Is What You See What You Get?: The “Filling In” Debate and its Implications for the Conception of Mind

by

Lyle Owen Crawford
B.A., University of Victoria, 2002

A Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of

MASTER OF ARTS

in the Department of Philosophy

Supervisor: Prof. Jeffrey E. Foss

ABSTRACT

The human visual system seems to create a richly detailed representation of the world with no distortions, intrusions or gaps, but these impressions are mistaken, as even simple experiments can demonstrate. Many of our perceptual deficits are so strikingly at odds with "the way things look" that some people suppose the brain must "fill in" the missing or corrupted information for consciousness. Contrasting positions in the debate over "filling in" are represented by the neuroscientist V.S. Ramachandran and the philosopher Daniel Dennett. Dennett criticizes Ramachandran's case for filling in, arguing that it is implausible and unnecessary, and that it presupposes a thoroughly discredited Cartesian concept of the self as a unitary, passive "viewer" of perceptions. These criticisms are well supported by evidence and argument. The tight link between the idea of filling in and the concept of the mind in which and for which filling in is supposed to occur means that the failure of this theory has radical implications for the nature of the mind. The mind is better described as a highly disunified, active process of inquiry.

Table of Contents

List of Figures

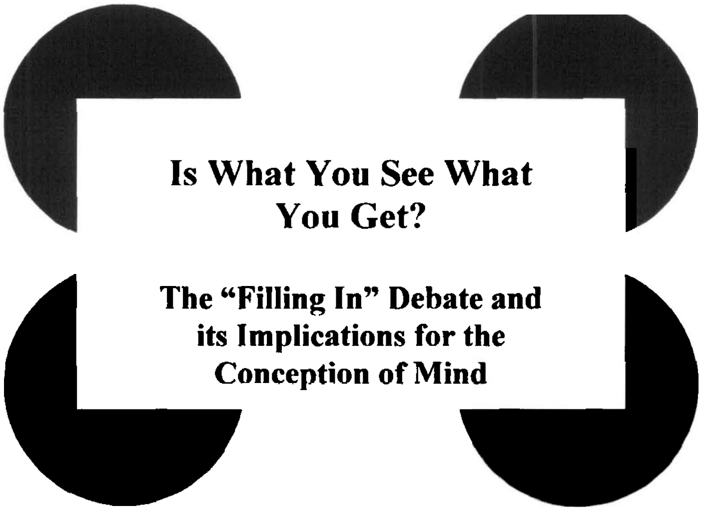

Is What You See What You Get?

The "Filling In" Debate and its Implications for the Conception of Mind

"It is remarkable concerning the operations of the mind, that, though most intimately present to us, yet, whenever they become the object of reflection, they seem involved in obscurity; nor can the eye readily find those lines and boundaries, which discriminate and distinguish them."

– David Hume, *An Enquiry Concerning Human Understanding*, §1[1]

"Get off the armchair!"

– Logo for Patricia Churchland's Experimental Philosophy Lab[2]

---

[1] Ed. Eric Steinberg. 2nd edition (Indianapolis/Cambridge: Hackett Publishing Company, 1993), p. 7

[2] http://philosophy.ucsd.edu/EPL/

Introduction: A Test Case for Conceptions of Mind

There is really nothing, it seems, that could be more intimately present to us than our own impressions of the world before our eyes. We are subject to a variety of well-known and not so well-known illusions, limitations and malfunctions, but the phenomenological integrity of a rich and *distinctly visual* experience seems on a footing about as sturdy as Descartes' *Cogito ergo sum.* In other words, whatever errors of judgement we may make in various circumstances natural or pathological, that the *way things look to us* is an experience of intrinsically visual information is so utterly basic that it seems to rest beneath even intuition, dubitable only, if at all, with the aid of abstract linguistic precision. The conception of mind as *that which perceives* sensory information of distinct kinds or modalities has been challenged in the debate over a fascinating, though at first glance philosophically innocuous, phenomenon. The phenomenon may be a general one, operating over different sense modalities, but it is most thoroughly documented and discussed in the case of vision.

Vision science in general seems to be a promising field for empirical investigations of mind and consciousness. As is sometimes noted, the retina is "an approachable part of the brain,"[3] the best opportunity to stimulate sensory neurons directly, precisely and naturally (i.e. non-electrically). Patricia Churchland remarks that so far more and better data exist for visual awareness phenomena – phenomena that "might reward the search for the neurobiological differences between being aware and

[3] Dowling, J.E. *The Retina: An Approachable Part of the Brain* (Cambridge, MA: Belknap/Harvard University Press, 1987)

not being aware in the awake, attentive animal"[4] – than for any other aspect of cognition. One such phenomenon is "filling in." Filling in is the standard term for a variety of ways in which the visual system apparently, and impressively, compensates for demonstrable deficits in the information that is collected by the retinas in our eyes. (It could, but does not usually, also refer to non-visual phenomena, for example, to equivalent auditory phenomena.) Things *looks like* they have more or less uniform (high) resolution and (full) colour detail throughout the visual field, and this field seems to contain no distortions, intrusions or gaps. But all of these impressions are mistaken, and they are all elements of vision for which it is alleged that the brain must and does fill in large amounts of appropriate, though not necessarily correct, information *for consciousness*. It seems as though the brain has a lot more work to do than just enjoying the show.

The term "filling in" is applied broadly, and this paper will consider types of filling in that differ from one another in what are arguably fundamental ways. The classic case is blind spot filling in, which occurs all the time in normal people with normal brains. Upon reflection, it can seem no less amazing than it must have in 1832, when the Victorian physicist Sir David Brewster concluded that it constituted a proof for the existence of God.[5] There are also examples found in people with certain types of brain damage, as well as other examples that occur in normal people but only under abnormal circumstances. As we will see, some of these examples are even more amazing.

---

[4] "Can Neurobiology Teach Us Anything about Consciousness?" Presidential Address to the *American Philosophical Association, Pacific Division* (March, 1993) *Proceedings and Adresses of the APA* (1994). Sec.3B. This is the consensus among cognitive scientists. See also, for example, Metzinger, T. (Ed.) *Neural Correlates of Consciousness: Empirical and Conceptual Questions* (Cambridge, MA: MIT Press, 2000), p.153

[5] Ramachandran, V.S. and Blakeslee, S. *Phantoms in the Brain* (New York: Quill/HarperCollins, 1998), p.273. Ramachandran raises the obvious question: however miraculous it may seem that the hole gets filled in, one is left to wonder why the "Divine Artificer" would have made a hole in the first place.

This thesis will discuss several ways in which visual experience seems to include filled in visual information, this information not being collected at the retina. It will then consider a debate, primarily between V.S. Ramachandran and Daniel Dennett, over the nature of filling in phenomena and, inevitably, the nature of the mind in which and for which filling in allegedly occurs. Ramachandran, on his own and with Churchland, argues that the brain really does fill in, in the full sense of adding visual information to the incoming veridical information so that the whole percept can become conscious as a distinctly visual experience. Dennett alleges that the intuitive "filling in" explanation, even when bolstered by informed neuroscience and the broader range of phenomena offered by Ramachandran, betrays a latent commitment to a Cartesian conception of mind, even when such a conception is explicitly disavowed. The brain is *not* just enjoying the show, says Dennett, but *nor is anyone else* at any special place within the brain– there is no one *for whom* the brain has to fill in. The view, or implicit assumption, that there is, when espoused by scientists who should (and do!) know better, he calls "Cartesian materialism."

Ramachandran and Churchland, already skeptical of the scientific tenability of many of Dennett's theories, are aware of his criticisms, and think that he simply prejudges neurobiological data yet to be collected and analyzed, whereas Dennett's position is, roughly, that nothing short of a complete overturning of the whole modern understanding of the brain could save "filling in." His complaints, along with the empirical data we will consider, force us to ask the traditionally "incoherent" question: might we sometimes be wrong not only about what we think we see (out in the world), but even about *what we see* (in our minds)? Filling in is thus not a mere curiosity of the visual system– it becomes a test case for different conceptions of mind, and calls into question those lines and boundaries which would distinguish even such seemingly basic operations of the mind as perceiving and judging.

This thesis argues that Dennett's criticisms of filling in are basically correct and non-trivial, although this will be qualified in some places since the variety of filling in phenomena means that some arguments that appear to be blanket criticisms can be misleading about some cases. Dennett's basic claim about the blind spot, that the brain ignores it rather than fills it in, will be defended with an analysis of "demonstrations" and other evidence offered in support of filling in. It will be argued that this evidence fails to establish that filling in occurs because, in light of certain reflections on how a visual system *would* be designed (by evolution) and how it *is* in fact designed, the best explanation of the introspective demonstrations is that the brain ignores the blind spot, and that basic intuitions about the spatial nature of visual perceptions are simply misleading. Further, both existing and plausibly imaginable neurological evidence fail to establish the "perceptual, *not* conceptual," or "interpolation, not (mere) inference," thesis argued for by Ramachandran and others.

Other examples of filling in, some of the "abnormal brain/normal context" variety, others of the "normal brain/abnormal context" variety, will also be challenged along Dennettian lines, although in most of these cases Dennett's idea of visual "labels" – a "paint-by-numbers" theory of visual representation – will prove more applicable. As with the blind spot, the responses to these examples will force us to consider some counterintuitive ideas about perception, and about our own seemingly secure judgements of "how it looks." We will see that vision extracts very basic features of scenes to construct representations, but these representations are not like painted pictures on an intuitive Euclidian canvas. They are not the sort of thing that needs to be filled in to "look like" they seem to, or rather, to account for our believing they look a certain way, since they do not really *look* like – visually appear to someone as – anything at all.

We will consider some surprising results of empirical investigations of the visual system. If interpreted by the intuitions that motivate claims of filling in the "normal" cases, these discoveries lead to the conclusion that our visual system fills in virtually everything we see, an incredible notion that, in any case, defeats the purpose of using a special term to pick out a special type of neural processing, and moreover has no theoretical advantage over Dennett's propositional labeling theory. One final set of experiments, using an "eye-tracker" apparatus, will count heavily against a major background assumption of filling in, that the brain uses information from successive fixations (of the fovea, or perhaps mental attention) to construct and maintain a rich, on-going visual representation of the world that is updated with information from the eyes and filled in with information from the brain. Not only are visual representations not the *kind* of thing that needs to be filled in– their content can be shown to be is so radically impoverished that the motivation for postulating a mechanism to produce extra visual detail loses its force. But the detail is manifestly there. We are not blind, something we can be confident of by our own and others' success at navigating and manipulating the world, if not, perhaps, by raw intuition to the contrary. The detail is "there" because it can, under evolutionarily significant circumstances, be checked (out in the world) as quickly as it can be asked about.

Is what we see what we get? The word "get" is ambiguous here. The upshot of this paper's conclusions about the filling in debate is that what we see is what we *retrieve*, not what we *receive*. We are left with a picture of a brain whose visual system creates propositional representational fragments on demand, to answer specific questions posed by the parts of the brain that need to know at any given moment. A visual representation is not and need not be a neurologically unified item just in virtue of being *visual* (and nor need it, at the deepest level, belong to a *sui generis* sensory modality). The filling in debate thus begins to reveal the mind as a highly disunified, active process of inquiry, not a passive perceiving *thing* to which integrated, non-propositional perceptions are delivered and presented.

What You See

Filling in obviously occurs in *some* sense or another. Most of the proof for this is right before our eyes (or perhaps right inside them, or even right behind them, depending on how one looks at it). Eventually we want to ask in what sense it occurs, but we can begin by considering the basic case for the claim that it does occur. The first premise in the argument for filling in is just: *this*– what it looks like. "What it looks like" is, as observed above, a more or less uniform (high) resolution and (full) colour visual field with no distortions, intrusions or gaps. This is quite striking in light of the second premise, which is that the visual field of a normal healthy person is actually characterized by highly non-uniform (veridical) resolution and colour content (having a useful degree of each only within a very tiny central "window"); furthermore, it does contain large distortions, intrusions and gaps. The conclusion seems obvious: the brain – somehow, somewhere – augments the incoming visual signal by performing quite sophisticated and extensive, but nevertheless fallible and discoverable, procedures of interpolation. In other words, it fills in what we see for us. Several notable features of vision are relevant here:

1) **Retinal blood vessels**. Most people who see a photograph of a retina must be confused. The most striking (macroscopic) feature is a web of blood vessels twisting over the surface. The famous "backward" design of the mammalian retina, with the receptors anchored in the back of the eye (so that light must travel through several layers of nerve fibres, ganglion cells, synapses and nuclei before it hits the receptor cells), means that these blood vessels occlude portions of the projected image, casting shadows

over the retina. Their presence can be revealed in a simple but dramatic demonstration: while looking at a plain surface, place a pen light at the corner of one eye and jiggle it up and down. When the light enters through the cornea at the right angle, the blood vessels suddenly become visible. They look, not surprisingly upon reflection, like a web of dark lines twisting across the visual field. Here is what a recent textbook on human vision, endorsed by leading vision and neuroscientists, remarks about this: "The reason we do not normally see them is that the brain adapts completely to their presence and fills in the part of the image over which the blood vessels cast their shadows... When you gently shake the penlight up and down... their shadows are now moving over different receptors, ones to which the brain has not adapted."[6] Though perhaps a somewhat cryptic explanation of how this filling in occurs (we will eventually come back to the mention of "adaptation"), this comment seems at least to state what the brain *has* to do, however it does it. There are gaps in the image cast over the retina; we perceive no gaps in the visual field; so the brain must fill them in for consciousness.

2) **The blind spot**. Most people know they have a so-called "blind spot," an area of the visual field of each eye in which they are functionally blind. The blind spot is due to the convergence of optic nerve fibres from all over the retina at the optic disk, where they must, due to the "backward" construction of the retina mentioned above, penetrate through the layer of receptors behind them to exit through the back of the eye. The very large number of nerves means that, although they are extremely thin, a surprisingly large gap is created: about 6° long and 4.5° wide– considerably larger than the central fovea

[6] Palmer, S. *Vision Science: Photons to Phenomenology* (Cambridge, MA: MIT Press, 1999), p.34 V.S. Ramachandran, among others, praises it highly, calling it "superb" and "encyclopedic in scope."

with which we do all our primary seeing (see (3) below)![7] During normal binocular vision, there is no part of the world before our eyes that completely falls into these gaps, since one eye sees whatever falls into the other's blind spot. What is puzzling is that, during monocular vision (with one eye covered or injured), there is no perception of a gap in the visual field, although the gap can be exposed quite easily, as in Fig. 1, by covering one eye at a time and maneuvering an object into the blind spot, whereupon it disappears from sight.

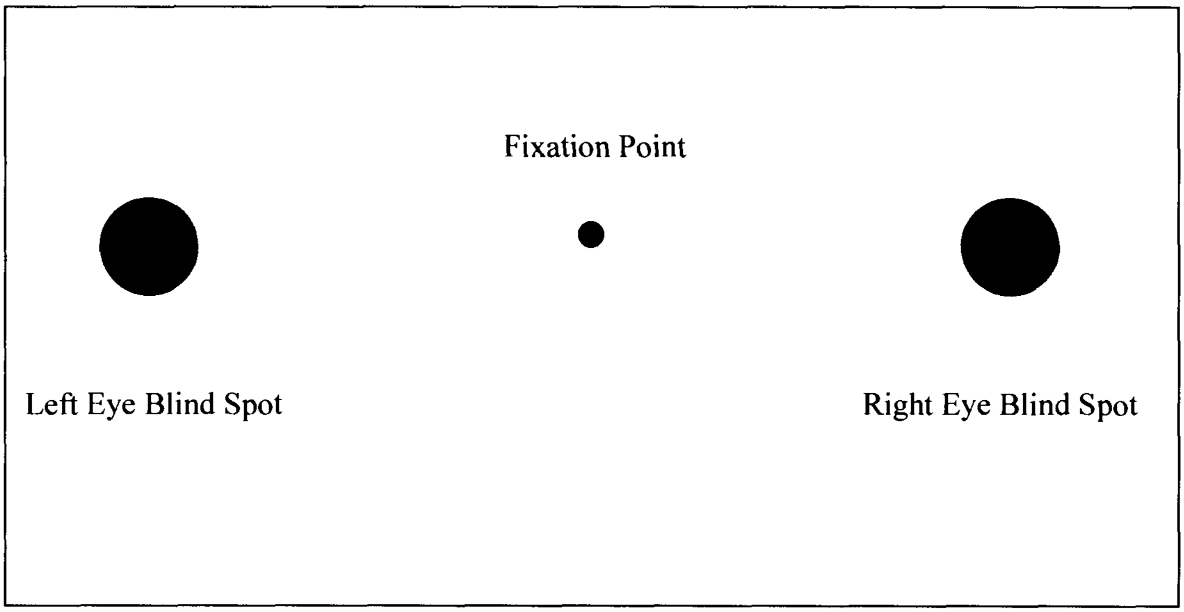

Fig. 1. By closing one at a time, fixating on the central point, and moving the page away from one's face, one can find the blind spot of each eye (the left or right circle will disappear).

It need not be "background" that is filled in, either. The horizontal bar in Fig. 2 can apparently be "completed" just as easily as the circle can be hidden.

[7] Churchland, P.S. and Ramachandran, V.S. "Filling In: Why Dennett is Wrong" in Akins, K. (Ed.). *Perception* (*Vancouver Studies in Cognitive Science, Vol.5*) (Oxford: Oxford University Press, 1996), p.132

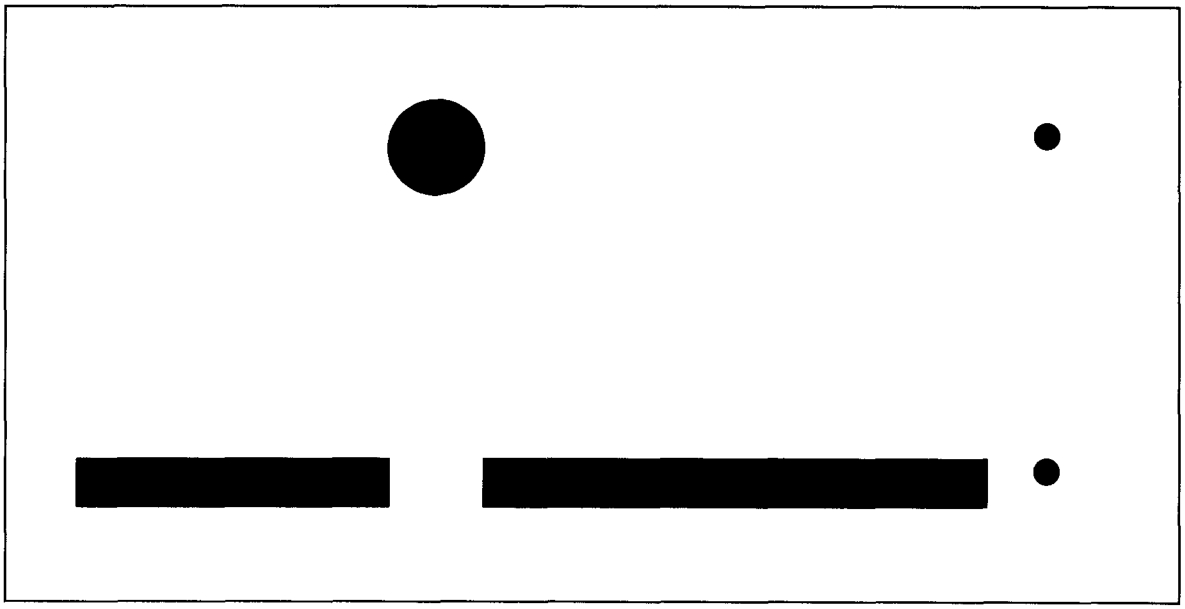

Fig. 2. Not only does the "background" (white) apparently get filled in, but so do "foreground" objects (the black bar).

Here our brains appear to be accomplishing a feat at least as impressive as the filling in of the web of blood vessel shadows across our visual field. And the standard neurological explanation of this feat is at least as cryptic: "Higher brain processes, probably in the visual cortex, seem to fill in the part of the visual field corresponding to the blind spot with appropriate information, which is then experienced consciously."[8] Again, cryptic though this is, how could it be otherwise? There is a large hole in the image projected onto the retina; we perceive no hole in our visual field; so the brain must fill it in for consciousness. The main difference between the optic disk and the blood vessels is that, while the latter can be revealed to us, using extreme angles and intensities of light, *as gaps or shadows* ("in front," as it were, of the world we see), the blind spot is never seen as a hole; it is, at most, a region of our visual field within which what we seem to see out in the world may not be veridical, and at any rate is not *caused by* optical information we collect from that region of the world.

[8] Palmer, p.35

3) **Photoreceptor distribution**. The retina is equipped with some fifty different kinds of specialized neurons[9]; however, vision is made possible by two main types of photoreceptor. It is common knowledge that the retina has rods, which give us high light and motion sensitivity but poor acuity and no colour sensitivity, and cones, which in bright light give us high acuity and all our colour sensitivity. What is less well known is the distribution of these cells across the retina. Fig. 3 plots receptor density as a function of degrees from the center of the retina. There are two obvious features to note about photoreceptor organization. First, at 0° there are no rods, but rod density increases to a fairly high peak at about 20° and tapers off beyond that. This is the basis of our peripheral vision, and is why it is easier to notice a faint light in the dark, or the flickering of a television or computer monitor, by looking past it than by staring straight at, as we are inclined to do. Rods work only in low-level lighting ("scotopic") conditions; they are almost totally unresponsive, and thus provide almost none of our visual acuity, in normal lighting ("photopic") conditions.[10] Second, more alarmingly, cones are found almost exclusively in the center of the retina. In fact, they are only dense enough to allow for fine operations such as studying text within a small cluster, the fovea, covering about 2° of visual arc, about the size of one's thumbnail at arm's length.[11]

---

[9] Marcus, G. *The Birth of the Mind: How a Tiny Number of Genes Creates the Complexities of Human Thought* (New York: Basic Books, 2004), p.71
[10] Palmer, p.29
[11] Churchland and Ramachandran, p.132

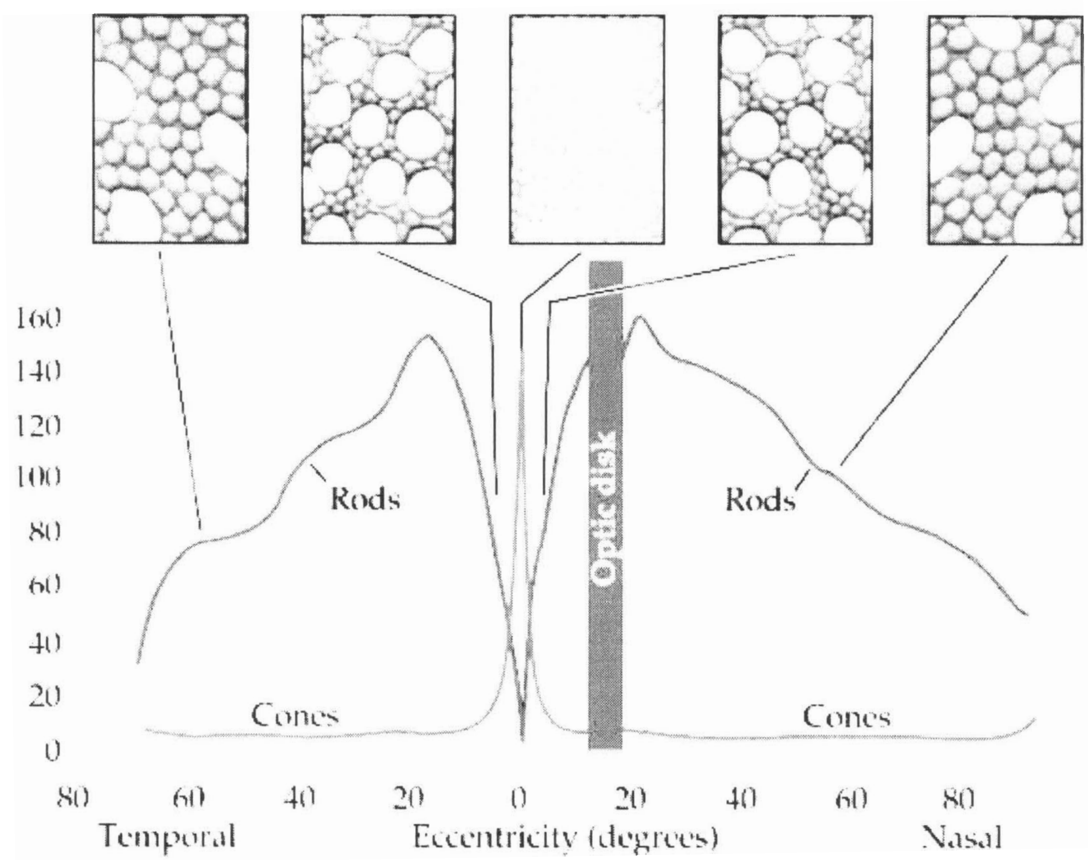

Fig. 3.[12] Receptor density (receptors per square mm x 10^3) as a function of eccentricity (degrees from center of fovea).

This organization yields a graph of visual acuity that seems strikingly at odds with our subjective experience of the world. Fig. 4 plots visual acuity (relative to foveal acuity) as a function of degrees from the center of the retina.

[12] http://www.physiology.wisc.edu/neuro524/vision.htm

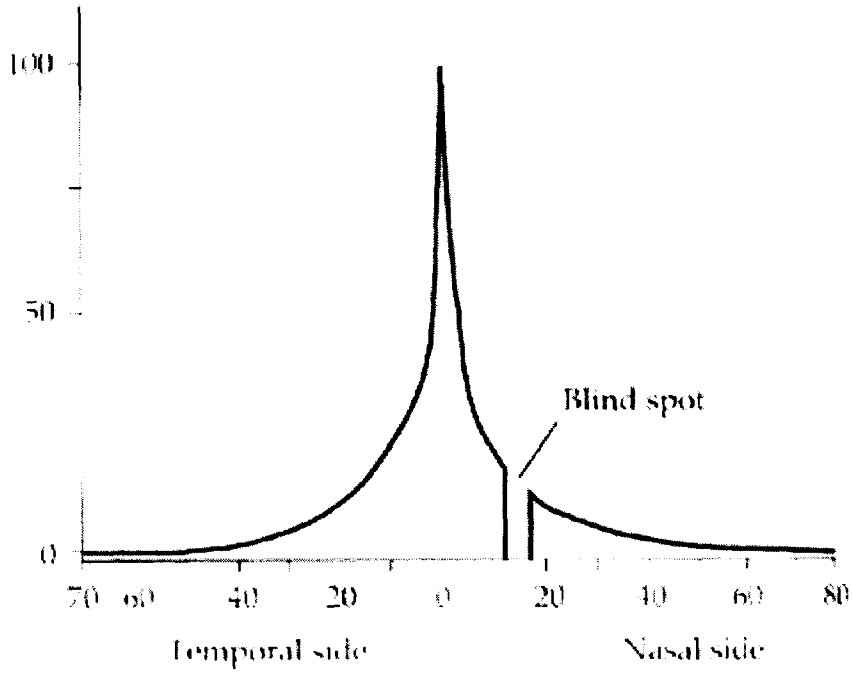

Fig. 4.[13] Visual acuity (percentage of maximum) as a function of eccentricity (degrees from center of fovea).

We might not be able to read fine text if we aren't looking right at it, but acuity subjectively seems to taper off with a much more gradual slope than this, and the impression of experiencing the whole visual field in colour is quite irresistible. Nevertheless, simple demonstrations of these limitations are possible. By controlling eye fixation and slowly moving some detailed or coloured object toward the center of the visual field, one can get a sense of how close to the center it must be before even basic discriminations can be made with confidence.[14] Although the apparent resolution and

[13] http://www.physiology.wisc.edu/neuro524/vision.htm

[14] As we will see, though, this type of demonstration only begins to reveal the lack of detail retrieval in peripheral vision, because one cannot control fixation beyond a certain degree. The eyes continually, involuntarily jump and jitter. The "baseline" movements are muscle tremors called the "physiological nystagmus," but there are also several kinds of more or less consciously controllable eye movements whose function is to acquire new information (Palmer, p.521-26). Experimental techniques that monitor and respond to these movements will yield more provocative results. For a discussion of how cognitive and sensory systems jointly and efficiently control eye movements, see Kowler, E. "What Movements of the Eye Tell us about the Mind," in Lepore, E. and Pylyshyn, Z. (Eds.) *What is Cognitive Science*? (Malden, MA.: Blackwell, 1999), pp.248-62

colour saturation of our peripheral vision is not generally claimed as an instance of filling in, perhaps because it is more difficult to quantify and because vague allusions to "higher brain processes" may sound quite unsatisfactory in this case, it does seem to be another example of the phenomenon: we "see" visual information not collected at the retina.

4) **Shapes, colours and motion**. Before moving on to the battles over filling in, we should note the range of specific types of visual perceptions with which filling in of some sort or another seems to occur. The "rectangle" around the title on page 1, and the "triangle" in Fig. 5, are examples of "amodal completion." The shapes aren't just implied or understood– it seems they are actually *seen*, in some frustratingly hard to describe sense. These sorts of boundaries or contours are said to be "subjective" or "illusory,"[15] although neither word quite seems to convey the effect. One can look right at the blank spaces, see that they are blank, and yet be unable to shake the impression of seeing the shape's boundaries run across them. A similar phenomenon occurs with colour when an illusory shape is represented on a grid by changing portions of the grid from black to red. We "see" the shape, as with the triangle, but this time we see it as a pink region (Fig. 6). This is known as "colour spreading." Like the perception of absent shapes, it is not an "optical illusion" in the narrow sense (i.e. not a light effect). One can look at any of the spaces and see perfectly well that it is white, but also be unable not to see the whole shape as pink.

[15] Amodal completion was studied in depth during the 1970- 80s by Kanizsa. See, for example, Kanizsa, G. "Subjective contours" in *Scientific American*, 234 (1976), pp.23-33

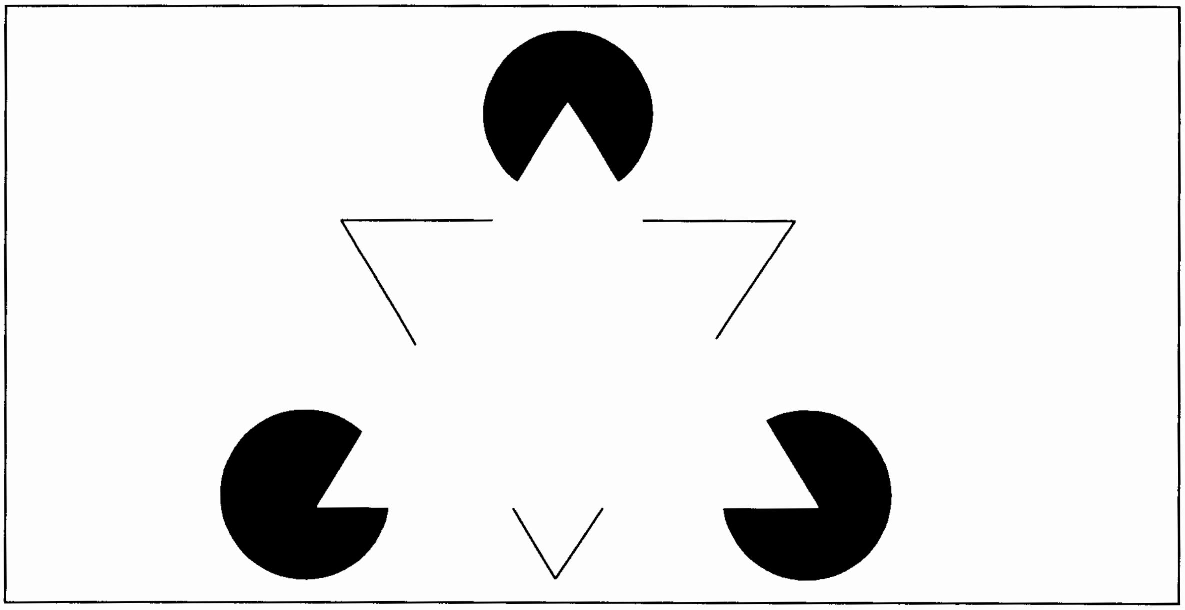

Fig. 5. Amodal completion, sometimes called "subjective" or "illusory" contours. The title graphic is another example.

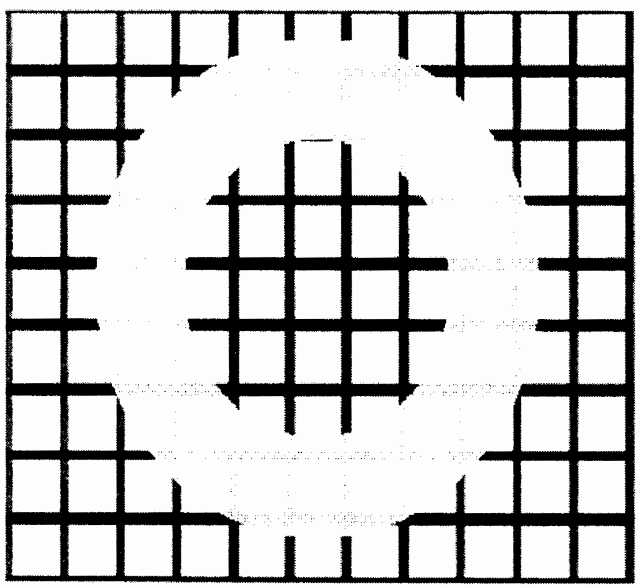

Fig. 6. If the grid lines behind the illusory circle were red, the circle would appear pink.

One final example of filling in before we move on to Ramachandran and Churchland is perhaps, upon reflection, the clearest case of all. Unlike the other examples, its temporal aspect is obvious. The subjective experience of motion can be produced from very detailed and vivid or very abstract stimuli even when nothing actually moves and no image follows a smooth trajectory. Movies are the most familiar example, with simple animation being a starker demonstration. Even completely disconnected and meaningless stimuli, such as a spot of light projected in two locations successively, will be perceived as motion within surprisingly broad spatial and temporal parameters.

Although probably to a lesser extent than the visual acuity example above, some of these phenomena are still difficult to quantify and do not provide the absolutely clear-cut instances of filling in that retinal blood vessels and the optic disk do. Nevertheless, they are commonly (and reasonably) understood along the same lines: we seem to see something that is not part of the image(s) projected onto the retina; therefore, the brain must provide it.

Off the Armchair

Ramachandran and Churchland both get their hands dirty in the lab, so we should take seriously the prediction they offer at the beginning of their discussion of filling in: noting the eye movements called "saccades," which rapidly aim the high-resolution fovea at various places in the visual field, they write that "presumably interpolation across intervals of time to yield an integrated spatio-temporal representation is a major component of what brains do. Interpolation in perception probably enables generation of an internal representation of the world that is useful in the animal's struggle for survival."[16] The internal representation is updated when possible with foveated information and filled in the rest of the time, a process Ramachandran supposes to be necessary in order "to make some aspects of the information explicit for the next level of processing."[17] Their aim is to show that during filling in, the brain is contributing *visual* information to an internal *visual* representation, not merely *judging* that, for instance, the region of the blind spot contains whatever it seems to contain. This, of course, is quite an intuitive thesis– after all, that's *the way it looks*. They defend it in part because the philosopher Daniel Dennett has urged a radically different way of understanding filling in, especially in the case of the blind spot.

According to Dennett, "the fundamental flaw in the idea of 'filling in' is that it suggests that the brain is providing something when in fact the brain is ignoring

[16] Churchland and Ramachandran, p.132. This is not a particularly idiosyncratic interpretation. Many cognitive scientists draw similar conclusions. For example, Bernard Baars explains that "we all have the illusion of seeing more than we do because the brain cleverly takes foveal snapshots of high-information regions in the visual scene and fills in the rest with plausible guesswork" (*In the Theater of Consciousness: The Workspace of the Mind* [Oxford: Oxford University Press, 1997], p.40).
[17] Ramachandran and Blakeslee, p.297

something."[18] In other words, an absence of information, no matter how "conspicuous," will not be perceived *as an absence* unless there is some kind of observer in the brain who is expecting the information, and that is just what there is not in most of the cases at hand. Roughly, all the brain needs to do is make the judgement "more of the same over there," and there is nothing left to explain, either of behaviour or of phenomenology. We will return to this line of criticism in the following section, but for reasons of organizational simplicity, we will now consider the rest of the case for filling in, as presented by Ramachandran and Churchland, with Dennett's claim in mind.

There are three important elements to note in Ramachandran's interpretation of filling in:

1) Officially, he holds no naïve or simplistic views about how the mind works. He would be the first to point out the fundamental lesson of modern neuroscience, that introspectively unified mental events are accomplished by *distributed* neural processing. Filling in is to be no exception. As we have seen, the term may refer to a range of phenomena. "Filling in occurs at several different stages of the visual process, and it's somewhat misleading to lump all of them together in one phrase. Even so, it's clear that the mind, like nature, abhors a vacuum and will apparently supply whatever information is required to complete the scene... [Blind spot filling in] appears to be a manifestation of a very general ability to construct surfaces and bridge gaps that might be otherwise distracting in an image."[19]

---

[18] *Consciousness Explained* (Boston: Back Bay Books/Little, Brown and Company, 1991), p.356. Dennett is not the only commentator with this interpretation of the blind spot. See, for example, Gazzaniga, M. *The Mind's Past* (Berkeley: University of California Press, 1998), pp.134-35
[19] Ramachandran and Blakeslee, pp.88-90

2) Having said this, Ramachandran explicitly banishes the homunculus theory of perception (a little Viewer in the brain who watches the proceedings), and declares himself to hold a "theory-neutral" definition of filling in, which is to say that he uses the term "simply as shorthand to indicate that the person quite literally sees something in a region of visual space from which no light or other information is reaching the eye."[20]

3) Points (1) and (2) make it clear that Ramachandran thinks filling in is to be explained in terms of positive contributions, "mechanisms that actively do something, as opposed to simply ignoring something."[21] He expects specific neural mechanisms will be discovered that are responsible for creating and relaying the visual information required for scene completion. For example, of the bar completion in Fig. 2 he speculates that "perhaps neurons in your visual system are making a statistical estimate; they 'realize' that it is extremely unlikely that two different lines are precisely lined up on either side of the blind spot in this manner simply by chance. So they 'signal' to higher brain centers that this is probably a single continuous line."[22]

Such a relay system is what Ramachandran has in mind when he says that visual filling in is "bottom-up" *perceptual* completion, not "top-down" *conceptual* completion.[23] The latter is equated to intellectual deduction, something that can at least be imagined to be otherwise, whereas the former is "forced upon us," as it were, and Ramachandran describes it in terms of the philosophers' famous "qualia" ("raw feels," the intrinsically subjective "what it is like" of sense perceptions). In other words, we

---

[20] Ramachandran and Blakeslee, p. 273
[21] Churchland and Ramachandran, p.147
[22] Ramachandran and Blakeslee, p.91
[23] Ramachandran and Blakeslee, pp.103, 110

"literally *see*" what isn't there.[24] Fig. 7 is supposed to illustrate this more vividly than the blind spot exercises we have encountered so far.

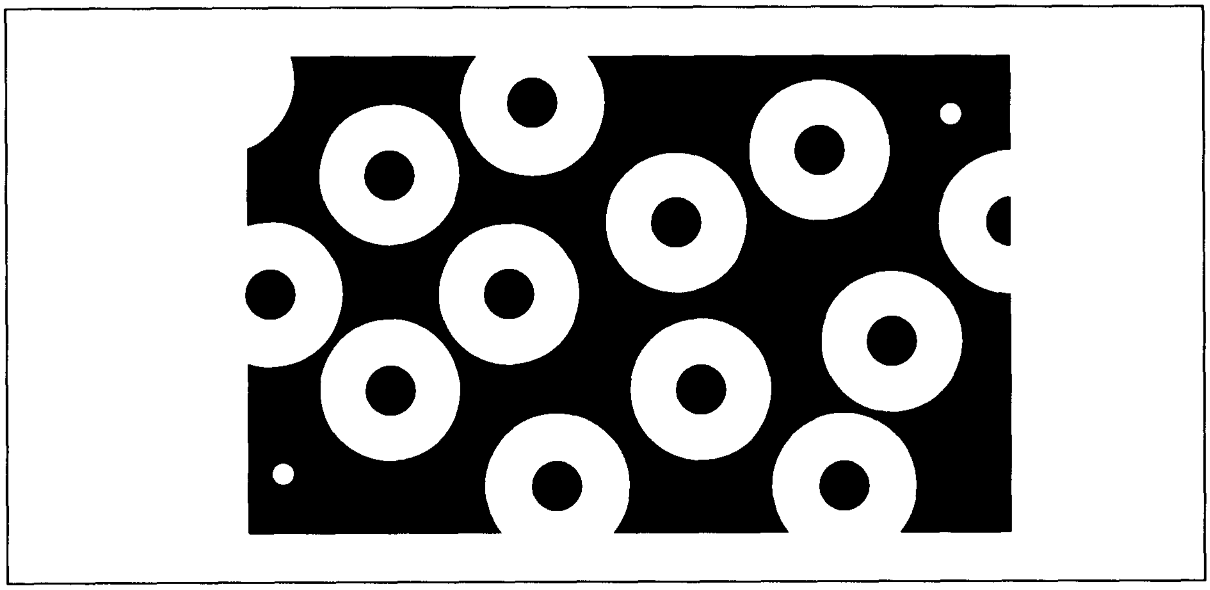

Fig. 7. Using either of the two fixation points (small white dots), one can place the center of a ring in one's blind spot, turning the ring into a disc, which then "pops out" from the others.

When either blind spot is aimed at the center of one of the rings, not only is that ring perceived as a solid disk, it triggers a "pop out" response, so that our attention is drawn to it as an exception, a deviation from the pattern. Ramachandran and Churchland draw two conclusions from this: 1) it refutes any view that the brain fills in only in the sense of making a judgement to the effect that there is just "more of the same" in the blind spot, since more of the same would presumably be another ring[25]; 2) it refutes any view that the brain is ignoring the central region altogether, since only the generation of extra "white" qualia could make the disk *look different* from its "black" qualia-generating

[24] Ramachandran and Blakesless p.236-37
[25] Churchland and Ramachandran, p.138

neighbours[26] (Churchland would probably not endorse "qualia" talk, and this is absent from their joint paper).

The blind spot continues to be a useful case because it seems to permit straightforward investigation of just what the filling in mechanisms can and cannot accomplish. For example, a comparison of Fig. 8 with Fig. 9 demonstrates that while a bar can be extended through the blind spot, it will be so only if it has to be in order to join up with another bar on the other side of the blind spot.

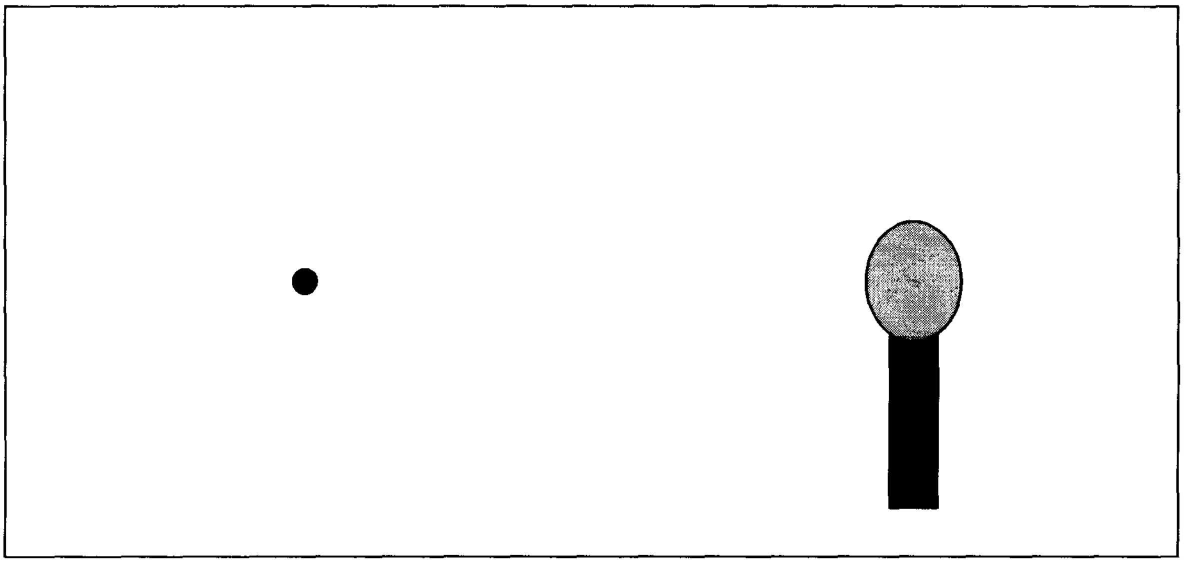

Fig. 8. If the grey oval is placed in the blind spot, the black bar does not seem to become any longer.

[26] Ramachandran and Blakeslee, p.237

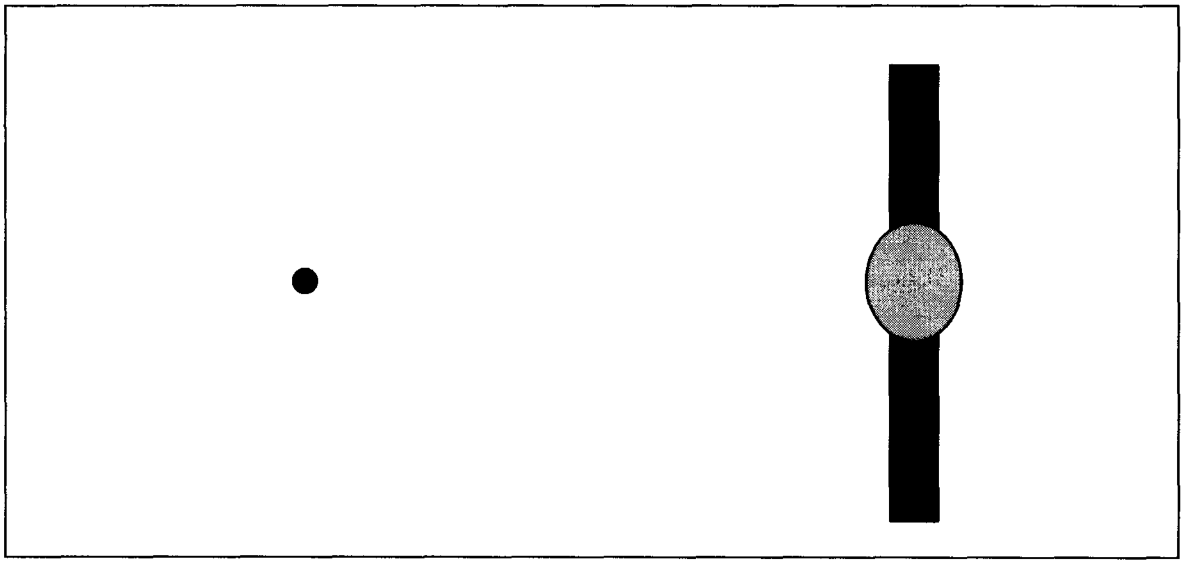

Fig. 9. If the grey oval is placed in the blind spot, the black bar seems to be completed across it.

Figs. 10 and 11 demonstrate an apparent primacy of vertical filling in effects. In Fig. 10, the brain "works harder" to match up the misaligned vertical bars than the misaligned horizontal bars. In Fig. 11, the brain assigns dominance to the vertical strip rather than completing the horizontal bar. (Ramachandran speculates that this may be due to the neural mechanisms of stereoscopic vision.[27])

[27] Ramachandran and Blakeslee, p.94

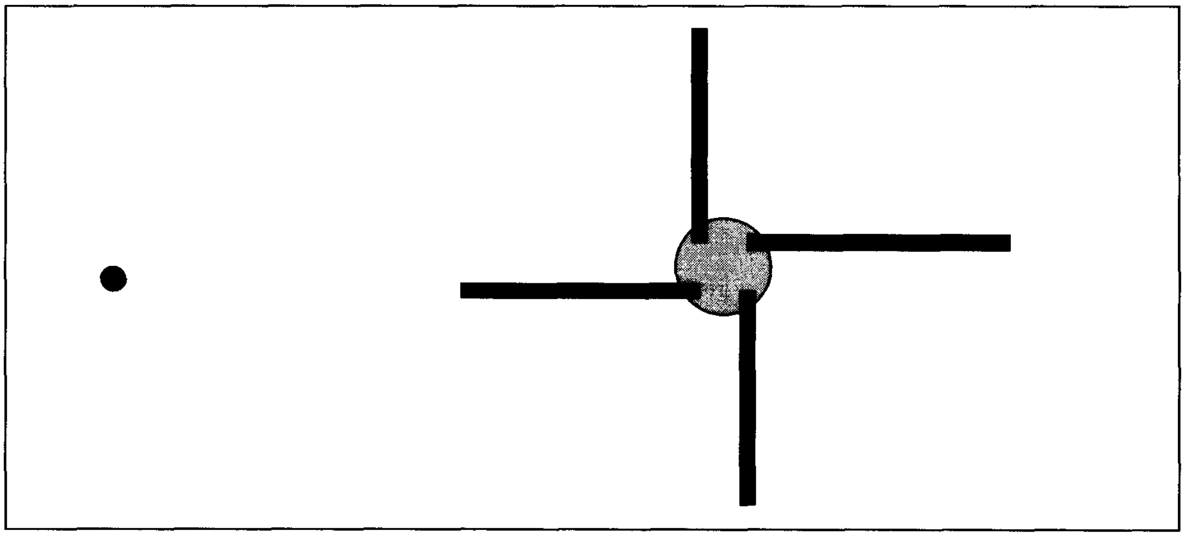

Fig. 10. If the grey circle is placed in the blind spot, the vertical line seems to complete, but the horizontal does not.

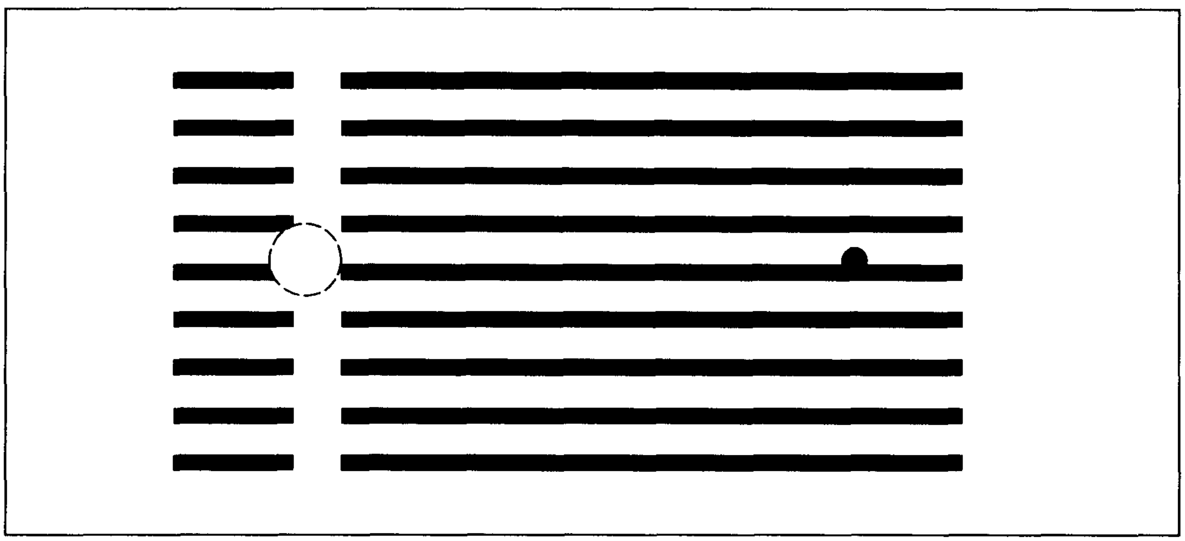

Fig. 11. If the circle is placed in the blind spot, the (illusory) vertical strip seems to take "priority" over the horizontal bar.

Blind spot filling in can be complex, but it is not uniformly so. Ramachandran notes how "clearly the neural machinery that allows completion across the blind spot cannot deal with corners"[28] (Fig. 12), and yet it can accomplish the apparently much more sophisticated task of filling in the spokes of a wagon wheel (Fig. 13).

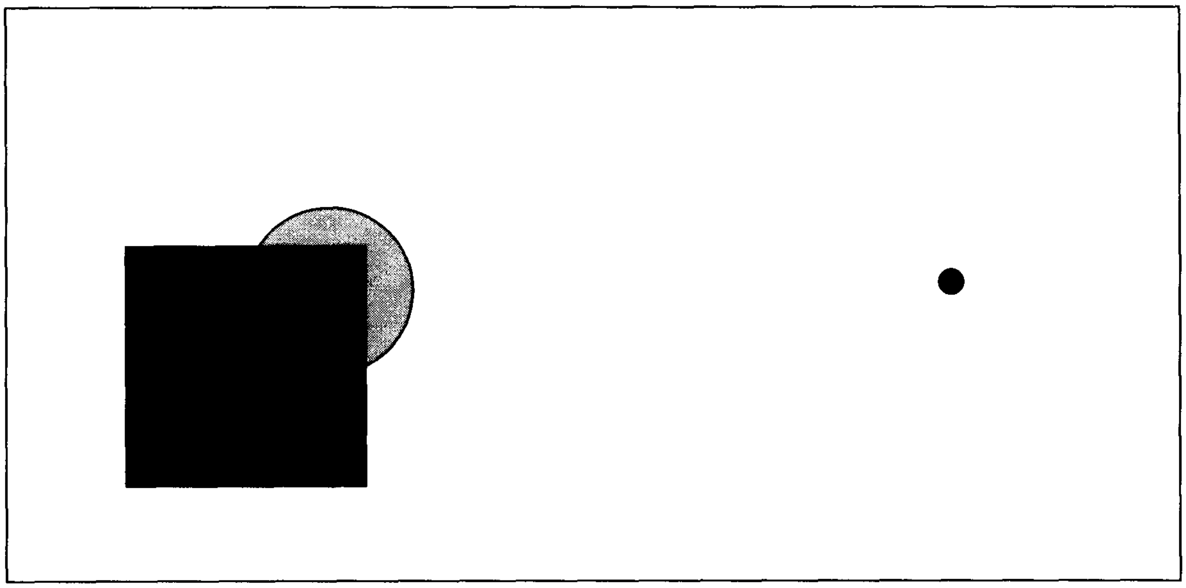

Fig. 12. The "filling in" mechanisms can't quite manage a corner.

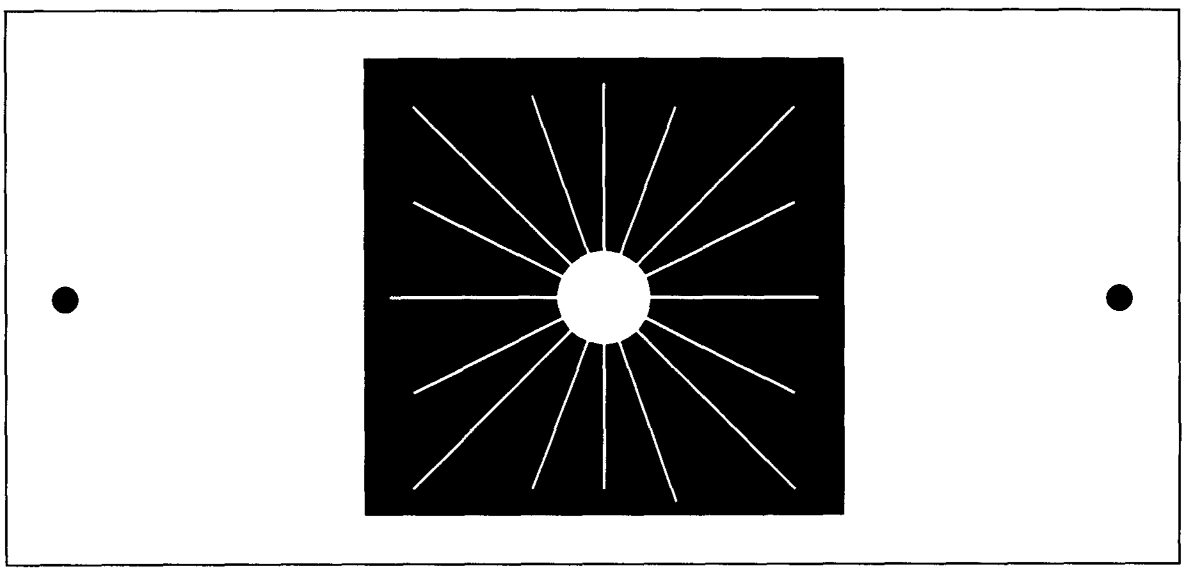

Fig. 13. The "filling in" mechanisms apparently can manage a wagon wheel.

[28] Ramachandran and Blakeslee, p.94

As usual in neuropsychology, some of the most interesting cases are provided by individuals with brain damage (and, to a lesser extent, simulated brain damage, created both in surgical contexts and with non-invasive techniques such as TMS – transcranial magnetic stimulation – which can produce temporary "virtual lesions"). Ramachandran has for some time investigated people with cortical scotomata. These are lesions to the areas of earliest processing in the visual cortex, mainly the area called "V1," the "striate" or primary visual cortex, but also the next layer, "V2." These receive lateralized input directly from each LGN (lateral geniculate nucleus – the main way station for each of the ingoing visual pathways after they have crossed at the optic chiasm) (see Fig. 14).

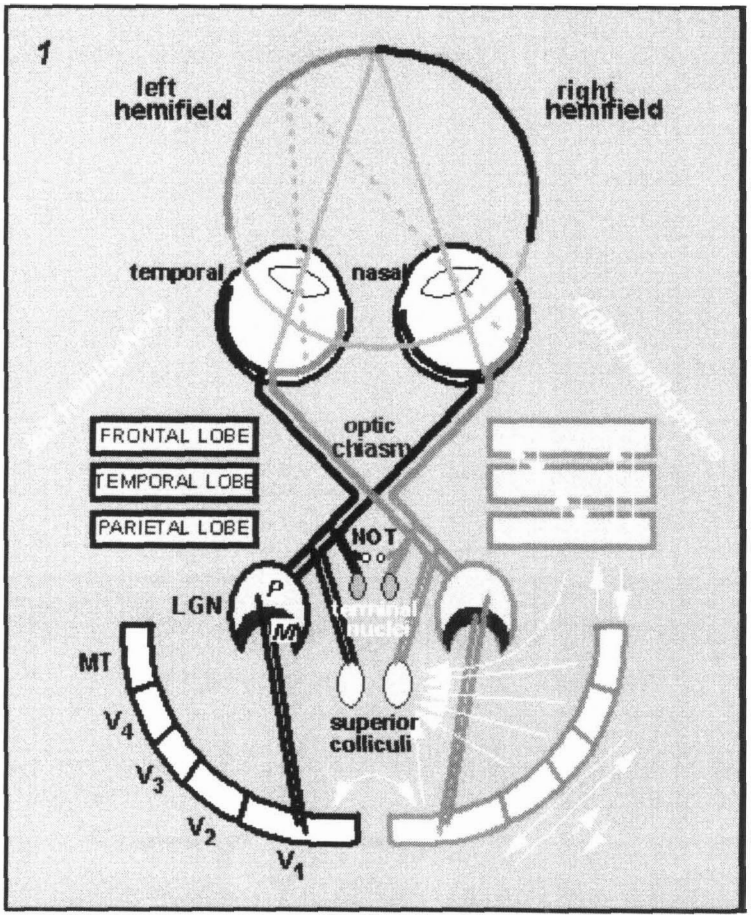

Fig. 14.[29] Schematic of the primate visual system. "Early and intermediate" visual processing is largely organized by "retinotopic mapping."[30] This means that the organization of cells in the visual cortex corresponds spatially – that is, it preserves general spatial *relationships* – to the organization of the photoreceptors in the eyes to which they are linked. Damage to these areas typically results in a scotoma, a functionally blind region of the visual field (for both eyes, since, as the diagram depicts, all information from each eye is hemifield specific and separated accordingly at the optic chiasm, so that all of the information from both eyes about one area of the visual field is transmitted to the same region of the primary visual cortex). Although more "focused" retinotopic mapping occurs at higher levels of visual processing,[31] beyond V2 the interconnections become so complex that single regions of the visual cortex do not correspond in the same way to single regions of the visual field.

[29] http://web.mit.edu/bcs/schillerlab/research/A-Vision/A2-1.html

[30] Farah, M. *The Cognitive Neuroscience of Vision* (Malden, MA: Blackwell, 2000), p.83. Retinotopic mapping is actually just one variant of the more general strategy of "topographic mapping," used throughout the brain as a highly efficient means of neural organization (Marcus, pp.157-63).

[31] Palmer, p.150

A cortical scotoma is in some ways similar to a new blind spot. People do not experience the field defect as a gap or hole, but can easily become aware, by demonstration or accident, that a region of their visual field does not contain veridical information. One interesting difference between a scotoma and the blind spot is that some of the effects observed with blind spot filling in are subject to curious delays when they are filled in through a scotoma. Visual completions often occur over a few seconds rather than instantly, and the effects, once produced, can sometimes persist for a similar length of time. For example, a scotoma aimed at a checkerboard pattern might seem to fill in with checkerboard pattern, but when the displayed pattern is made to flicker, the filled in pattern inside the scotoma takes a few seconds to begin to flicker with it. [32] In another test, a pattern of closely spaced small *X*s filled in across the scotoma, but a pattern of larger *X*s spaced farther apart did not (Fig. 15).

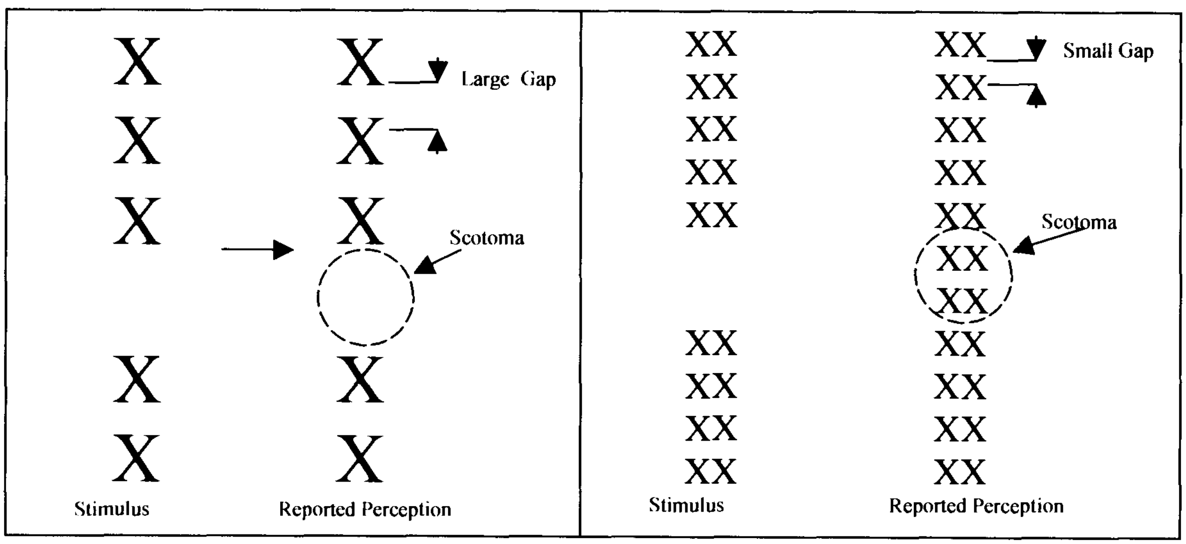

Fig. 15. The column of small, closely spaced *X*s seems to complete, but the column of large, widely spaced *X*s does not.

[32] Churchland and Ramachandran, p.140-43

Ramachandran and Churchland argue that these scotoma effects strengthen the case for fully *visual* filling in, citing both the gradual nature of some of the completions and people's emphatic descriptions of what it *looks* like (or rather, that it looks like something at all, as opposed to people merely *thinking that* something is the case). Citing both the "visual character" of the phenomena and their absence in equivalent tests with "laser-induced *retinal* scotomata," they predict that the mechanisms responsible will be found in the visual cortex.[33]

Ramachandran and Richard Gregory have described a surprising phenomenon that they have called "artificial scotomas."[34] These are effects resembling those produced by blind spots and cortical lesions, but observed with perfectly normal and healthy visual systems. Under sufficiently controlled conditions, subjects will report complex filling in of regions in their peripheral vision. For instance, a small grey patch on a screen displaying static noise ("snow") can be made to fill in after several seconds of the subject fixating elsewhere on the screen, and the effect can even persist for a few seconds after the screen goes blank. As further evidence for the amazing sophistication of some filling in phenomena, Ramachandran reports a similar effect with text (Fig. 16). As with the static, however, this occurs only in low-resolution peripheral vision, where subjects are unable to say exactly *what* text they see (or, for that matter, discern individual flickers of static), but only report *that* they see text (though they still insist that that they *see* it).[35]

[33] Churchland and Ramachandran, p.143
[34] "Perceptual filling in of artificially induced scotomas in human vision" in *Nature* 350 (1991), pp.699-702
[35] Churchland and Ramachandran, p.145

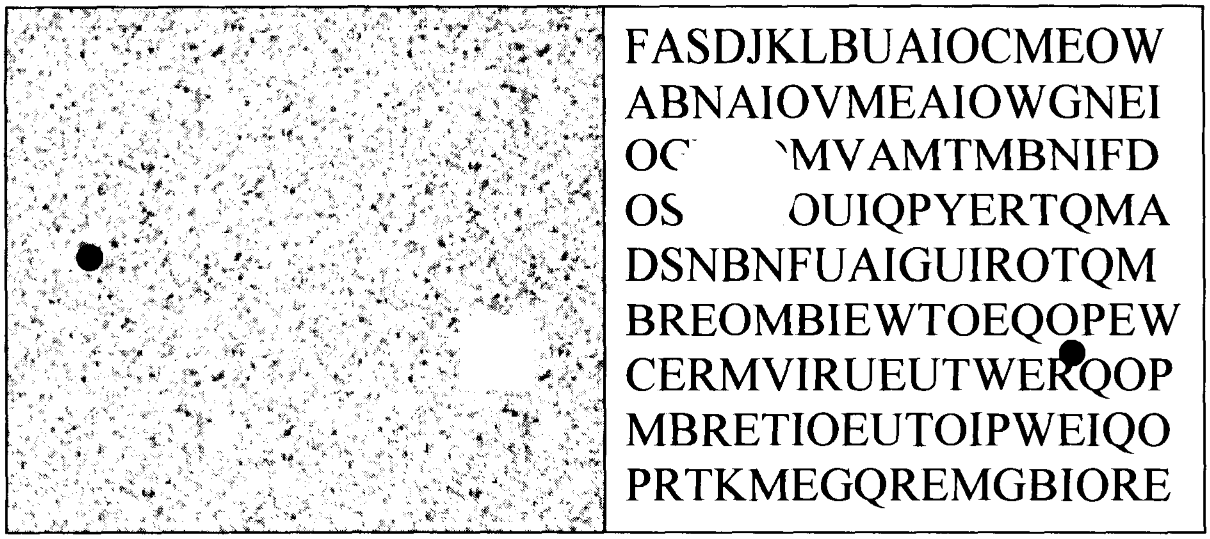

Fig. 16. Under carefully controlled conditions, so-called "artificial scotoma" effects can be created with static or text.

Churchland and Ramachandran go on in their paper to discuss some further experiments and interesting features of cortical physiology, but the effects outlined so far convince them and others that filling in is real. In particular, they are confident hypothesizing that "(1) the brain has mechanisms for interpolation, some of which may operate early in visual processing; (2) brains sometimes visually represent completions, including quite complex completions and (3) such representation probably involves those interpolation mechanisms."[36] They believe, contra Dennett, that the evidence against "the brain ignoring things" interpretation of filling in is compelling, and that there is nothing mystical or Cartesian about "the brain providing things." In the following sections, we will consider a few more criticisms they direct specifically against Dennett, but the general case for filling in has been presented.

[36] Churchland and Ramachandran, p.154

No One to Complain

That filling in does occur seems well established, at least insofar as a range of effects intuitively deserving of the name has been documented and people's sincerity in reporting them is not in doubt. However, the conception of the mind for which things are said to be filled in may be highly misleading, notwithstanding the officially "theory-neutral" account we have from the more careful neuroscientists. Here is the moral our vision textbook draws from the filling in experiments: "Our visual experiences are *derived* from receptor activity, but this activity itself does not appear to be conscious. Awareness appears to arise somewhere farther along a complex chain of neural information processing, rather than in the receptors... What people experience visually is not the pattern of receptor firings in the retina, but activity at some higher level in the brain."[37]

This seems unobjectionable. After all, (almost) no one will want to say that an eyeball sitting on a table is conscious. Once we have taken this step, though, and gone hunting for consciousness by following the neural trail (in effect, saying at each step: "Well, this can't be it– maybe we'll find it a bit further in/up..."), it seems very likely that we will fall into the notorious "explanatory gap" between mere "physical" processes and true "mental" processes. This apparent gap is responsible for the feeling among many people, expert and lay, that when it comes to mental phenomena, we can know everything and yet nothing. For instance, in his book about "the scientific search for the soul" (which is largely about vision science), Francis Crick remarks that "there are two rather surprising aspects of our present knowledge of the visual system. The first is how

[37] Palmer, pp.616-17

much we already know– by any standards the amount is enormous... The other surprising thing is that, in spite of all this work, we really have no clear idea how we see anything."[38]

One might summarize Daniel Dennett's philosophy of mind as the principled attempt to expose as bankrupt the intuition that consciousness is hidden– in particular, that it is some special, unified miracle that happens somewhere in the mists that obscure the peak of a neurological mountain whose slopes are labeled "afferent" and "efferent." Filling in, then, is an obvious target, insofar as the above interpretation of it postulates some privileged place beyond the retina (or LGN, or striate cortex...) where true visual *awareness*, as opposed to mere processing, "arises." The fundamental reason Dennett rejects filling in is that it seems to demand vast amounts of biologically pointless visual processing all done for the viewing pleasure of homunculus *who isn't there*. Thus he thinks talk of filling in is careless at best, but more often a "dead giveaway of vestigial Cartesian materialism,"[39] his derisive term for the alleged implicit assumption by many modern scientists that there must be, if not an immaterial soul mysteriously interacting with the brain, at least some precise time and place of "real consciousness," a privileged

---

[38] *The Astonishing Hypothesis* (New York: Scribner, 1994), pp.23-24. Evolutionary psychologist Steven Pinker is a notable example of a scientist who thinks we have explained, or are in a good position to explain, most specific features and operations of the mind, including vision, but who still regards consciousness *qua* "sentience" as a "riddle wrapped in a mystery inside an enigma," something science has made literally "no progress" at explaining and which may even be permanently "cognitively closed" to the mind of *Homo sapiens* (*How the Mind Works* [New York: Norton, 1997], pp.60, 558-65).

Patricia Churchland, of course, does not endorse this view. As "eliminative materialists," she and Paul Churchland explicitly reject any understanding of consciousness in causal terms, which these expressions of skepticism/frustration implicitly suggest. "Sentience" is not, on their view, some extra thing that has to be *caused* by neural activity. The Churchlands call this the Beatty Crocker theory of consciousness, after a Beatty Crocker Microwave Cookbook whose introduction explains that microwaves heat food by hitting molecules, which increases their kinetic energy... making them rub together with greater friction, which as everyone knows causes the irreducible entity *heat*! (Churchland, P.S., 1994, 2C and Churchland, P.M. *The Engine of Reason, the Seat of the Soul: A Philosophical Journey into the Brain* [Cambridge, MS: MIT Press, 1995, p.207])

[39] Dennett, 1991, p.344

place in the brain that has the perspective and the qualia we attribute to the mind. The following sections will draw on Dennett's arguments, and sketch a generally Dennettian perspective on filling in.

In Figs. 3 and 4 above, we saw how sharply cone density and visual acuity fall off outside the fovea. We uneasily acknowledged this as an apparent case of filling in, but it was much more difficult to say in exactly what sense it is. Unlike the case of the blind spot, where it seems for all the world that we just *see* something in a region for which no information is collected at the retina, all we could say about the resolution of our visual field is that it seems we see more of the world than we do. But this impression is not really an impression *of* anything uniformly resolved (and coloured), as attending to text or other fine details while fixating elsewhere quickly reveals. It is more like a confidence that nothing is really missing from our peripheral vision, even if our discriminatory abilities are demonstrably weaker there. We will consider two general aspects of the wide spacing of photoreceptors in peripheral vision. The first aspect, considered next, will be the inter-receptor spaces themselves; this will lead naturally into a discussion of the blind spot. The second aspect, considered near the end of this paper, will be the sheer amount of detail we have the impression of continuously apprehending.

Ramachandran explicitly denies that the spaces between retinal receptors are filled in for consciousness, although this denial is tucked away in parentheses, in an endnote. Even there, he observes only that "we don't say" that this is what is happens– "there is, after all, no homunculus – that little man in the brain – watching an internal mental screen who would benefit from such filling in."[40] He almost makes it sound absurd, but inter-receptor filling is not a ludicrous notion at all. It is probably the easiest

[40] Ramachandran and Blakeslee, p.273

type of filling in to imagine a mechanism for, since interpolation could be accomplished using only extremely local information (from the nearest receptors) and an unsophisticated procedure such as simple averaging of the two or three nearest colour/intensity stimuli. But in fact, there is no such mechanism, and as he points out, *nor is there any need for one.* Ramachandran states this clearly enough but does not really explain it, which is why its implications for the more "obvious" cases of filling in might be overlooked. The tone and placement of this comment suggest a fairly trivial reminder, but the lesson is profound.

As noted, the retina is "pixilated" with discrete photoreceptors that vary considerably in their physical proximity. And yet we do not see the world as pixilated or fractured. The reason, however, is not that the spaces between the pixels are filled in on some "perfect resolution" image in the mind. It is just the opposite– there is nowhere for these pixels to be displayed *as* pixels, since *the primary visual cortex to which the signals are sent via the LGN is correspondingly "pixilated."* There is no one for whom the spaces need to be filled in; the visual cortex is the only one "looking."

There is a disproportionate allocation, called the "cortical magnification factor," of cortical cellular resources to retinal areas of high receptor density; "the central area of the visual field, which falls on or near the fovea, receives proportionally much greater representation in the cortex than the periphery does."[41] But it isn't just that a bunch more cells are collectively assigned to the central "area"– at the first layers of the visual cortex, there are "automatically" more cells corresponding to the fovea for (roughly) the simple reason that they are individually at the receiving ends of retinal cells. Receptors in the retina, *however far apart they are*, are linked to cells in the visual cortex that are

[41] Palmer, p.37-38

responsible for receiving and processing their signals. And the gaps between the receptors… aren't. That is all there is to say. We don't see the gaps because they extend from the retina all the way down the line– the gaps are *part of us*, not "in front" of us.

There are two important things to note at this point. First, any insistence that, if this is indeed how things work, it just means that there must be someplace *even further* into the brain where information from the primary visual cortex is filled in has nowhere near the same intuitive pull as the claim that information from the retina must be filled in. It is easy enough to regard the retina as "not really *me*," but much harder to regard a good chunk of one's brain, the primary visual cortex, as "not really *me*."

The further into the brain one follows a signal path, the more arbitrary it will be to place the *self* on one side or the other. This is obviously not to deny that things downstream from the primary visual cortex play an important role in a phenomenon usefully called "consciousness." As with the eyeball sitting on a table, presumably just an eye attached to a striate cortex is not going to be conscious in any meaningful sense.[42] But this is different from saying that there is something further down the line that has to *see* some kind of picture that gets passed on to it (and that has, in effect, higher resolution "vision" than the sensory apparatus itself). Interconnections and the division of labour become extremely complex beyond this point, and intuitions about the unity of perceptions can lead us completely astray. For instance, even such basic discriminations such as "what" and "where" are made by totally different subsystems, in this case along

[42] Crick and Koch deny that V1 has a direct role in visual awareness. Their general theory of consciousness is noted for hypothesizing the crucial importance of 40Hz neural activity throughout the cortex, and for linking visual awareness with higher/later layers of the visual system (5 and 6), which are connected, in each hemisphere, to the intralaminar nucleus, a sort of routing station in the phylogenetically ancient brain structure called the thalamus (see, for example "Are we aware of neural activity in the primary visual cortex?" in *Nature*, 375 (1995), pp.121-123).

the ventral (temporal lobe) and dorsal (parietal lobe) pathways respectively.[43] The intuition that consciousness is lodged in some particular "higher level of the brain," far removed from the robotic retina (or robotic motor cortex), which merely collects raw materials (or takes orders), begins to seems less certain. We will have more to say about this.

Second, we are finally in a position to appreciate the force of Dennett's main criticism of blind spot filling in, that the brain is not providing, but rather ignoring something.

> Since your brain has never had to deal with input from this area of your retina [the optic disk], it has devoted no resources to dealing with it. There are no homunculi responsible for receiving reports from the area, so when no reports arrive, there is no one to complain. An absence of information is not the same thing as information about an absence. In order for you to see a hole, something in your brain would have to respond to a contrast: *either* between the inside and outside edge – and your brain has no machinery for doing that at this location – *or* between before and after: now you see the disk, now you don't [as in the standard demonstration of the blind spot given in Fig.1].[44]

---

[43] Palmer, p.38. Actually, although the "what"/"where" view has been standard among neuroscientists, including Ramachandran, Farah, and others, since Ungerleider and Mishkin assigned (all) spatial coding to the dorsal "stream" in the early 1980s, it may not be correct. Goodale has argued against the neural unity of spatial coding since the early 1990s; he argues that the division of labour is really between systems for "allocentric" and "egocentric" frames of reference. The former, now assigned to the ventral stream, is supposed to be of more recent evolutionary origins, since it supports abstractions, mental manipulations, and other higher cognitive functions, while the latter is more important for automatic visual control of action. See, for example, Goodale, M., Milner, A., Jakobson, L., and Carey, D. "A neurological dissociation between perceiving objects and grasping them" in *Nature*, 349 (1991), pp.154-56 and Goodale, M. and Murphy, K. "Space in the Brain: Different Neural Substrates for Allocentric and Egocentric Frames of Reference" in Metzinger, T. (Ed.) *Neural Correlates of Consciousness: Empirical and Conceptual Questions* (Cambridge, MA: MIT Press, 2000), pp.189-202

[44] Dennett, 1991, p.324. This must be qualified by pointing out that there does have to be an *area* of V1 that, in some sense, corresponds to the blind spot. As we saw in Fig. 14, each hemifield is mapped only once, with input from both eyes, and the two blind spots do not overlap. This does not mean, however, that there are cellular resources waiting for input from receptors that would be in the optic disk. More will be said about this.

The idea that the brain could be just ignoring a large hole in the visual field is at first hard to fathom, even if one accepts that the brain is just ignoring the spaces between the retina's photoreceptors. The reason is obvious enough: the optic disk is *so much bigger*. And it's *right there*! How could we possibly miss it unless it were filled in? This response, however, betrays the Cartesian assumptions Dennett is attempting to defeat. It construes the perceptual activity of the visual cortex in essentially *optical* terms.[45] The lesson of the inter-receptor spaces was that, as they get larger, visual *acuity* falls, but that is all. As Ramachandran points out, it is not that we "miss" the spaces between photoreceptors due to their small size, the way we miss the spaces between the coloured dots that comprise an image printed in a magazine. The point is that there is *no one looking where the spaces are*. There is no one who even *might* discern the gaps (say, by squinting hard!), and so no reason to expect to see shadows bobbing around in front of the world. In this respect, the blind spot is just a region of massively reduced acuity. And insofar as the case for filling in supposes that we would otherwise *see* a gap or shadow in the blind spot, it does appear to be an instance of Cartesian materialism.

As mentioned above, Dennett actually says more about the blind spot than this. He adds that "the brain doesn't have to 'fill in' for the blind spot, since the region in which the blind spot falls is already labeled (e.g. 'plaid"... or just 'more of the same')."[46] This is the way Dennett thinks vision works in general. To construct a "bitmap" representation of the world, in which the colour and intensity values at every point are somehow neurologically coded, would demand vast amounts of processing. But this is

[45] We do often conceptualize vision in terms of *looking through* some sort of optical instrument in the face. A vivid, if not entirely serious, example of this is the popular depictions of how the world looks to a fly or other animal with compound eyes (it's like seeing everything through a kaleidoscope– what a headache!).
[46] Dennett, 1991, p.355

totally unnecessary, since there is a far simpler way to accomplish the same job, and no one to complain about the shoddy work. A "colour-by-numbers" system that just assigns labels to whole regions would be a far more plausible design.[47] "An obvious question is: are the [labeled] regions... 'filled in' or not? In one sense, they are, since any procedure that needs to be informed about the color of a region can, by mechanical inspection of that region, extract that information. This is purely informational filling-in."[48]

We will encounter Dennett's ideas of "labeling" again, but for now it must be acknowledged that there does seem to be at least some tension between the "ignoring" and "labeling" claims, and Churchland and Ramachandran also seem slightly unsure of how to construe them.[49] If the brain truly ignores the blind spot region, why would it need even to label it? Dennett connects labeling with the brain's practice of "jumping to conclusions."[50] Perhaps his point is that the brain still needs to *decide* what is out in the world in front of it (in a way that is admittedly different from the way it decides what is out in the world *behind* it), and Dennett's point is that it does not make its decisions, either for the blind spot or for peripheral vision generally, by (unconsciously) constructing an internal visual representation to (consciously) look at, but rather by just making a rough guess once and leaving it at that unless confirmation is required. If Dennett does think that the blind spot is really a "hole" in "visual space" that is

---

[47] Dennett, 1991, pp.347-52

[48] Dennett, D. "Filling in Versus Finding out: A Ubiquitous Confusion in Cognitive Science" in Pick, Van den Broek, Knill (eds.) *Cognition: Conceptual and Methodological Issues* (American Psychological Association, 1992)

[49] Churchland and Ramachandran, p.134. See also Akins, K. and Winger, S. "Ships in the Night: Churchland and Ramachandran on Dennett's Theory of Consciousness" in Akins, K. (Ed.). *Perception* (*Vancouver Studies in Cognitive Science, Vol.5*) (Oxford: Oxford University Press, 1996). In a mildly odd passage that begins as a correction of Churchland and Ramachandran's reading of Dennett, but seems to end more as a correction of Dennett's reading of himself, Akins and Winger observe that "Dennett assumes that the visual system *is* providing something, namely, an inference with an abstract propositional content" (p.181).

[50] e.g. Dennett, 1991, p.355

permanently labeled "more of the same," the following account, which takes the "ignoring" claim seriously, will diverge from his.

Although it is a little difficult to see how best to integrate the "more of the same" theory with the rest of the perspective developed above, Churchland and Ramachandran are wrong to conclude that the ring/disk test (Fig. 7) decisively refutes it. Dennett has not made any specific claims about what rules or parameters determine what "the same" is. As he points out, the brain might label the ring whose center falls in the blind spot as a disk simply because the label-generating mechanism has access only to very local (in this case, within the space occupied by that one ring) and simple information,[51] just as the proposed filling in mechanism would. So the issue for Dennett will eventually become how to understand the "visual character" of a label/judgement, in other words, the so-called "qualia" problem.

The characterization of the blind spot as a region that is "simply ignored" is an empirical claim. Theoretically, it could turn out to be false, if there really were some (neurobiological equivalent of a) homunculus watching the visual cortex as though it were a TV screen. As a theory of perception, any sort of homunculus theory faces both a dearth of empirical support and the notorious problem of an infinite regress, but in fact no one claims or expects to find a "homunculus module," or takes the notion seriously at all. The usual strategy seems to be that taken by Ramachandran and Churchland: to skirt the conceptual problems and instead offer exercises that seem to demonstrate directly that *some* kind of mechanism of visual interpolation must exist, since intuitively only that could account for the impressions produced, namely that one "literally *sees*" the various

[51] "Seeing is Believing – Or Is It?" in Akins, K. (Ed.). *Perception* (*Vancouver Studies in Cognitive Science, Vol.5*) (Oxford: Oxford University Press, 1996), p.171

"completion" demonstrations. More than this, the demonstrations are supposed to take initial steps toward revealing the specific abilities and characteristics of the mechanism(s). It is far from clear, however, that any of these demonstrations succeeds at these goals. Although there is indeed some sense in which one "sees" all the things to which Churchland and Ramachandran direct our attention, in every case, this is *exactly what we should expect* to see if our brains truly ignore the blind spot.

Shape recognition is an enormously complex matter, and researchers do not have a finished theory; however, a basic "atomistic" model is well confirmed:

> Certain neurons [in the primary visual cortex are called] **simple cells** because their responses to complex stimuli can be predicted from their responses to individual spots of light... A simple cell's response to a larger, more complex pattern of stimulation can be roughly predicted by summing its responses to the set of small spots that compose it... The vast majority have an elongated structure, firing most vigorously to a line or an edge at a specific retinal position and orientation... An early step in spatial image processing is to find the lines and edges in the image... Higher-level properties, such as shapes and orientations of objects, might then be constructed by putting together the many local edges and lines that have been identified by their detector cells in V1.[52]

The brain determines a shape based on the input it gets, or does not get, from these simple cells. In Fig. 7, if the center of a ring falls into the region of the blind spot, the brain receives no information about an inside edge (defined by "white qualia"). This information (i.e. positive data plus lack of disconfirming data) is all it takes for the brain to determine (wrongly) that the object is a disk. So the "pop out" effect can be explained in terms of available information and rules for interpreting it.

---

[52] Palmer, p.151. The landmark studies of simple cells were done in the 1960s by David Hubel and Torsten Wiesel, beginning with "Receptive fields of single neurons in the cat's striate cortex" in *Journal of Physiology* (*London*), 148 (1959), pp.574-91

The same lesson applies to Figs. 8 and 9, the bar extension demonstration. In Fig. 9, the brain is presented with two bars with no white edge separating them. All the brain gets is black and more black, and nothing contradicting black. Again, it does not have to decide what to put in a hole– it just has to decide *what the object is*, and on the basis of all the evidence it has, it can conclude only that it is a solid bar.

Figs. 10 and 11 are more interesting because they show the brain dealing with imperfect evidence; so we do, as Ramachandran claims, learn something about the tolerances and/or priorities of the visual system. It seems the brain is indeed more "eager" to find (or less sensitive to discrepancies in) vertical alignment. But what, in Fig. 10, can Ramachandran claim– that the brain somehow reaches into "visual space" and shoves the vertical bars together before filling in the gap between them? The "aligned bar" experience is, to be sure, a visual one, but what else would it be? It is, after all, a judgement of the *visual* system. What there is no reason to conclude, however, is that there is literally, out (or would it be "in"?) in visual space, an aligned, completed bar *being perceived.* In other words, there is no reason to think that the brain perceives a completed "isomorphic" representation of its own making, rather than just determines what it *would* represent if there were someone for whom to represent it.

Fig. 11 is more interesting still, because the brain gets *competing* pieces of evidence (horizontal bar vs. vertical illusory strip). It actually isn't clear how much this demonstrates about the vertical priority of completion effects, since rotating the figure 90° does not seem to produce a completed black bar– it seems harder to say what one sees. Apparently, though, the edge detector cells that respond to the illusory strip (as in the title graphic and Fig. 5), at least in its vertical orientation, get priority in forming the brain's decision. This is an interesting discovery about the visual system and how it prioritizes evidence, but there is no reason to think it is a discovery about a process of interpolation.

Ramachandran marvels at both the sophistication and idiosyncrasies of the filling in mechanism, and thinks it a curious fact that it is clever enough to complete the center of a wagon wheel but not the corner of a square. A second look at Figs. 12 and 13, however, shows that again this is just what we should expect to see if there is simply no one looking at the region of the blind spot. A wagon wheel with no center (not an *occluded* center region– simply no center region *at all*) *already is a completed wagon wheel.* It is just eight bars in different orientations, each one presenting itself to the brain in the same way as in Fig. 9, with no information about a boundary interrupting it. The bars do not need to be extended in visual space for an atomistic shape discrimination system to determine that it sees a wagon wheel.

A square with no corner, on the other hand, is not already a square, or any other shape. The brain has highly indeterminate evidence, and people accordingly report that they just can't describe or decide what they see (except with obviously inadequate terms such as "smeared"– it does not actually look like *a* smeared *thing*). The experience still has a "visual character," but that is about all that can be said of it. This is a very revealing result– not, as Ramachandran suggests, a mere detail about the quirky rules or abilities of filling in. That someone can have a *visual* experience that nevertheless does not *look like* any particular thing is precisely the sort of startling fact, easily dismissed *a priori* as "incoherent," that one needs to ponder before deciding that Dennett's rejection of intrinsically visual qualia is ("obviously," "intuitively," etc.) inadequate.

If these redescriptions of the blind spot demonstrations seem frustrating or obstinate, it must be asked exactly why they seem so. Ramachandran himself indulges in the tendency to think, and speak, of an image that we see as an image in "visual space." We cannot help but imagine visual space as a place corresponding to the nice Euclidean space occupied by objects in the world and presented in 2D blind spot diagrams. Of course, a large chunk of *that* space could not be "ignored," in the radical sense we mean here: simply not being anything. It could be covered over, or perhaps it could be shrunk to a point (severely distorting the space around it), but it could not just not be there. The problem with following this intuition is that there is no such thing or place as visual space. *There is nothing to which our spatial/geometric intuitions have to apply.* Put differently, "visual space" is not a *place* in which we make observations. Its "contents" are not observed; its contents are the observations, or judgements, themselves. In other words, its contents are propositional, not geometric, and there is no reason why some of the many judgements the brain makes about what is in front of its eyes could not be such that they would specify an "impossible" image (e.g. a disk with area less than πr^2, or a bar whose ends line up with the ends of an adjacent bar but whose total length is less than that bar, etc.) *if* they ever had to be rendered onto some kind of display.[53]

[53] We don't "see" the impossible images in the case of the blind spot, but as we've learned, victims of scotomata often do. A friend of the author who suffered a minor scotoma tried to describe it. An elongated, tapered deficit in the upper left quadrant of her visual field made any detailed or attended object in that region disturbingly incomplete. The best way she could describe the experience was in terms of "impossible geometry." It was as though a narrow triangular piece had been cut from a sheet of paper and then the paper's new edges connected. Geometrical intuition tells us that one of three things must occur: 1) the paper becomes a three dimensional shape (a cone); 2) the paper remains flat and so tears at one or many other places, leaving one or many holes; 3) the paper remains flat but stretches (supposing it could stretch) to avoid tearing or lifting up into a cone, and this distorts everything across its surface. There are no other options in real space, yet none of them occurs "in visual space." Beyond this description of what it did *not* look like, words failed her.

With this, though, a perhaps even stronger intuition is engaged: surely filling in could not be a matter of the brain merely making judgements, for how could the brain's "determining that..." *look like* anything at all? How could a judgement, or any combination of judgements, be a true *perception*? Isn't this a sort of "category error"?

The thing to realize about this objection is that it is not an objection specifically against the Dennettian elimination of filling in. It is a general problem of how to understand visual awareness. There are not two different things the visual system can do: judging and seeing. It does not receive little pictures– it receives streams of raw data: nerve impulses with nothing intrinsically "visual" about them. *Everything* the visual system does is judging, and like it or not, all these judgements together *are seeing*. The failure to come to grips with this may be the source of theories of "filling in" understood as "interpolation" or "completion." These interpretations of the filling in demonstrations only make sense on the assumption (or, in theory, discovery) of a process of "rendering," of re-representing back into something intrinsically *visual*. In other words, they assume the brain has to create a passive, picture-like perceptual object in order to see. If it doesn't do this, then "interpolation" fails to pick out any special class of visual phenomena. It will just mean: "all the things the brain does with the evidence it has – sometimes more, sometimes less – to arrive at its decisions," which just describes seeing in general. But the proponents of filling in have a list of distinctly non-standard visual phenomena they mean to specify with the term.

Dennett's lesson applies again: there is *no one to complain* that the judgements have not been rendered back into a properly visual form. If this seems just incredible, it is worth allowing two further points to curb our intuitions.

First, we should be wary of assuming that we know or can imagine what different brain processes would seem like to the brain doing the processing, as neurobiologically informed critics of Dennett admit:

> There is no reason to think that it should seem that your brain jumps to a conclusion. After all, suppose that your brain actually filled in [what you seem to see] through some internal spatial map. It would not seem that way to you either. In general, it does not seem to you, the person, that your brain does anything one way or the other when you perceive; you simply perceive the world (and, on occasion, misperceive it, as in the blind spot demonstrations).[54]

This should not surprise us. A creature for whom perception of the world seemed mediated by conscious brain processes would be unlikely to leave descendents. A perceptual apparatus, from sense organs to the neural mechanisms devoted to processing their input, needs to be "transparent," so that immediate experience is of the *world*, not of the perception.

Second, we should be wary of intuitions about what it is possible to represent perceptually as opposed to propositionally. In Dennett's sarcastic paraphrase, the intuitive contrast runs something like: "after all, perceptions are like pictures, beliefs are like sentences, and a picture's worth a thousand words."

> But these are spurious connotations. *There is no upper bound on the richness of content of a proposition.* So it would be a confusion – a simple but ubiquitous confusion – to suppose that since a perceptual state has such-and-such richness, it cannot be a propositional state, but must be a perceptual state (whatever that might be) *instead.*[55]

[54] Pessoa, L., Thompson, E. Noe, A. "Finding Out About Filling In: A Guide to Perceptual Completion for Visual Science and the Philosophy of Perception" in *Behavioral and Brain Sciences* (Dec 1998) 21(6): pp.728-48
[55] Dennett, 1996, p.162

Richness of representational content cannot by itself reveal to us that some mental activity is a "perception, *not* a judgement." Of course, there are no literal sentences in the head any more than there are literal pictures in the head. The point is that if propositional content can be indefinitely rich, then there is no *function* it cannot perform. We can explain what we do or believe, even about our own minds, without invoking things we are tempted to call "true perceptions" (i.e. *non-propositional* representations). To insist that true perceptions are necessary to explain *what we see* is nothing but deference to raw intuition. And as we have seen, it is just this insistence that motivates the interpolation theory (the "value-added" interpretation of filling in).

We have glimpsed but not yet fully stated Dennett's cardinal rule about the brain, perhaps the organizing principle of his whole philosophy of mind. In Dennett's own, unflaggingly expressive prose:

> Once the brain has made a particular discrimination (for instance, of motion, or of the uniformity of color in an area), the task of interpretation is *done*; no further presentation of the "evidence" on which it is based is required, and hence no further *rendering* of the evidence for the benefit of the Judge. Applying the Thrifty Producer Principle ["If no one is going to look at it, don't waste effort providing it"], we can see that any further representations, with special representational properties, would be gratuitous.[56]

This is the first of two principles we will encounter that articulate a fundamental *economy* of perceptual processes. Technically, of course, that the brain *does* do only what it needs to do is another empirical claim, and subject to empirical refutation. Churchland and Ramachandran take issue with what they see as Dennett's tendency to "confidently

[56] Dennett, 1992

prejudge what the neurobiological data at the cellular level will look like" and accuse Dennett (who is on record as stating that "Biology is Engineering"![57]) of "reasoning more like a computer engineer... than a neurobiologist who realizes how much is still to be learned about the brain."[58] However, they never quite claim that the "Thrifty Producer Principle" as Dennett formulates it is false, and the "neurobiological data at the cellular level" would have to be extremely surprising, given what is known already, to refute it.

The neurobiological data that are documented are single-neuron activation studies in which lateral interactions among cortical cells apparently allow for cellular compensations of certain deficits. In the Gattass Effect, cells in the area of V1 corresponding to the blind spot of one eye (see Fig. 14 and note 44) can become activated, when only that eye is open, as though they had been presented with the stimulus roughly corresponding to what is described as being filled in. In the Gilbert Effect, cortical cells adjacent to an area covered by a retinal scotoma (in monkeys) or artificial scotoma can enlarge their "receptive fields" so they respond to a larger area of the visual field. Churchland and Ramachandran are unsure of what to make of these effects, which they awkwardly call "modification/interpolation[s]" of the cells' responses, and venture only that "it is unlikely the results are irrelevant" to settling the filling debate.[59] The results are interesting and undoubtedly important, but they do not add much to the case for filling in. At best they seem to indicate neural correlates of Dennett's "purely informational filling in." Ramachandran himself has characterized this cellular activity in propositional terms (cells make an estimate and "signal *that*" their estimate is the case).

---

[57] *Darwin's Dangerous Idea.* (New York: Touchstone, 1995), ch.8
[58] Churchland and Ramachandran, p.147
[59] Churchland and Ramachandran, pp.147-53

At some point, some mechanism or other in the visual system has to make a determination about what is out in the world. The case for filling in is not won just by discovering that some such mechanism turns out to be in V1. Unless the determination results in new visual information being added to a visual representation that is isomorphic with the things it represents and therefore in need of having literal holes filled, these are cases of inference, not interpolation. Akins and Winger point out that neuronal response studies "do not address the issue at the right level of abstraction, the level of *neural representational structure and content*," in other words, they do not address the issue of "what makes a particular kind of representation a visual one."[60] Neither the assurances of subjects that their experiences have a "visual character," nor the inference or even discovery that these experiences are associated with activity in the *visual* cortex, can establish the "perceptual, *not* conceptual" thesis. This sort of evidence is just not radical enough to answer Dennett's point. It does not support the intuition behind most descriptions of filling in, that some mechanism makes a determination and on the basis of that determination constructs a unified representation, with special perceptual properties of which we are conscious.

There must, moreover, be a substantial evolutionary presumption against such a mechanism. Evolution is cheap and dirty– it cuts any corner it can and cares only about results. Of course, the results it cares about are always short-term, and sometimes this leads to gross *in*efficiencies or other bad design, such as the ludicrous and lamentable arrangement of the male prostate,[61] or as noted above, the "backwards" design of the

[60] Akins and Winger, 1997, p.183

[61] Many older men learn to their dismay that the urethra passes *through* the prostate rather than around it. Apparently the result of a sheer historical accident – one "misstep" in evolution after which the only small

mammalian retina. These are instances where, in effect, evolution "takes a wrong turn" after which the only statistically available (reachable by random mutation) advantages are always one step further on, rather than "backing up" for a redesign. The brain is an expensive organ to produce and to sustain. Neural tissue is by far the most metabolically demanding tissue in the body.[62] So the brain it is not a luxury– the reason it has evolved is that it does a good job of moving its brain-genes-carrying body around in the world. On the presumably uncontroversial assumption that rudimentary brains were doing a good job of moving bodies around before any of them had consciousness that could benefit from filling in, it is exceedingly difficult to imagine how a filling in mechanism could evolve. The reason is that, historically and phylogenetically, abilities to make sensory discriminations increase in minute advances, in step with abilities to make use of those discriminations in the proverbial "four *F*s" (fleeing, feeding, fighting and reproducing). So at any point in evolution, we would have to expect that the statistically available advances have been small refinements of either some discriminatory ability, or of some way to make fast and reliable use of those discriminations.

Of course, as mentioned above, we would also expect brains to evolve in such a way that perceptual experiences are roughly continuous and uniform; but the easiest way to do this would be, in effect, just to make the brain happy (functional) with the information it has. Interpolation, the neural creation of extra information *on the basis of* existing discriminatory abilities, cannot be any more *useful* (as opposed to "pretty to look at") than the initial discriminations themselves, and would only waste valuable resources

improvements were refinements of the same fundamental design – this causes great sorrow to the many men whose prostates become enlarged at some point later in life.

[62] Comparative estimates range, but one *low* estimate is sixteen times the energy consumption of muscle per unit of weight (Leonard, W. "Food for Thought: Dietary change was a driving force in human evolution" in *Scientific American* ["Special Edition: New Look at Human Evolution," Aug. 2003], p.66).

and processing time. Not only would this be a neural apparatus for which it is very difficult to imagine any selection pressure, but supposing that there is a need for it groundlessly assumes that the mind has some previously set degree and scope of what Dennett calls "epistemic hunger." Epistemic hunger, however, varies as much as the ability to satisfy it. Surely a lesson we learn from contemplating the simpler brain-endowed creatures of this world is that *what it takes*, informationally, for an organism's experience to be unified and useful can vary tremendously. Presumably we don't want to say that the less information captured by an animal's visual system, the *more* filling in it does!

In any case, Ramachandran does endorse an evolutionary logic of the minimum possible processing, but he seems to think it yields a presumption in *favour* of filling in: "interpolation saves an enormous amount of computation."[63] How? By utilizing a trick he suddenly introduces with the term "surface interpolation," in which the colour, texture, etc. of a whole region is represented, apparently, by a single datum, which is all that is necessary to "make explicit" a determination for further processing. The problem with this argument is obvious: he has simply postulated a Dennettian "more of the same" label. If he allows that *grainy yellowness at* X_1-X_2, Y_1-Y_2 is all it takes, then he has granted purely informational filling in. He cannot counter "Yes- but we *see* it," since that is just Dennett's point.

Evolutionary considerations by themselves cannot decide the fact of the matter. The brain does whatever it does. But they are another reason to hesitate before deploying our intuitions against a position such as Dennett's. At any rate, no one claims to have found the kind of telltale "bad engineering" – a sort of "prostate of the brain" – that

[63] Ramachandran and Blakeslee, p.104

would indicate that evolution has accomplished sensory refinements by crafting an ever more sophisticated mechanism for generating visual information for itself to perceive, rather than by (something like) the more efficient means of linking the discriminatory mechanisms directly to those mechanisms that need to use the discriminations.

With this ground covered, we can comment briefly on a few of the other examples of filling in we have seen. As we noted above, illusory contours and colour spreading are common but somewhat ambiguous examples of filling in. With the Dennettian account we have been exploring, we can better understand the ambiguous nature of these effects. Do we see the "triangle" in Fig. 5 or not? It is (or rather would be) stable, well lit, unobscured, etc., so it's not that we are unable to visually discern what is in front of us. It may not *feel* like it, but the problem is that the brain cannot decide whether it is a triangle; in other words, the visual cortex is receiving competing information, none of which is becoming dominant. The edge detectors with which the brain assembles *all* the shapes we see are firing away with multiple vectors that together strongly imply a simple, coherent shape occluding other shapes behind it. And yet the gaps are right there. What should we expect the visual cortex's indecision to *look like*? First of all, we *should* expect it to "look like" something– we should expect an experience with a "visual character" to be "forced upon us." The visual cortex's decisions, even when they are dead wrong, look like things, so why not its indecision? But do we see the triangle or not? The question is misguided, because there is no one in the brain by whom the triangle must be decisively either seen or not seen. Some parts of the brain "believe" there is a triangle and others do not.

We get a similar but somewhat more complex account of colour spreading. Here we must recall again that different sorts of discrimination are accomplished by mechanisms at different places in the brain. Roughly: shapes, including "illusory contours," are detected by mechanisms in V1 and V2, while most decisions about colour are made further along in V4.[64] As we have seen, global retinotopic mapping is not maintained past the early to intermediate stages of processing. One aspect of this organization is that, whereas the receptive fields of neurons in V1 correspond nearly to single points of the visual field, the receptive fields of neurons in V4 cover several degrees of visual angle, and so do not make discriminations about locations as precise as the do the cells in V1.[65]

Normally, when shapes are detected on the basis of real boundaries between colours, the brain has no trouble deciding where the colours start and stop, and labels can be unproblematically attached to the regions perceived. When the boundaries are illusory, though, the precision of the shape discrimination and the imprecision of the colour discrimination together yield judgements such as the surprising "pink ring." In fact, colour spreading does not occur only with illusory shapes. A whole grid of red lines on white yields the same judgement, for the same reason. The line detector cells in V1 can decide almost exactly where the lines of the grid are, but the colour detector cells in V4 are much less precise. With a single line, the imprecision of the colour judgement is not very noticeable, perhaps because one line is about the least ambiguous form the visual system could have to label. However, as the size of the grid spacing shrinks below

[64] Farah (pp.50-53, 32-33) puts the receptive fields of V4 cells at 3°. Measurements of receptive field sizes for cells in different areas vary, but it seems well established that they are smaller in V1 and larger in later areas.

[65] Farah, p.52

the size of the receptive field of V4 cells, the colour identity of the spaces becomes the subject of more heavily competing information (more "red" and less "white"), and the background seems to get increasingly pink. But the "pink background" is not a single perception of a colour "blend"– it is not a coloured *thing* filled in or seen anywhere, by anyone, in the brain. It is the mismatch between the spatial precisions of two kinds of discrimination. Once again, the perceptual experience is the brain's indecision, not the apprehension of any neurologically unified representation of "pinkness."

An adequate discussion of the perception of continuous motion could easily lead into a significant digression. But we can observe that Dennett's account does encompass the subjective experience of motion at a very fundamental level. Churchland and Ramachandran regard the most basic feature of this account as "implausible from a scientific perspective," but offer little comment on why.[66] Dennett argues that "it *does not follow* from the fact that we are equipped to make sequence judgments about events in our experience that there is *any* occurrence in real time of a sequence of neural representations of the events in the order judged."[67] In other words, over very short intervals of time, such as those separating movie frames or flashes of light on a wall, there may be no way to say what comes "before consciousness" and what comes "after consciousness," because "consciousness" is a collection of neural processes spread out in space and (therefore) *time*.

This is highly relevant to "filled in motion" for reasons that are obvious enough upon reflection. In order to fill in a trajectory between two points in time, the brain would already have to know about the second point. So, it is natural to suppose, there

[66] Churchland and Ramachandran, p.155. Dennett discusses some of Churchland's earlier criticisms of his views on "Time and Experience" in 1991, ch.6
[67] Dennett, 1996, p.163

must be some sort of "buffer" where perceptions are briefly stored and edited. But, we might want to ask, which *side of consciousness* is this buffer on? Is the editing accomplished (the filled in trajectory plotted) and then presented to consciousness, or does everything just stream through the "Gate of Awareness" and then get edited as necessary for recall, so that by the time a brain is asking itself what it has seen, the "real memories" have already been adjusted? Dennett's point is that the question itself is a lapse into Cartesian materialism. Beneath the times typical for neural processing, trying to pinpoint the exact time of an experience is meaningless. The brain supports a "stream of consciousness" by continually making judgements about the contents and order of its experiences, but (at the "hardware" level) the discriminations are distributed and there is no one for whom any absolute, continuous sequence must be arranged.

Ramachandran's examples of filling in of scotomata, both cortical and "artificial," are interesting because, in one sense at least, there *is* someone to complain. If someone has a lesion in V1, there are still cells further on "expecting" input from that region, although beyond the retinotopic map of the primary visual cortex the absence of signal is experienced "as a lack, not as a positive area of black."[68] And since the lesion is in the visual cortex, not the eye, it completely destroys one region of the visual field, so the deficit is not overcome by having both eyes open (see Fig.14). For these reasons, it is much easier for scotoma victims than it is for normally sighted persons (with only a regular blind spot for each eye) to make mistakes about what is in the world around them and become aware of their deficit.

---

[68] Dennett, 1991, p.325. See note 53.

As we have seen, in some ways the filling debate is largely about what counts as empirical evidence.[69] The problem with the scotoma effects is simply that neither the reports nor the short delays establish that the effect is "perceptual, not conceptual." The visual cortex, part of which is in the novel position of never getting any disconfirming evidence against which its guesses may be checked, is making do with what it has. The parameters revealed in experiments such as Fig. 15 (the size the *X*s and of the gaps between them) tell us only about the circumstances under which the brain makes one judgement rather than another. The results have a visual character, as we would expect, but that does not mean that something like an image has been produced for the brain to look at. For instance, there is no indication that subjects are able to "count" the *X*s they say are filled in (and so Fig. 15, understood as a literal *copy* of a picture that scotoma patients *see*, may be misleading in this respect).

The "artificial" scotomata are different, since they are caused by the brain's gradually "changing its mind" about the contents of a small region left mostly unchecked out in the informationally sparse peripheral vision, rather than by a cortical lesion, but these also fail to establish the "perceptual, not conceptual" hypothesis. The filled in "text" and "static" are vivid for the subject, but lack all specificity. This looks more like a job for Ramachandran's "surface interpolation." Or, as Dennett puts it, "the illusory *content* is *that there is twinkling in the square*."[70] Neither the richness of the representation nor the conviction that it is *seen* establish anything more. If it seems just incoherent to say that one cannot see a *patch* of static without seeing myriad individual twinkles *of* static, one should consider the filled in text. The details of literally filled in

[69] See Akins and Winger, p.183
[70] Dennett, 1992, p.167

text would be much larger and more discernable than the contents of a similar size patch of static, and yet even the text is *nothing in particular–* it cannot be "read." There is only a "representation to the effect that there [is] no gap in the text, but just more of the same… The effect is, of course, perceptual, but that does not mean it is not conceptual, not propositional."[71]

We have now considered most of the alleged examples of filling in, and seen some ways in which disputes over their correct interpretation become disputes about the nature of mind generally. Some of the examples (e.g. blind spot filling in) have been less ambiguous than others (e.g. illusory shapes and colour spreading). In the section to follow, we will consider what are perhaps the most spectacular filling in effects of all. One is perfectly, remarkably quantifiable, and the other less quantifiable than any case we have been considering here. They are both so extreme, however, that the "interpolation" model of filling in looses its intuitive plausibility, and we are forced to contemplate the perceiving mind in a totally different way.

---

[71] Dennett, 1992, p.169

"Filling in" Over the Top

In 1963, John Krauskopf made a discovery that shows how misleading it is to think of the eye as a sort of biological video camera.[72] As we have observed, although we are not really aware of it, the eyes move continuously in a variety of random and directed ways. This means that the images of everything we see, even with our best attempts to fixate on a single point, are constantly moving over our retinas. Krauskopf investigated the effects of canceling this image movement. Using an ingeniously designed apparatus consisting of a tiny projector mounted on a molded contact lens, he was able to present the eye with a "stabilized image," in particular a stabilized orange circle against a larger unstabilized green circle. Once the circle is stabilized, so that its boundary no longer moves over the retina as the eyes jitter and jump, *all perception of it ceases*. Although subjects stare straight out at a large, bright orange circle covering the whole center of their field of vision, they see only a large green circle. If the discovery of edge and line detector simple cells in the striate cortex were not enough to convince us of the perceptual importance of information about boundaries, Krauskopf's results certainly should.

Before considering interpretations of this phenomenon, we can understand something more about why it occurs. Notwithstanding the eye's superficial similarity to a pinhole camera, it does not "take a picture" of the world. The eye is fundamentally different from a camera because the retina is designed to extract not *images*, but *structure* from the world, and eye movements are an integral part of this design. The retina's function is not to passively record the colour and intensity of light coming from every

[72] Krauskopf, J. "Effect of retinal image stabilization on the appearance of heterochromatic targets" in *Journal of the Optical Society of America*, 53 (1963), pp.741-744

point of an image and then send that off to be "developed" in the visual cortex. A great deal of visual processing is done right at the retina. Photoreceptors in the eye do not record "real" (absolute) brightness, but something called the "delta-brightness," which is the *difference* between the brightness at that receptor and the *average* of the brightnesses at all the immediately surrounding receptors. This strange procedure "provides a portrait of the overall structure of the locally *relative* brightness levels across each of the two input images [one from each eye]... What the fixation cells at the output layer actually detect is *correspondences in structure* between the right-eye and left-eye images, rather than correspondences in absolute brightness levels."[73] Perhaps the most important feature of this system is that delta-brightness detectors are maximally sensitive to "*changes* in the structure-of-brightness-levels... already discovered,"[74] i.e. motion relative to the eye.

The advantages of this design include a high degree of "noise" rejection (in the averaging of brightness values), and the continued effectiveness of stereoscopic vision even with global right/left brightness disparities.[75] The general evolutionary advantage of having a simple, bottom-up, "automatic" system for finding and tracking structure is obvious: structure is a property of *things*, and in the struggle for survival, *things*, and in particular *moving* things, are very important to know about. This is so important that an organism would be ill-advised to leave the basic task of tracking structure to high-level mechanisms receiving huge arrays of raw, unprocessed brightness values. Indeed, even very humble organisms with very few or no high-level mechanisms to speak of have an overwhelming stake in getting the job done quickly and reliably.

---

[73] Churchland, P. M., p.77-78
[74] Churchland, P. M., p.238
[75] Churchland, P. M., p.78

One consequence of using such a reliable procedure for tracking structure across time is that the necessary delta-brightness detectors, always "actively looking for structure," inevitably exhibit a "gradual adaptation to an unchanging scene."[76] We have already seen something of the importance of simple cells, but so-called "complex cells" are actually far more populous in the striate cortex, and these cells are characteristically (edge) *motion sensitive* and *position insensitive*.[77] From the retina on, the system is designed for "optimal processing of motion. By the time the information gets to the visual cortex, very few cells produce sustained responses to unchanging stimulation. Thus, the temporal characteristics of cortical cells in area V1 appear to explain the disappearance of stabilized images." [78] Normally, images are never stabilized, since even the eyes' slightest movements, the physiological nystagmus, are enough to move brightness-structure boundaries over the retina's microscopic photoreceptors. Only techniques such as that employed by Krauskopf stabilize well-defined images[79] enough for adaptation– with one exception, obvious in retrospect. The very first example of filling in we encountered was the retinovascular system, the surprisingly numerous and prominent blood vessels snaking their way across the retina. The shadows of these blood vessels are naturally, permanently stabilized images.[80]

[76] Churchland, P. M., pp.238-39. A striking example of colour adaptation occurs with perception of a "Ganzfeld" – a sort of "stabilized *non*-image" – a completely uniform, featureless field of colour with no shadows, contours, etc. A Ganzfeld can be achieved with special goggles popularly used for meditation (although these more often cycle or flash), or crudely with a coloured ping pong ball cut in half and placed over the eyes with a light source behind them. Over a short period of adaptation, the field colour seems to fade, until subjects call it simply "gray."

[77] Palmer, p.153

[78] Palmer, pp.522-23

[79] A partial adaptation effect can be achieved with normal, deliberate fixation if the inner "circle" is defined by a smooth transition rather than a sharp boundary.

[80] Palmer claims that we fail to perceive not only retinal blood vessels but also the *blind spot* for this reason (p.273-74). However, the blind spot is not really a "stabilized image" at all. As we have seen, it is not a *thing* in the visual field that is immobile *with respect to* our photoreceptors. It is just an *arrangement of* photoreceptors, with no one looking in the gap.

The above account of structure/motion detection seems to give a reasonably clear idea of why we do *not* see stabilized images; but, one may be tempted to ask, what about what we *do* see in their place? Supplying all that "greenness" seems to be quite an impressive feat of filling in. Indeed, it is *so* impressive, it is a wonder that those who construe it as filling in, even people willing to suppose that the brain fills in complex designs such as plaid through the blind spot, do not become a little embarrassed about it. Claims that the brain fills in the colour of a large circle occupying the whole center of the visual field, or that it performs its marvelous interpolations not just in one regularly shaped region, but in a twisting web of gaps (which are different for everybody) across the whole visual field, are really incredible. And awkward, since, following an otherwise exquisitely detailed description of the neurological basis of vision, the "astonishing" Krauskopf Effect is noted in our textbook with the comment that "the physiological mechanisms that fill in sensory experience after adaptation to unmoving contours are not yet known."[81]

The defence of filling in gets even more strained at this point. "Krauskopf's results are just what one would expect if the perceived color and boundaries of regions were *determined* exclusively by edge information [italics added]."[82] If an account of vision in which the visual system does its best to find boundaries, but only *determines* the colour content of a region, rather than inspecting and then representing the whole thing, sounds familiar, it should. It is Dennett's "colour-by-numbers" (labeling) theory.[83] All that is left over is the strong intuition that labels are not *what it looks like*, and the as yet

[81] Palmer, pp.522-23
[82] Palmer, p.273
[83] Churchland and Ramachandran "suspect that neither the bit-map metaphor nor the colour-by-numbers metaphor is even remotely adequate," but say disappointingly little about why. As usual, "much more needs to be known... about the detailed nature of neural computation" (p.153-53).

undiscovered physiological mechanisms that would, in effect, have to *fill in everything we see*.

Churchland and Ramachandran also discuss the Krauskopf Effect, and call it a "remarkable" demonstration of "subjective filling in"[84] (a somewhat odd locution for an eliminative materialist– is Churchland losing her nerve?). They go on to report a "stunning" related phenomenon in which a boundary between two coloured stripes is stabilized, but neither colour region is dominant in the way the larger green circle was in Krauskopf's experiment. Adaptation occurs for the boundary between the stripes, but not for the boundaries defining the whole object against its background. The result is that "the colours begin to fill in across the stabilized border. At this point, some observers describe what they see as a new and unnamable colour that is somehow a mixture of red and green."[85]

"Somehow a mixture," but not a *literal mixture*: not a pigment *blend* we might call "gred." But this is how Churchland and Ramachandran interpret the reports. They seem to accept that people are literally seeing *a* (single) new colour, "hitherto unperceived" somewhere over the rainbow. This is just as groundless as insisting that *the thing* one literally sees at the "smeared corner" in Fig. 12 is a new and hitherto unperceived *shape* (although, lacking two specific competing forms, we would be hasty to call it, say, a "squircle"). These experiences are the brain's *indecision*, not experiences *of* some unified visual representation constructed as a "compromise" between competing stimuli. People may be more likely to say that they can't *describe* what they see than

[84] Churchland and Ramachandran, p.146

[85] Churchland and Ramachandran, pp.146-47. Although not confirmed in their report of the subjects' responses, one imagines that the impression is at least similar to the impression one has when wearing red-blue 3D glasses and looking at a well lit white background ("red and blue all over").

they are to say that they can't *decide* what they see, but this should not surprise us. After all, it's not the whole brain (or the parts of the brain responsible for formulating reports) that can't come to a decision, but only a specialized subsystem within the larger system.

Is the label theory plausible? Leaving aside for now the question of whether the brain could be "satisfied" with a label for any given visual representation, we might wonder whether the brain could really assign and keep track of a whole world of labels. We have seen that unsupported intuitions about what is plausible may be a very poor guide when investigating the brain. However, this is an instance where an intuition ("surely the brain does not continuously keep track of all those labels, even if that would be vastly easier than filling in") may actually be correct. If it is correct, though, it is correct in an even less intuitive way, revealed in experiments even more surprising than the image stabilizations. The label theory is not at all refuted by these experiments, but rather turns out to demand much less, in terms of sheer information processing, from the brain managing the labels than we might think.

The "commonsense" alternative to labeling ("we don't *have labels*– we just *see*") is not really in the offing– it doesn't specify any neurobiologically substantive theory. The proponents of the "interpolation" model of filling in, on the other hand, do have a substantive hypothesis about what the brain does. As they point out, their view entails an "integrated spatio-temporal representation"[86] of the world carried around in the head by each perceiving subject. The function of the sensory organs would be to supply information with which the brain can generate and update this internal representation, a *visual* information structure that is, fundamentally, *how we see.* The whole point of filling in would be to "patch up" the spatial and temporal gaps in this representation. It

[86] Churchland and Ramachandran, p.132

would be useful, then, to determine just how much information is stored in this on-the-fly representation and see whether that corresponds to the intuitions about our visual experience that motivate the theories of filling in.

In the late 1970s, an experimental technique was developed with the aim of studying the function and content of this representation.[87] As we have learned, the eyes are constantly jumping around, primarily, though not exclusively, with semi-consciously directed saccades, whose purpose is to aim the high-resolution fovea at new areas of interest. As John Grimes chronicles it, the idea was that there must be some sort of "integrative buffer... conceived of as a short-term memory store where the visual contents of the previous fixation were maintained until they could be fused with the contents of the current fixation."[88] The experimental apparatus is called an "eye movement-contingent display system," which links an "'eye-tracker,' a machine that monitors the position and movement of the eyes, to a computer controlling a text display."[89] Such a system is useful because it allows a change in the display to be accomplished entirely within the brief saccadic "leap," during which a phenomenon variously understood but usually called "saccadic suppression" renders us functionally blind.[90]

The first idea was to determine how sensitive and detailed the integrative buffer is by alternating, during saccadic suppression, between two texts of identical content but differing form, such as:

---

[87] McConkie, G.W. and Zola, D. "Is visual information integrated across successive fixations in reading?" in *Perception and Psychophysics*, 25 (1979), pp.221-24
[88] Grimes, J. "On the Failure to Detect Changes in Scenes across Saccades," in Akins, K. (Ed.). *Perception* (*Vancouver Studies in Cognitive Science, Vol.5*) (Oxford: Oxford University Press, 1996), p.92
[89] Grimes, p.92
[90] Palmer, p.523

1. CaN nEuRoBiOlOgY tEaCh Us AnYtHiNg AbOuT cOnScIoUsNeSs?
2. cAn NeUrObIoLoGy TeAcH uS aNyThInG aBoUt CoNsCiOuSnEsS?
3. CaN nEuRoBiOlOgY tEaCh Us AnYtHiNg AbOuT cOnScIoUsNeSs?
4. cAn NeUrObIoLoGy TeAcH uS aNyThInG aBoUt CoNsCiOuSnEsS?

The surprising finding was that the alternating text resulted in no disruption of a subject's ability to read it. The truly amazing finding was that the alternating text resulted not even in *awareness* that something odd was happening with the display. The nature of a saccade allowed for even more dramatic demonstrations. Saccades are "ballistic movements," which means that the eyes' trajectory, and hence their next fixation point, cannot be changed once the saccade is initiated.[91] One implication of this is that, with a sufficiently sensitive and well-calibrated eye-tracker, and a fast enough display system, a displayed text can be custom configured for each eye fixation. An entire text can be displayed in fovea-sized (3° at the display distance from the eye) chunks conveniently placed just where they are needed as a person's eyes skip across the screen, such as:

1. Can neurobiology teach xx xxxxxxx xxxxx xxxxxxxxxxxx
2. xxx xxxxxxxxlogy teach us anything xxxxx xxxxxxxxxxxx
3. xxx xxxxxxxxxxxx xxxx xx xxxthing about consciousness?

Subjects' complete ignorance of the fact that they are staring at a whole screen of *X*s with only two or three actual words jumping around to meet up with their fixation points, rather than what it looks like, a page of fixed, meaningful English text, seems to force upon us a highly counter-intuitive conclusion: "people either have little or no access to, or make little use of, the actual peripheral visual information gleaned from prior fixations."[92]

Importantly, the phenomenon is a general one, and has nothing directly to do with lexical processing. On the assumption that "there must be some limit to what a person

[91] Palmer, p.523
[92] Grimes, p.94

can overlook," since "the strength of our visual experience is so powerful and seems so veridical,"[93] further experiments were conducted with rich and detailed scenes rather than text. The expectation was that the threshold for awareness of a disruption during visual inspection of a scene of objects (this being a more primary sort of visual processing task than reading) would be very low. Most people would not suspect, or perhaps even believe, that major changes could be made to a scene they were studying intently without them even knowing *that* a change had occurred, let alone what the change was. Intuitively, that a change occurs during saccadic suppression should make little difference to one's ability to detect it, if the change is at all conspicuous. In other words, if something right before our eyes is one way and then an instant later another way, we expect that a sudden incongruity between the new image and the integrated representation maintained in the brain will alert people to this fact, even if, by an odd and precise arrangement of timing, they technically don't see the change as it occurs. But this is not what happens. For a whole range of prominent scene manipulations, "subjects simply had little awareness of the image changes, and it was difficult to watch their oversights without considerable amusement."[94]

How major are "major changes"? Here are some examples, followed by the detection *failure* rate: "Two men change hats. The hats are of different colours and styles." – 100%; "In a crowd of 30 puffins, 33% of them are removed." – 92%; "In a marketplace, brightly coloured fruits switch places among their four respective baskets" – 75%; "The swimsuit of one of four people posing for a swimsuit advertisement changes from bright pink to bright green. The subjects begin viewing the picture while fixated on

[93] Grimes, p.100
[94] Grimes, p.101

that person." – 58%; "Two cowboys sitting on a bench exchange heads [and thus faces, to which the brain's most sophisticated object recognition mechanisms are devoted]." – 50%.[95] These and the many other examples are all the more astounding in light of two facts: 1) many of the "detections" were described only as perception of a "flicker" of some kind, or even just a vague feeling that something was different; 2) many of those detections are attributable to the inherent delays in the cathode ray tube display system, which must scan line by line down the screen to update the image, and so cannot always complete the new scene before a short saccade comes to an end and the eye re-fixates. Even the true detections, in which the change is noticed and correctly identified, are often described in *non-visual* terms– subjects happen to *note* to themselves that something in the scene is one way, and later note that it is another way. A typical report of a "detection" would be: "You feel like you're hallucinating. It looked like something happened, but you didn't really see anything. Sometimes you just push the button when that happens."[96]

There is nothing in our heads we could accurately envision as a picture of the world, being touched up from saccade to saccade. Our extremely limited range of useful visual acuity doesn't just stop us from reading out of the corner of our eyes – it translates into astounding failures that belie the intuition that things in our peripheral vision are still *there*, even if they're difficult to use. Furthermore, *even a good deal of the information retrieved by the fovea* is *not stored* as part of an on-going, internal representation of the world. The most blatant changes out in the world need cause no disturbance in our heads. If a person can fixate directly on a bright green parrot occupying 25% of the scene they

[95] Grimes, p.102
[96] Grimes, p.103

can see, and then *fail even to suspect that something has happened* when, after a 20 *millisecond* saccade, they are fixating on a bright red parrot,[97] we must be prepared to consider some quite radical revisions to our intuitive notions about the nature and function of visual representations in the brain.

It does seem unlikely that we will have to do away entirely with the notion of an internal representation of the world we see (although some in the AI crowd would go this far[98]). Presumably the scenes displayed could not be changed to just anything; notably, even the most radical manipulations "did not change the meaning of the scene."[99] Of course, different aspects of a scene might be more, less or differently meaningful to different subjects thinking different thoughts about what they see at different instants, and this would explain why some changes were detected by some subjects but not by others.

What it seems that we do have to give up is the intuition that our internal representation of the world, to whatever extent we have one, is a *visual* representation–something, roughly, that "looks like the world looks" (in other words, something that would *need* to be filled in). If such prominent features can be shown to vanish without a trace within the shortest interval of time that is part of the visual system's standard functioning, so that the change can be detected only by non-visual cognitive mechanisms, we must ask what basis we really have for accepting the intuition that we carry in our brains anything like a picture of the world we see. Since we know that visual acuity is simply not good enough to take whole "snapshots" that could be strung together for consciousness, a non-propositional visual representation would have to be maintained

[97] Grimes, p.105
[98] See Pylyshyn, Z. *Seeing and Visualizing: It's Not What You Think* (Cambridge, MA: MIT Press, 2003), p.256. Pylyshyn comments that this extreme anti-representationalism "is demonstrably false even though it contains an important insight."
[99] Grimes, p.108

and updated within some sort of pre-conscious buffer. And yet we also have good reason to doubt that any such representation could last long enough (say, a few hundred milliseconds) for even one or two updates. If it weren't for short-term memory contents of the (propositional) form "I am looking at *X*," it seems that we could miss just about anything.

As Grimes, a defender of the view that some kind of visual representation constitutes "the interface between the environment and the conscious experience of the viewer," is himself compelled to muse: "perhaps the internal representation is based more on information carried by the visual objects, rather than the details themselves (i.e. not based upon the visual details of the image of a dog, but instead upon the observation *that* the image is of a dog)."[100] Without becoming too distracted by Grimes' talk of conceptual discriminations being the *basis* of a (presumably bona fide *perceptual*) representation for "conscious experience," we can see that we have finally, if awkwardly, arrived at the second principle of the fundamental *economy* of perceptual processes: "there is no need to store all of the visual details of the world in some internal memory; the world itself provides for the continuity of the details... Such an economy is sustainable because should the details be needed again for some reason, they are there, available for reacquisition."[101]

This is the conclusion at which Dennett arrives when he considers what seems to be our uniformly detailed perception of uniformly detailed scenes: "no matter how vivid your impression is that you see all that detail, the detail is in the world, not in your

---

[100] Grimes, p.108
[101] Grimes, p.108

head."[102] This principle, however, has been named not by Dennett but by the cognitive scientist Andy Clark. Like the "Thrifty Producer Principle," the "007 Principle" has a broadly evolutionary rationale: "In general, evolved creatures will neither store nor process information in costly ways when they can use the structure of the environment and their operations upon it as a convenient stand-in for the information-processing operations concerned. That is, know only as much as you need to know to get the job done."[103] Clark's understanding of the brain as an "organ of environmentally situated control"[104] and the valuable lesson learned from artificial intelligence, that "the world is its own best representation," lead him to a view of the eye-tracker experiments, and of visual representation generally, that is highly consonant with the Dennettian perspective. "A compelling hypothesis is that the visual system is not even attempting to build a rich, detailed model of the current scene but is instead geared to using frequent saccades to retrieve information *as and when it is needed* for some specific problem-solving purpose… [T]he neural representation of worldly events may be less like a passive structure and more like a recipe for action."[105]

Perhaps the most impressive, and certainly the most vague, case of filling in from our original survey was the apparently substantial contribution of visual detail to our peripheral vision necessitated by the surprising (lack of) distribution of cones across the retina. The experiments with eye-trackers vividly demonstrate the full extent of our "impairment," and thus the true magnitude of the would-be filling in task: the brain would need to manufacture not just details, or even a lot of details, but complete content

[102] Dennett, 1991, p.355
[103] *Being There: Putting Brain, Body, and World Together Again* (Cambridge, MA: MIT Press, 1997), p.46
[104] *Mindware: An Introduction to the Philosophy of Cognitive Science* (New York: Oxford University Press, 2001), p.95
[105] Clark, 2001, pp.91-93

that is fabricated from scratch, and yet bizarrely "inaccessible." There is plainly something wrong with any account on which it would be necessary to suppose that the brain fills in virtually everything we see, and proponents of filling in are not going to endorse such a view unless perhaps they have theological aspirations.

Nor do they have any reason to endorse such a view, since our demonstrable, near-total lack of access to the thing that is supposed to *be* filled in – a rich, *visual* (i.e. "purely *perceptual*") representation of the world – renders the hypothesis that we have such a representation superfluous. In other words, this strange aspect of vision obviates the need for such a mind-bogglingly, even magically, sophisticated system. If the brain can be untroubled by massive ("perceptual") changes in a scene so long as the scene's most basic ("conceptual") meaning is left intact, we have good reason to reject the view that what we strictly, technically *see* ("are conscious of") has perceptual properties anything like those of a picture. And if we reject the theory that we hold in our heads at any instant the kinds of details that a picture has, there is no longer a need for processes of visual interpolation to provide them– the world is its own best representation.

With this, we can see that the 007 Principle is really a corollary of the Thrifty Producer Principle: a visual representation of the world need only allow us to *act now* and find new information as we can use it. There is no one passively gazing at the scenery; rather, there is a brain-in-its-environment supporting a dizzyingly complex project of questions and answers. In other words, the "perceiving self" turns out to be not a Cartesian "I," watching the show behind a veil of ideas, but an active process, spread out in space and time, of the brain's making just those discriminations it needs to make in order to satisfy specific epistemic hungers and to control the organism.

Conclusion: What You Get

Where is filling in supposed to occur? We have seen that a main lesson we are to learn from filling in is that what we are conscious of is "some higher level of neural activity" than the retina.[106] Just how high a level "is conscious" will fix the range within which interpolation must occur. This is a somewhat uncomfortable position for the proponents of filling in. On one hand, we have the repeated claims, about the blind spot and retinal blood vessels, that "higher-level processes fill it in appropriately... some higher-level process fills in these regions as well."[107] If one accepts filling in, this seems plausible, especially in the case of the blind spot, which unlike the blood vessels is not in front of, but actually *part of*, the retina. On the other hand, Ramachandran is quite clear: "filling in... is not a high-level cognitive process."[108] It is a "bottom-up" *perceptual* process that is forced, not based, upon high-level brain activity. These views do not quite contradict each other (how much room is there between "higher" and "high"?), but they do point to the strain any neurobiological theory will feel if it looks too closely for consciousness.

Depending on one's intuitions or knowledge about consciousness and the brain, one is apt to look for a filling in process in different ways: either "it can't be too 'high' up, since if incomplete visual data makes it that far, it will surely be too late– too much information has been *used* by this point for real consciousness to happen any later" or "it *has* to be high up, since there's nothing at the bottom sophisticated enough to actually *do* all that reconstructive work (it's just delta-brightness vectors, primitive property/object individuation, etc.)." Ramachandran seems to reject the latter reasoning, since he equates

---

[106] Palmer, p.616
[107] Palmer, p.616
[108] Ramachandran and Blakeslee, p.96

high-level processes with conceptual or imaginative processes over which we have some control. As we have seen, his expectation is that mechanisms that "operate early in visual processing" interpolate visual information to a visual representation maintained in the brain across time. Whatever else "interpolation" and "filling in" could mean, they have to indicate some kind of process of inserting new information between existing data. A mechanism that, whenever it was "asked," simply interpreted whatever data were available one way or another – Dennett's "purely informational filling in" – will not fit the bill.

The corollary of the "interpolation" hypothesis, that an interpretation could not itself *be* the "filling in," but must instead be the *result* of filling in, is not well supported by the observation that we don't have free imaginative control over these experiences. The intuition that a "mere decision" would somehow have to be rendered into something "perceptual, *not* conceptual" to explain either the fact that we cannot *help* what we see, or fact that we *see* it is just that: an intuition. There is really no reason why we should expect to have imaginative control if Dennett's view were correct, and in light of how little conscious access a brain has to its own workings, we probably should expect no more freedom to imagine that a "filling in" interpretation made somewhere in the visual cortex is wrong than we have freedom to imagine that a foveated detail of a well lit scene is other than it seems to be. And so it is strange that Ramachandran and Churchland want to ask subjects "Does it actually *look* that way to you, or do you just *believe* it looks that way to you?" They take people's reports ("I don't *believe* I just believe it looks that way…") at face value, even though it is perfectly clear that on Dennett's account there would be no way – indeed, no one – to distinguish between the two possibilities. This appears to be just obstinate. It doesn't settle the issue one way or the other unless people are presumed to be infallible regarding the contents of their own minds, and Ramachandran and Churchland are the *last* people who would presume this.

The intuition that seems to be the real root of the interpolation hypothesis is suggested by Ramachandran's endorsement of "qualia," in which he bestows consciousness on "intermediate" processing[109]– far enough away from either sensory or motor processing to seem plausible. This is the line in the sand, as it were. For Dennett, "'what it is like' is a matter of how things are judged to be."[110] In some ways, the filling in debate could never be settled without settling the qualia debate, and it is no accident that in *Consciousness Explained*, Dennett's no-compromise, sustained attack on the notion of qualia immediately follows his comments on filling in.[111] Covering this debate would call for a very long digression, but here we can at least see the terms of the impasse. It may be that we are owed a better story about how to understand the "visual character" of a judgement or label, but insofar as this is an empirical, not a metaphysical, problem, the claim that these things *can* and *do* have this character is at least corroborated by the most basic aspects of brain organization (the lack of a central homunculus for whom any "decoding" or "rendering" of raw data would be required in order for perceptions to be *seen*). These are things about which Ramachandran is officially in full agreement. The point is that, although it is admittedly difficult to imagine "from the

[109] Ramachandran and Blakeslee, p.245
[110] Akins and Winger, p.184
[111] Ch. 12 "Qualia Disqualified." See also "Quining Qualia" in Marcel, A and Bisiach, E. (eds.) *Consciousness in Contemporary Science* (New York: Oxford University Press, 1988), pp.42-77. Churchland and Ramachandran's allegations of Dennett's "sweeping behaviourism and instrumentalism" (p.155) cannot be adequately addressed without extensive reference to these arguments.

inside," and perhaps, as the mantra goes, "more evidence is needed,"[112] a visual character is just the sort of character we should expect visual judgements to have in *and for* a well designed brain with no one inside it to complain about shoddily constructed representations consisting of crude labels and instructions for how to get new information quickly.

(Insofar as this *is* a metaphysical problem, it gets us nowhere– it generates no testable hypothesis about the brain. This is the last direction Churchland and Ramachandran want the debate to go. And as Dennett shows, even the raw intuitions at stake in such a debate – intuitions about the "*real seemings*" of things – become exceedingly tangled if not hopelessly vacuous in carefully constructed thought experiments.)

At times, it is difficult to decide what the debate between Ramachandran (and Churchland) and Dennett is fundamentally about, and some have argued that they talk past each other and really agree on the most substantive philosophical issues.[113] But the charge leveled by Churchland and Ramachandran, that Dennett is just playing "word-

[112] One place more evidence might come from is the study of the unnerving phenomenon of anosognosia, or "unawareness of deficit." The most common form of the condition occurs when a stroke in the right hemisphere of the brain causes both paralysis of the left side of the body and a *complete* (*genuine*) *denial* of this paralysis, no matter what "demonstrations" of the paralysis are arranged. A less common form of anosognosia is Anton's syndrome, in which people rendered cortically blind by a stroke are simply not able to know this about themselves. Ramachandran argues persuasively that "what is damaged in these patients [with anosognosia] is the manner in which the brain deals with a discrepancy in sensory inputs concerning the body image" (Ramachandran and Blakeslee, p.141). While these cases do not appear to be instances of (what would be) "global filling in," since it is not clear that sufferers believe they see anything in particular (although see Metzinger 2003), they do offer a valuable perspective on how we should view our own inability not to believe that our visual experience has either the content or the "character" that it seems to. More generally, what is already known about anosognosia suggests that there is merit in approaches to subjectivity that construe the problem in terms of the brain's model of the body. See for example Ramachandran and Churchland's fellow neurophilosophers Antonio Damasio (*The Feeling of What Happens*: *Body, Emotion and the Making of Consciousness* [New York: Harcourt, 1999]) and Thomas Metzinger (*Being No One*: *The Self-Model Theory of Subjectivity* [Cambridge, MA: MIT Press, 2003]).
[113] See Akins and Winger.

police" and has made, at best, a "semantic" but not an empirical contribution,[114] is not justified. Dennett's point that talk of "filling in… is *not* a safe bit of shorthand, or an innocent bit of temporizing, but a source of deep confusion and error"[115] is born out when scientists adopt what they take to be a theory-neutral definition of the term but are then forced to postulate tremendous computational feats accomplished by unknown mechanisms at an unknown stage in visual processing– and all, in the end, for an unknown audience. Churchland and Ramachandran are obviously not scientifically naïve, but both the term they adopt and the neurobiological features it commits them to – features that are conspicuously undiscovered, functionally and conceptually gratuitous, or demonstrated by experiment not to exist in the relevant way – indicate that Dennett is not being pedantic with his complaints about latent (materialist) Cartesianism.

Dennett never tires of mocking any theory of consciousness that insists on making the experiencing self some kind of "nugget of specialness" ("Why would you think any more of yourself if you turned out to be a sort of mind-pearl in a brain-oyster?"[116]). In a full-blown dualist theory, this insistence is about identifying a precise time and place of *interaction*. In a materialist account, this insistence takes the much more seductive form of gradually pushing consciousness further into the mysterious inner realms of the brain, as more is understood about the outer realms but nothing is found that seems to qualify as genuine *experience*. Dennett, and to a large extent Clark, try to stop this trend before it starts. For them the mind is not cut off from the world and its details, trying to keep up with what's going on out there by viewing the brain's frantically constructed

[114] Churchland and Ramachandran, p.156
[115] Dennett, 1992, p.33
[116] Dennett, 1991, p.367

representations. The structure of the world is part of a perception, something the brain is actively engaged with, not merely having an impression of.

So Dennett's criticisms have substantial merit. His primary claims, however counterintuitive, are well supported by evidence and argument, and they result in a significantly different conception of mind. The introspective evidence for the classic (alleged) case of filling in, the blind spot, can be understood in terms of the brain ignoring things rather than providing them. And many of the other cases of filling in are well accounted for by Dennett's theory of propositional labeling– a theory in mild tension but not contradiction with the "ignoring" claim, since ultimately what the brain "ignores" is the fact that its representations consist of mere labels, and moreover very few of them at any one time. The fundamental economy of the brain's information processing is both an evolutionary presumption and a stark fact revealed in surprising empirical investigations of the visual system's rules and tricks.

Various aspects of our intuitive conception of the mind must be sacrificed to this new perspective: intuitions about the integrity, geometry and isomorphism of visual representations, and even deeper intuitions about the passive nature of representations. Insofar as we have two competing conceptions of mind in the views that 1) the true perceiving is going on somewhere deep in the brain, after all the raw data is put together, and considerably patched up, into an intrinsically *visual* form that is the real object of perception, and 2) all there is or needs to be is a distributed series of judgements, our tour of the conceptual and empirical issues of the filling in debate favours the latter. What we see is what the brain actively inquires about, its own and the world's answers to specific, functional questions.

Bibliography

Akins, K. and Winger, S. (1996) "Ships in the Night: Churchland and Ramachandran on Dennett's Theory of Consciousness" in Akins, K. (Ed.). *Perception* (*Vancouver Studies in Cognitive Science, Vol.5*) (Oxford: Oxford University Press)

Baars, B. (1997) *In the Theater of Consciousness: The Workspace of the Mind* (Oxford: Oxford University Press)

Churchland, P.M. (1995) *The Engine of Reason, the Seat of the Soul: A Philosophical Journey into the Brain* (Cambridge, MS: MIT Press)

Churchland, P. S. (1994) "Can Neurobiology Teach Us Anything about Consciousness?" Presidential Address to the *American Philosophical Association, Pacific Division* (March, 1993) *Proceedings and Adresses of the APA* (1994). Sec.3B

Churchland, P.S. and Ramachandran, V.S. (1996) "Filling In: Why Dennett is Wrong" in Akins, K. (Ed.). *Perception* (*Vancouver Studies in Cognitive Science, Vol.5*) (Oxford: Oxford University Press), pp.132-57

Clark, A. (1997) *Being There: Putting Brain, Body, and World Together Again* (Cambridge, MA: MIT Press)

---------. (2001) *Mindware: An Introduction to the Philosophy of Cognitive Science* (New York: Oxford University Press)

Crick, F. (1994) *The Astonishing Hypothesis* (New York: Scribner)

Crick, F. and Koch, C. (1995) "Are we aware of neural activity in the primary visual cortex?" in *Nature*, 375, pp.121-123

Damasio, A. (1999) *The Feeling of What Happens: Body, Emotion and the Making of Consciousness* (New York: Harcourt)

Dennett, D. (1988) "Quining Qualia" in Marcel, A and Bisiach, E. (eds.) *Consciousness in Contemporary Science* (New York: Oxford University Press), pp.42-77

---------. (1991) *Consciousness Explained* (Boston: Back Bay Books/Little, Brown and Company)

---------. (1992) "Filling in Versus Finding out: A Ubiquitous Confusion in Cognitive Science" in Pick, Van den Broek, Knill (eds.) *Cognition: Conceptual and Methodological Issues* (American Psychological Association)

---------. (1995) *Darwin's Dangerous Idea.* (New York: Touchstone)

---------. (1996) "Seeing is Believing – Or Is It?" in Akins, K. (Ed.). *Perception (Vancouver Studies in Cognitive Science, Vol.5)* (Oxford: Oxford University Press), pp.158-71

Dowling, J.E. (1987) *The Retina: An Approachable Part of the Brain* (Cambridge, MA: Belknap/Harvard University Press)

Farah, M. (2000) *The Cognitive Neuroscience of Vision* (Malden, MA: Blackwell)

Gazzaniga, M. (1998) *The Mind's Past* (Berkeley: University of California Press)

Goodale, M., Milner, A., Jakobson, L., and Carey, D. (1991) "A neurological dissociation between perceiving objects and grasping them" in *Nature*, 349, pp.154-56

Goodale, M. and Murphy, K. (2000) "Space in the Brain: Different Neural Substrates for Allocentric and Egocentric Frames of Reference" in Metzinger, T. (Ed.) *Neural Correlates of Consciousness: Empirical and Conceptual Questions* (Cambridge, MA: MIT Press), pp.189-202

Grimes, J. (1996) "On the Failure to Detect Changes in Scenes across Saccades," in Akins, K. (Ed.). *Perception (Vancouver Studies in Cognitive Science, Vol.5)* (Oxford: Oxford University Press), p.92

Hubel, D. and Wiesel, T. (1959) "Receptive fields of single neurons in the cat's striate cortex" in *Journal of Physiology (London)*, 148, pp.574-91

Hume, D. (1993) *An Enquiry Concerning Human Understanding* (Eric Steinberg, Ed.) 2nd edition (Indianapolis/Cambridge: Hackett Publishing Company)

Kanizsa, G. (1976) "Subjective contours" in *Scientific American*, 234, pp.23-33

Kowler, E. (1999) "What Movements of the Eye Tell us about the Mind," in Lepore, E. and Pylyshyn, Z. (Eds.) *What is Cognitive Science*? (Malden, MA.: Blackwell), pp.248-62

Krauskopf, J. (1963) "Effect of retinal image stabilization on the appearance of heterochromatic targets" in *Journal of the Optical Society of America*, 53, pp.741-744

Leonard, W. (2003) "Food for Thought: Dietary change was a driving force in human evolution" in *Scientific American* ["Special Edition: New Look at Human Evolution," Aug. 2003], pp.62-71

McConkie, G.W. and Zola, D. (1979) "Is visual information integrated across successive fixations in reading?" in *Perception and Psychophysics*, 25, pp.221-24

Marcus, G. (2004) *The Birth of the Mind: How a Tiny Number of Genes Creates the Complexities of Human Thought* (New York: Basic Books)

Metzinger, T. (Ed.) (2000) *Neural Correlates of Consciousness: Empirical and Conceptual Questions* (Cambridge, MA: MIT Press)

Metzinger, T. (2003) *Being No One: The Self-Model Theory of Subjectivity* (Cambridge, MA: MIT Press)

Palmer, S. (1999) *Vision Science: Photons to Phenomenology* (Cambridge, MA: MIT Press)

Pessoa, L., Thompson, E. and Noe, A. (1998) "Finding Out About Filling In: A Guide to Perceptual Completion for Visual Science and the Philosophy of Perception" in *Behavioral and Brain Sciences* (Dec) 21(6): pp.728-48

Pinker, S. (1997) *How the Mind Works* (New York: Norton)

Pylyshyn, Z. (2003) *Seeing and Visualizing: It's Not What You Think* (Cambridge, MA: MIT Press)

Ramachandran, V.S. and Blakeslee, S. (1998) *Phantoms in the Brain* (New York: Quill/HarperCollins)

Ramachandran, V.S. and Gregory, R. (1991) "Perceptual filling in of artificially induced scotomas in human vision" in *Nature* 350, pp.699-702