registered *in order* to direct attention to it.

Here then are the two premises:

- (1). If a feature pops out, then one is not aware of it until one focuses attention on it.
- (2). If a feature pops out, then its presence is detected by the selection mechanism, and it is registered in the salience map as the next target, before attention can be directed to it.

From these we get a simple conclusion:

If a feature pops out, it can be registered and detected without awareness.

This is unremarkable until one notices that (as argued in section IV) the features F and G can include various pairs of phenomenal properties. F is "apparent motion in a horizontal direction" and G is "apparent motion in a vertical direction". That this F pops out in a field of those G 's indicates that the difference between "being appeared to horizontal-wise" and "being appeared to vertical-wise" can be registered and detected, and can serve to direct attention, even though one is not (yet) aware of any of the stimuli that appear in any of those ways, or of being appeared-to in any such way. But what we have, then, is a state in which the subject is being appeared-to, in a particular fashion, without thereby being aware of any aspect of that appearance, or of being in such a state. So we have a state of "phenomenal consciousness" without the consciousness. Being appeared-to and being aware-of split apart. We have the former without the latter. States of phenomenal consciousness are not states of one thing; the notion splits in twain. "Phenomenal" shows up--it can pop out--prior to consciousness.

The suggestion may seem outlandish and weird, so it may be worth reviewing the stages of the argument again, using one of the other examples. Talk of "subjective contours" clearly employs the traditional terminology of phenomenal properties. The pac-men create the appearance of a square, with bright contours; but there are no line segments on the page. So a Kanizsa square is a phenomenal square: a mere appearance-of a square. Now we abruptly present a subject a novel stimulus array in which there are many groups of notched circles, but in only one such group are they aligned so as to generate the Kanizsa illusion. That group will "pop out"--it will draw attention to itself, in more or less constant time, no matter how many distractor groups are also present. So, the psychologists conclude, subjective contours are basic features that are processed preattentively, in parallel over the entire visual field. The system registers the appearance of a subjective square in region delta, prior to, and without, the ministrations of selective attention. Indeed, it is because that appearance is unique that attention is about to be drawn to that region. So, in these delicate moments, the subject is being appeared-to in a determinate fashion--so as be presented the appearance of a Kanizsa square in region delta--yet is not aware of those contours or of that illusory square. Being appeared-to can occur in a determinate way without the subject being aware-of that way.

I earlier posed the question: where, in the architecture of visual processing, are we to find states that have all the properties of states of phenomenal consciousness? Now we have an answer. If states that have phenomenal character split apart from those that yield awareness, then there is no one place in that architecture in which states have all the properties of states of phenomenal consciousness. Furthermore, states in the preattentive portion of the architecture can satisfy the sufficient conditions for the traditional notions of phenomenal character: states therein can be states of being appeared-to in a determinate way. Phenomenal character is found early, in preattentive feature processing. States in virtue of which the subject is aware of features of the ambient array presumably arise later, in the "attentional" or "post-attentional" portions of the architecture. They are certainly not yet found in the preattentive bits.

VI. Philosophical implications of divorce

Phenomenal character can be entirely divorced from awareness. I will call this the "divorce" hypothesis. It has implications for theoretical arguments found within both philosophy and psychology. Take the philosophical discussions first. A widely held assumption, common to all sides of various philosophical debates, must be abandoned. If the divorce hypothesis is true, then there is no necessary connection between phenomenal properties and consciousness. What havoc this finding wreaks within recent discussions! For philosophers almost invariably assume that understanding phenomenal character and understanding consciousness are two labels for the same endeavor.

Consider this sampling of recent discussions. As exhibit number one, here is Chalmers describing "the hard part" of the mind-body problem:

A mental state is conscious if there is something it is like to be in that mental state. To put it another way, we can say that a mental state is conscious if it has a *qualitative feel*--an associated quality of experience. These qualitative feels are also known as phenomenal properties, or *qualia* for short. The problem of explaining these phenomenal properties is just the problem of explaining consciousness. This is the really hard part of the mind-body problem. (Chalmers 1996, 4)

"The hard problem" is a wonderful phrase for marketing purposes, and it has engendered a sprawling literature. But we can now see that it has been ambiguous since its inception. Is the hard problem the problem of explaining phenomenal properties, or is it the problem of explaining consciousness of phenomenal properties? They are not the same thing! So when Chalmers says "Color sensations stand out as the paradigm examples of conscious experience" (Chalmers 1996, 6) one must say: not always. To understand the phenomenal properties of chromatic appearance, it is not necessary that we solve the problem of consciousness. It is conceptual error to claim that sensory appearance is necessarily tied to conscious experience. Similarly, it is simply a mistake to say

the phenomenal concept of mind ... is the concept of mind as conscious experience, and of a mental state as a consciously experienced mental state. (Chalmers 1996, 11)

"Phenomenal" and "conscious" are *two* concepts, and they can be divorced. The preattentive visual architecture provides examples in which they *are* divorced.

Now to be fair to Chalmers, these propositions are not entirely his fault; the notion of "phenomenal" is itself ambiguous, and there *are* readings in which phenomenal properties are by definition confined to states of awareness. (They might be understood as the manner of appearance of one's own mental states when one is conscious of them, for example.) Much of the blame can be laid at the feet of Thomas Nagel, author of the fecund phrase "what it is like to be a bat". Nagel took this to express a difference between states of consciousness and other sorts of states, yet it has been widely understood as picking out the sprawling class of states that possess phenomenal character. Now Nagel was successful in articulating one of the murky features of our notion of consciousness: there is something it is like to have state M if and only if one is conscious of M. So the misunderstanding is predictable: if "what it is like" characterizes a phenomenal property, then phenomenal properties become properties exclusively of conscious mental states.

The phrase "phenomenal consciousness" was introduced into recent discussions by Ned Block, and some of these assumptions were found in that introduction. For example, while contrasting phenomenal consciousness (P-consciousness) and access consciousness (A-conscious) he says:

it is in virtue of its phenomenal content or the phenomenal aspect of its content that a state is P-conscious, whereas it is in virtue of its representational content, or the representational aspect of its content, that a state is A-conscious. ... The P-conscious content of a state is the totality of the state's experiential properties, what it is like to be in that state. (Block 1997, 383)

He also urged that "phenomenal properties cannot be P-unconscious" (406). As with Chalmers, one can chalk this up to a different notion of "phenomenal content": the sort that Block has in mind is not the simple sort we find while looking innocently at Kanizsa squares. Instead Block's sort must explain what it is like to be in the state, e.g., one's own mode of presentation of the mental state itself. In fact Block has more recently allowed that "phenomenality" can be divorced from consciousness, so he would allow amendments for both such claims.

VII. Perceptual psychology after divorce

The possibility of divorcing phenomenal character from awareness also has some startling implications for some of the models and theoretical disputes among perceptual psychologists. Candidates for divorce are far more numerous than the three described in section IV.

For one, there are other preattentive processes that clearly involve phenomenal properties as traditionally understood. The examples of "shape from shading" and pictorial depth cues raise a number of hotly contested issues. For example: to what extent can perceptual grouping processes--the capacity to

perceive a set of elements *as* a group, a unity, or a structured whole--proceed preattentively? Is it only form primitives that can be sensed preattentively, or can more complex structural relations of overall shape also be found? One surprise with shape-from-shading was that features of such complexity could be handled in early, automatic, parallel processes; it had been thought that 3d shape was a higher-level feature whose analysis was complex, and done later. But perceptual grouping processes in general and these shape effects in particular invoke classical phenomenal properties. When an ambiguous figure shifts in aspect--it looked like a duck, and now it looks like a rabbit--what changes is nothing in the stimulus itself, but how it appears. And so all the results that suggest preattentive capacities for grouping and segmentation could be added as additional examples of potentially preattentive phenomenal properties.

Divorce also allows some otherwise odd results in priming and masking studies to be redescribed in simple terms. One relatively direct way to approach the study of perception without awareness is to "mask" some stimuli, and test whether they still have a "priming" effect of some kind. Masking is used to eliminate awareness of the stimulus; priming to show that it is, nevertheless, perceived. A stimulus might be presented for a very short duration, and immediately followed by a different stimulus, which perhaps shares a boundary with the first one. With durations cut short enough, subjects will not report being aware of the first stimulus at all. It may nevertheless affect later perceptions in ways that indicate it was perceived.

Many such "priming" effects can be found; for the purposes of this paper the most interesting derive from the congruity or incongruity of colors. Complementary colors, or more specifically, opponent colors (red *v.* green, yellow *v.* blue) are counted as "incongruous". Whereas red and orange share a red component, and are "congruous". If one is asked to name the color of a patch, and an incongruous color is flashed along the way, there is interference in producing the name: subjects slow down, and errors increase. Flashing a congruous color does not produce the interference. A related effect is called "Stroop" interference. The task is to name the color of a patch, and the interference is produced by also presenting an incongruous color *word*. The color patch might be green, but if you flash the word "red" many subjects will call it "red". What is interesting is that both sorts of results are also found if the presentation of the interfering stimulus is masked to the point where subjects do not report being aware of it (see Cheesman and Merikle 1985, 1986; Cohen, Ivry, Rafal, and Kohn, 1995).

Incongruity of color is a good example of what the tradition would call a relation between appearances. "Being a complement of" is a phenomenal relation. Interference is explained by the incompatibility of appearances. If this account is to be viable, then the subject must apprehend the appearances. We get a delayed response, for example, because first the screen looks green, then it looks red; and green and red are incongruous. Now suppose we mask the first stimulus to the point where the subject is at chance in guessing whether

there was or was not a stimulus present. We get a reduced effect size, but one that is still statistically significant. The preferred conclusion seems, at first, shocking. A subject can be appeared-to greenly, without being aware of the green, or of the stimulus that appears green, or even of the fact that there was a stimulus at all. But with the wrinkles ironed out of our conceptual wardrobe, these results can now be clothed in style. They are just another example of being appeared-to without being aware-of. The subject is being appeared-to in a determinate fashion--to get the results, the character of that appearance must be, specifically, *green*--but the subject is not aware of that appearance. The sensory state has phenomenal character of which the subject is not conscious. This shows, as definitively as anything could, that color sensations are *not* the paradigm example of conscious mental states. Here we have what seem to be chromatic sensory states, but (apparently) no awareness of the color of the masked stimulus. Both conjuncts are subject to some doubt, but if confirmed, we have a case of color sensation with no consciousness thereof.

Third and last in the list of potentially redescrivable phenomena are some syndromes found in the neuropsychology clinic. Blindsight and the neglect syndrome provide some of the most compelling evidence for dissociations between perception and awareness. In both, perceptual capacities seem to be retained even though awareness of what is perceived is lost. The two syndromes pose rather different research problems. In blindsight the surprise is not so much that damage to area V1 causes a scotoma--a region of visual perimetry within which the patient cannot report upon the features, or even the presence or absence of, visual stimuli. The surprise is that even in the absence of "acknowledged awareness", patients proved capable of successful forced choice discriminations: of location, direction of motion (vertical v. horizontal), form primitives (X v O), and, most surprising of all, color. The focus of research has been on how this residual perceptual function is possible given the damage to V1. In contrast, the retention of perceptual capacity in neglect is not so problematic, and there the issue is rather why there is such a puzzling loss of awareness of stimuli in portions of the visual field (see Driver & Vuilleumier 2001).

To take these in turn. Residual chromatic discrimination in blindsight was demonstrated by an arduous procedure of forced-choice discrimination of pairs of stimuli presented in the patient's scotoma. The experiments are done in what Weiskrantz calls the "pure unawareness" mode: the patient will report no awareness of any feature or any event within the scotoma. They cannot tell when a stimulus is present or absent, much less anything about its features, and must be prompted to respond. Nevertheless, over a long, long series of what seemed to the patients to be forced-choice guesses, chromatic discrimination can be demonstrated, and--the real stunner--sensitivity to different wavelengths can be shown to be close in form to a normal "spectral sensitivity" curve (see Stoerig and Cowey 1989, 1991, 1992). In daylight (photopic) intensities, with cones active, the visual system is most sensitive to wavelengths around 550 nm. In scotopic conditions, (night vision, with just the rods active), the peak

sensitivity shifts to roughly 500 nm. Cowey and Stoerig demonstrated the same shift in spectral sensitivity for stimuli presented within the scotoma of blindsight patients, and they demonstrated that in photopic conditions the spectral sensitivity shows peaks and valleys similar to the normal curve, indicating the contribution of more than one type of cone. So even though patients may not be aware of the presence of visual stimuli in the scotoma, when prompted to guess, their responses indicate some capacities to discriminate among wavelengths similar to normal ones. As Weiskrantz remarks:

The latter capacity--discrimination of colours--presses credulity to the limit, because in those tests--which by their nature were very time-consuming and lasted for several days--the subjects uniformly and consistently denied seeing colour at all, and yet performed reliably above chance, even between wavelengths falling relatively close together. Moreover, the fine-grained features of the spectral sensitivity curves of these subjects (carried out, again, by forced-choice guessing) suggested that wavelength opponency, that is, colour contrast, was intact (Stoerig and Cowey, 1989, 1991, 1992). The subjects seemed to be able to respond to the stimuli that would normally generate the philosophers' favourite species of 'qualia', namely colours, but in the absence of the very qualia themselves! (Weiskrantz 1997, 23).

I have avoided the term "qualia" in this paper, preferring here the traditional notion of "phenomenal properties", but even in the latter terms we get what might seem a paradoxical result. The shift between photopic and scotopic sensitivity is widely thought to explain a perceptual effect known as the "Purkinje shift". In the daytime the red flower of a rose might look as bright as its green leaves, but as dusk falls the relative brightness of different colors appears to shift. The red flower will come to look black, and the green leaves will appear noticeably brighter than the flower. This is a shift in sensible appearance, in how things look; the change in relative brightness is a change of phenomenal properties. But the standard explanation for it appeals to the differing spectral sensitivity curves of photopic and scotopic vision. Photopic sensitivity peaks at 550 nm, and the red and green surfaces are equally far down from that peak, and so appear equally bright. Scotopic vision is maximally sensitive at about 500 nm, and the green leaf happens to reflect light maximally in that range. Whereas nighttime vision is quite insensitive to wavelengths reflected at the red end of spectrum. So at dusk the green leaf will look brighter than the red of the rose.

The Stoerig and Cowey results indicate that these mechanisms are at least partially intact in blindsight. The wavelength discriminations that survive could manifest the same shift in sensitivity between scotopic and photopic conditions. So it is conceivable that even within a scotoma, the green leaf looks brighter than the red rose in scotopic conditions, while they look equally bright in photopic conditions. If, as argued earlier, it is possible to be appeared-to in a particular way without being aware of any aspect of that appearance, then it is possible that *G* looks brighter than *R*, even though the subject is not aware of *G*, aware of *R*, or aware of any feature of *G* or *R*. Perhaps blindsight is precisely a case in which being appeared-to is dissociated from being aware-of.

Blindsight might or might not entail the absence of qualia, depending on whether one considers qualia to be phenomenal properties or the awareness of phenomenal properties. The latter sort are missing, but the former might still be present. The discrimination data summarized in spectral sensitivity curves would be taken to show that surfaces that appear equally bright to this subject in the daytime will not appear so at dusk, and that the one which reflects a preponderance of wavelengths around 500nm will look brighter than the one that reflects longer wavelengths. Now we find those same discrimination data replicated for stimuli presented within the scotoma of a patient with blindsight. Suppose being appeared-to can be divorced from being aware-of. It might then still be true of a subject that one stimulus looks brighter than another, though that subject has blindsight, and is not aware of either stimulus, or of any aspect of their appearances. Perhaps the problem is not loss of the appearances at all, but instead loss of awareness of them. Put in terms of "qualia", perhaps these subjects still have chromatic qualia, but are not aware of them.

Hemi-neglect provides an even better example of this possibility. Unlike blindsight, neglect can be demonstrated even in the absence of any damage to primary sensory pathways (Driver & Vuilleumier 2001, 40). Recent accounts of neglect attribute the symptomatology to damage to mechanisms of selective attention, rather than to sensory loss. Indeed, Driver and Vuilleumier (2001) propose that the visual processing in neglect is similar to the preattentive processing of normals. The conjunctions and dissociations of symptoms are quite intriguing, and provide a good review of the architecture of early vision.

The basic syndrome is as follows:

a lack of awareness for sensory events located towards the contralesional side of space (e.g. towards the left following a right lesion), together with a loss of the orienting behaviors, exploratory search, and other actions that would normally be directed toward that side. Neglect patients often behave as if half of their world no longer exists. In daily life, they may be oblivious to objects and people on the neglected side of the room, may eat from only one side of their plate, read from only one end of a newspaper page, and make-up or shave only one side of their face. The spatial bias towards one side can also be apparent in many simple paper and pencil tests. When required to search for and mark all target shapes on a page, the patients may cancel only those towards the ipsilesional side. When bisecting a horizontal line, they may err towards that side, and when drawing from memory, or copying a picture, they may omit details from the contralesional side. (Driver & Vuilleumier 2001, 40)

Neglect is not all or none, but is graded, with the probability of neglecting a stimulus increasing as it is presented further and further in the contralesional direction. But "directions" and "the left side" are relative to a reference frame, and the reference frames in neglect are both variegated and surprisingly sophisticated. Suppose (to simplify the discussion) we have a patient with damage to right parietal cortex, who tends to neglect "the left side". This is not a loss of visual sensitivity to the left side of space; it is not hemianopia. Patients can still see things on the contralesional side; the problem is difficulty in shifting and maintaining attention to events on that side. Patients with "extinction" for example might still be able to see and consciously report on a

light isolated on the left side, if it has no competition on the right; the problem arises only when two lights are presented simultaneously, with the one on the right then "extinguishing" awareness of the one on the left. This can happen even if both lights are presented in the right (intact) hemifield. So the reference frame cannot be retinotopic. (The latter is clear anyway, since it would not explain why the left side of plates, newspapers, and one's own face might be ignored. A patient with hemianopia can compensate simply by shifting the eyes. Some neglect patients have no visual field cut at all!) Furthermore, stimuli affecting the same retinal point will be neglected to varying degrees, depending on the position of the eyeball, the head and the trunk. A stimulus might be neglected when the eyes are aimed forward in their sockets yet detected when they are shifted sideways, even though the stimulus is produced so as to affect the same retinal area on both occasions. Similar results are found with head motions and with trunk motions. So "the left side" might be the left side of the head, the left side of the trunk, or the left side of external space (Driver & Vuilleumier 2001, 46).

The spatial coordinate systems defining which side is "the left side" can be rather sophisticated! Even more intriguing is that some of these coordinate schemes must be object-centered: what is ignored is the left side of *an object*, and so it all depends on what the neglect patient sees *as* an object. For example, if a picture shows one plant that has two flowers, a neglect patient who is asked to copy it may copy only the flower on the right. But suppose we occlude parts of the picture, so that the same two flowers are shown detached from one another, on independent stems, and ask the patient to copy the picture. He or she will carefully draw two flowers, but neglect details on the left side of each one (see Bisiach and Vallar 2000, 481).

What happens if an object is tilted? Unless one tilts the head, the "left side" of the object will be tilted relative to every body-centered coordinate scheme. Driver (1996, 318) showed that neglect patients continue to neglect the "left side" of objects, even when the object is tilted, and the axis of symmetry defining the left side of the object is tilted relative to every body-centered coordinate scheme. He calls this "axis-based" neglect.

Neglect must be something other than a sensory deficit. In order to determine which bits of an object make up its left side, one must be able to perceive its entire contour. To neglect the left side of a plate, the left side of the flower, or the left sides of the two flowers, one must have some perception of the midline of these objects, and so one must perceive their entire contours. So neglect seems to require intact perception of contours, combined with a deficit in allocating attention. Driver and Vuilleumier propose that the main deficit is a bias in the competitions for selective attention. Neglect is most clearly manifest when multiple items are competing for selective attention; the bias favors items on the ipsilesional side, in whatever reference frame defines that "side".

Grouping effects can decrease neglect, on this account, because they decrease the competition for selective attention. The left side of the entire

plant will be ignored, if the entire assemblage is perceived as one plant; if instead its flowers are perceived as two detached flowers, the left sides of both such objects are neglected. Given our earlier concerns it is particularly fascinating to note that Kanizsa subjective figures can secure this advantage as well. Mattingley, Davis and Driver (1997) did experiments on neglect patients very similar to the Davis and Driver 1994 experiments on preattentive pop-out. In some cases the "pac men" were arranged so that the notches in the circles were aligned, and a subjective square would be apparent to normal viewers. In other cases the notches were misaligned, or the pacmen had circular borders, which would defeat the appearance of the subjective square. The results are dramatic:

extinction was virtually abolished when such arcs were removed, as the bilateral events then yielded a single subjective object (a bright white rectangle, apparently superimposed on the black circles) due to modal surface completion. This suggests that extinction is reduced when the concurrent target events can be linked together into a single subjective object, becoming allies rather than competitors in the bid to attract attention. (Driver & Vuilleumier 2001, 52)

Subjective contours and subjective objects are paradigm examples of entities whose properties are phenomenal, and so it is gratifying to find such creatures of appearance in a place where awareness is less than full.

To common sense a neglect patient seems oblivious or insensible to the contralesional side of the current goings-on, where "current" changes with the context. Yet now we see that what appears to be blindness can be produced by failures in the mechanisms that regulate competition for selective attention. If a subject is truly incapable of shifting attention to anything visible on the far left side, however "left side" is defined, then in many respects that subject is as good as blind to developments in that region. If a problem with selective attention implies that visual sensory states registering such developments cannot ever win in the competitive allocation, then the subject remains perpetually unaware of such phenomena. Yet sensory capacities remain; this would be a kind of blindness deriving from a failure in attentional mechanisms.

What we have, then, might be the full preattentive sensory capacities without the capacities to engage the selection that leads to post-attentive processing. Like the sensory states found in preattentive processing, neglect patients could have states in which they are being appeared-to in a determinate fashion, even though they are entirely unaware of those appearances. Unlike blindsight patients, the neglect patient might have no visual scotomata, and primary sensory channels might be entirely intact.

VIII. Conclusions

Different authors mean different things by the term "phenomenal property". The traditional notion is one used to characterize sensory illusions and perceptual effects, and is already familiar, well-behaved, and well-ensconced in the models of perceptual psychologists.

Empirical models of perception provide explanations of at least some of

these phenomenal properties. Some of these explanations cite neural mechanisms as explananda.

Phenomenal properties as traditionally understood are almost invariably accompanied with awareness. That is, to speak plainly, usually people are aware of how things appear to them. They will tell you if you ask. Furthermore, any example of a phenomenal property identifiable from the first person perspective is necessarily one of which the author is aware. But there is no conceptual necessity--no necessity linking the concepts of--states of being appeared-to and states of being aware-of.

So it is conceptually possible that there exist states of being appeared-to which are not, in any way, conscious states. Some of the states identified in the preattentive processing of apparent motion, subjective contours, and pictorial depth cues show that this possibility is real. Some of the states in blindsight and hemi-neglect may be examples of the same thing.

It is therefore a conceptual error to identify the problem of understanding phenomenal properties with the problem of consciousness. If that conjunction is identified as "the hard problem" of consciousness, then the solution to the hard problem of consciousness, like the solution to the problem of life, will be seen in the vanishing of the problem.

Psychologists already employ tools that are adequate to the understanding of phenomenal properties as traditionally understood. They provide explanations of some perceptual effects, and for all we know some of these putative explanations could be sound.

Philosophers who use the term "phenomenal properties" to beat up on psychologists mostly mean something higher-order, such as: how mental states appear to someone who is conscious of them. These higher-order notions, while endlessly fascinating, have yet to be adequately clarified.

If we confine ourselves to the task of explaining phenomenal properties as traditionally understood, then *enormous* progress has already been made within the relatively short lifespan of the discipline of experimental psychology. An optimist can be forgiven for believing that the tools for its solution may already be at hand.

References

- Allport, Alan (1993). Attention and control: Have we been asking the wrong questions? A critical review of twenty-five years. *Attention and Performance 14* (3): 183-218.
- Bisiach, Edoardo and Vallar, Giuseppe (2000). Unilateral neglect in humans. In Rizolatti, Giacomo (ed.) *Handbook of Neuropsychology*. 2nd edition. Vol 1. Series Editors Francois Boller and Jordan Grafman. Amsterdam: Elsevier, 458-502.
- Block, Ned (1997a). On a confusion about a function of consciousness. In Block, Ned, Flanagan, Owen, and Güzeldere, Güven (eds.) (1997). *The Nature of Consciousness: Philosophical Debates*. Cambridge, Massachusetts: MIT Press, 375-415. (Originally appeared in *Behavioral and Brain Sciences 18* (1995): 227-288.)
- Block, Ned (2003). Consciousness, Philosophical Issues About. In Lynn Nadel (ed.) *The Encyclopedia of Cognitive Science*. London: MacMillan. Volume 1, 760-770.
- Broadbent, Donald E. (1958). *Perception and Communication*. Oxford: Pergamon.
- Chalmers, David J. (1996). *The Conscious Mind*. New York: Oxford University Press.
- Cheesman, J. and Merikle, P. M. (1986). Distinguishing conscious from unconscious perceptual processes. *Canadian Journal of Psychology 40*(4): 343-367.
- Cheesman, J. and Merikle, P. M. (1985). Priming with and without awareness. *Perception and Psychophysics 36*: 387-395.
- Chisholm, Roderick (1957). *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press.
- Cohen, A., Ivry, R. B., Rafal, R. D. and Kohn, C. (1995). Activating response codes by stimuli in the neglected visual field. *Neuropsychology 9*(2): 165-173.
- Davis, Greg and Driver, Jon (1998). Kanizsa subjective figures can act as occluding surfaces at parallel stages of visual search. *Journal of Experimental Psychology: Human Perception and Performance 24*(1): 169-184.
- Davis, Greg and Driver, Jon (1994). Parallel detection of Kanizsa subjective figures in the human visual system. *Nature 371* (27 October) 791-793.
- Desimone R. and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience 18*: 193-222.
- Di Lollo, Vincent, Enns, James T. and Rensink, R. A. (2000). Competition for consciousness among visual events: the psychophysics of re-entrant visual processes. *Journal of Experimental Psychology: General 129* (4): 481-507.
- Driver, Jon (1996). The neuropsychology of spatial attention. In Harold Pashler (ed.), *Attention*. London: University College London Press, 297-340.

- Driver, Jon and Vuilleumier, Patrik (2001). Perceptual awareness and its loss in unilateral neglect and extinction. *Cognition* 79: 39-88.
- Duncan, John (1998). Converging levels of analysis in the cognitive neuroscience of visual attention. *Philosophical Transactions of the Royal Society of London B* 353: 1307-1317.
- Enns, J. T. (1990). Three dimensional features that pop out in visual search. In D. Brogan (ed.) *Visual Search*. London: Taylor and Francis, 37-45.
- Enns, J. T. and Rensink, R. A. (1990). Influence of scene-based properties on visual search. *Science* 247 (Feb. 9, n4943): 721-723.
- Enns, J. T. and Rensink, R. A. (1991). Preattentive recovery of three dimensional orientation from line drawings. *Psychological Review* 98(3): 335-351
- Enns, J. T. and Rensink, R. A. (1993). A model for the rapid interpretation of line drawings in early vision. In David Brogan, Alasair Gale, and Karen Carr (eds). *Visual Search II*. London: Taylor and Francis, 73-90.
- Farah, Martha (1990). *Visual Agnosia: Disorders of Object Recognition and What They Tell Us about Normal Vision*. Cambridge, Massachusetts: MIT Press.
- Fodor, Jerry A. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Goodman, Nelson (1977). *The Structure of Appearance*. 3rd edn. Boston: Dordrecht Reidel.
- Humphreys, Glyn W., Keulers, Nicole, and Donnelly, Nick (1994). Parallel visual coding in three dimensions. *Perception* 23 (4): 453-470.
- Ivry, Richard B. and Cohen, Asher (1990). Dissociation of short and long-range apparent motion in visual search. *Journal of Experimental Psychology: Human Perception and Performance* 16 (2), 317-331.
- Julesz, Bela (1984). A brief summary of the texton theory of human vision. *Trends in Neuroscience* 7: 41-45.
- Julesz, Bela and Bergen, J. R. (1983). Textons, the fundamental elements in preattentive vision and perceptions of texture. *Bell Systems Technical Journal* 62: 1619-1646.
- Livingstone, M. S. and Hubel, D. H. (1987). Psychophysical evidence separate channels for the perception of form, color, movement, and depth. *Journal of Neuroscience* 7: 3416-3468.
- Luck, S. J., Vogel, E. K., and Shapiro, K. L. (1996). Word meanings can be accessed but not reported during the attentional blink. *Nature* 383: 616-618.
- Mattingley, J. B., Davis, G., & Driver, J. (1997). Preattentive filling-in of visual surfaces in parietal extinction. *Science* 275: 671-674.
- Marr, David (1982). *Vision*. San Francisco: W. H. Freeman and Company.
- Milner, A. D., and Rugg, M. D. (1992). *The Neuropsychology of Consciousness*. London: Academic Press.
- Nagel, Thomas (1974). What is it like to be a bat? *Philosophical Review* 83: 435-50. Reprinted in Nagel 1979, 165-180.

- Nagel, Thomas (1979). *Mortal Questions*. Cambridge: Cambridge University Press.
- Neisser, Ulric (1967). *Cognitive Psychology*. New York: Appleton Century Crofts.
- Posner, Michael I. (1994). Attention: the mechanisms of consciousness. *Proceedings of the National Academy of Science USA*, 91 (August): 7398-7403.
- Posner, Michael I. and Petersen, E. E. (1990). The attentional system of the human brain. *Annual Review of Neuroscience* 13: 25-42.
- Posner, Michael I. and Rothbart, Mary K. (1992). Attentional mechanisms and conscious experience. In A. D. Milner and M. D. Rugg (eds.), *The Neuropsychology of Consciousness*. London: Academic Press, 91-111.
- Ramachandran, V. S. (1988). Perceiving shape from shading. *Scientific American* 259: 76-83.
- Ramachandran, V. S. (1987). Visual perception of surfaces: A biological theory. In S. Petry and G. E. Meyer (eds). *The Perception of Illusory Contours*. New York: Springer Verlag, 93-108.
- Raymond, J. E., Shapiro, K. L., and Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: an attentional blink? *Journal of Experimental Psychology: Human Perception and Performance* 18: 849-860.
- Raymond, J. E., Shapiro, K. L., and Arnell, K. M. (1995). Similarity determines the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance* 21: 653-662.
- Rensink, R. A. and Enns, J. T. (1995) Pre-emption effects in visual search: Evidence for low level grouping. *Psychological Review* 102(1): 101-130.
- Rensink, R. A. (2003). Visual attention. In Lyn Nadel (ed.), *The Encyclopedia of Cognitive Science*. (London: MacMillan), vol 4., 509-515.
- Rensink, R. A. (2000). Seeing, sensing, and scrutinizing. *Vision Research* 40: 1469-1487.
- Sellars, Wilfrid (1963). *Science, Perception, and Reality*. London: Routledge & Kegan Paul.
- Stoerig P. and Cowey, A. (1989). Wavelength sensitivity in blindsight. *Nature* 342: 916-918.
- Stoerig, P. and Cowey, A. (1991). Increment threshold spectral sensitivity in blindsight: Evidence for colour opponency. *Brain* 114: 1487-1512.
- Stoerig, P. and Cowey, A. (1992). Wavelength sensitivity in blindsight. *Brain* 115: 425-444.
- Stoerig, P. and Cowey, A. (1991). Increment threshold spectral sensitivity in blindsight: Evidence for colour opponency. *Brain* 114: 1487-1512.
- Treisman, Anne (1988). Features and objects: The fourteenth annual Bartlett memorial lecture. *Quarterly Journal of Experimental Psychology A* 40: 201-237.
- Treisman, Anne (1993). The perception of features and objects. In A. Baddeley and L. Weiskrantz (eds.), *Attention: Selection, Awareness, and Control: A*

- Tribute to Donald Broadbent*. Oxford: Clarendon Press, 5-35.
- Treisman, Anne (1996). The binding problem. *Current Opinion in Neurobiology*, 6: 171-178.
- Treisman, Anne (1998). Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society of London B* 353: 1295-1306.
- Treisman, Anne and Gelade, Garry (1980). A feature-integration theory of attention. *Cognitive Psychology* 12: 97-136.
- Umiltà, Carlo (2000). Visuospatial attention. In Rizzolatti, Giacomo (ed.) *Handbook of Neuropsychology*. 2nd edition. Vol 1. Series Editors Francois Boller and Jordan Grafman. Amsterdam: Elsevier, 2000, 393-425.
- Weiskrantz, Lawrence (1997). *Consciousness Lost and Found*. Oxford: Oxford University Press.
- Wolfe, Jeremy M., Cave, K. R., and Franzel, S. L. (1989) Guided search: an alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance* 15 (3): 419-433.
- Wolfe, Jeremy M., Friedman-Hill, S. R., Stewart, M. I., and O'Connell, K. M. (1992a). The role of categorization in visual search for orientation. *Journal of Experimental Psychology: Human Perception and Performance* 18(1): 34-49.
- Wolfe, Jeremy M. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin and Review*, 1(2): 202-238.
- Wolfe, Jeremy M. (1996a) Extending guided search: why guided search needs a preattentive "item map". In Arthur F Kramer, Michael G. H. Cole, and Gordon D. Logan (eds.) *Converging operations in the study of visual selective attention*. Washington, DC: American Psychological Association, 247-270.
- Wolfe, Jeremy M. (1996b). Visual search. In Harold Pashler (ed.), *Attention*. London: University College London Press, 13-73.
- Wolfe, Jeremy M. and Bennett, Sara C. (1997). Preattentive object files: shapeless bundles of basic features. *Vision Research* 37 (1): 25-43.