8

Ships in the Night: Churchland and Ramachandran on Dennett's Theory of Consciousness

Kathleen A. Akins and Steven Winger

I. Dennett's Views on Filling In

In Daniel Dennett's exchange with P.S. Churchland and V.S. Ramachandran, much of the discussion circles around one general issue – the relationship between sub-personal neural representations and the phenomenology that we as persons experience. I sit in my chair enjoying a particular Brahms trio; inside my head, at numerous cortical and subcortical sites, a variety of neural processes occur; some of these processes, we think, are responsible for my experience. But what is the relationship between what I hear and what goes on? What kinds of neural representations – what kinds of neural processes, structures and contents – constitute this type of auditory event?[1]

A special case of this general issue arises for human vision, for the kinds of representations, processes and contents that underlie our ordinary visual experience of the world. When we look at the world, say, through the window to the trees and garden below, the information received through our eyes is transient (e.g., it is interrupted by blinks and saccadic eye movements), serial (we look, first, at one small point in visual space, say at the small pond on the right, then at another point, at the clump of irises beside it) and varies in its informational texture (e.g., from the fovea to the periphery, the eye receives information with diminishing spatial and frequency resolution). But the world we *perceive*, however, the world of which we are conscious, does not bear the mark of its representational origins; even given the fragmented and serial nature of visual information, somehow we come to experience objects and properties as stable across time, as unified, as existing independently of our perceptions and actions, and so on. Focusing on first one part of the garden, then another, we nonetheless perceive the landscape as a whole, existing through time in front

of us; the garden does not come into existence, piece by piece as our eyes traverse the landscape, nor are the objects within it – the pond and the trees – seen as fragmented or disunified. Moreover, when we ask about our visual experience itself, we take our perceptions of the world to be panoramic: at any moment, it *seems* to us that the visual field encompasses a single, large arc of the world in front of us, one that is equally rich in detail across its entire broad field. In other words, we both perceive the world as permanent and stable and experience our visual phenomenology as everywhere rich in information and non-gappy. The particular case of our general question, then, is about the representations and processes that sustain this kind of conscious visual experience.

It is in the context of providing an answer to this question about visual experience that Dennett warns his readers against various forms of filling in. Most generally, he cautions us that, in constructing a theory of vision, we should not accept uncritically either our introspective understanding of what needs to be explained (our views about 'what it is like' to see) or our intuitive assumptions about the kinds of representational structures that must underlie visual experience. More specifically, he argues against three separate kinds of filling in. (What makes Dennett's discussion of filling in so confusing is that an answer to the general question about the relationship between representational vehicles, representational content and conscious experience is merely presupposed, taken as a background assumption. It is never explicitly stated. Here we will begin with a discussion of visual filling in; later, in part III, we will take up Dennett's implicit views on the broader question.)

(1) "Neural Pigment" and "Mental Figment" versus Neural Symbols

First, Dennett dismisses a naïve view about how our visual experiences must come about – the view that properties of the world are represented by like properties in the brain, and that these representations, in turn, give rise to phenomenological experiences with similar characteristics. Thus, a naïve person might imagine that viewing a red apple gives rise to a red coloured representation in the brain (a representation coloured with neural "pigment"), which results in a visual image *of* a red apple, an image which is itself "red" in some metaphorical sense (as Dennett would say, an image that is "coloured" with some mental property, "figment"). Less naïvely, one might imagine that looking at an armchair results in a two-dimensional picture of the armchair at various places in visual cortex, pictures that result in a visual experience of the armchair which is, itself, somehow chair-

shaped or, more vaguely, "spatially extended". As Dennett points out, there is no reason why representational vehicles must themselves exhibit the properties represented – why color needs to be represented by coloured symbols, why the spatial relations of the world must be mimicked by the spatial relations between various symbols in the brain, and so on. Nor is there good reason to think that there are mysterious mental properties, "sort of" colours or "sort of" shapes that exist in a special mental medium. Thus, claims Dennett, we need not posit any visual filling in in the following sense – neural representations need not exhibit the properties of the world they represent nor does visual experience require a mysterious filling in with mental figment in visual experience.

As Dennett well knows, there are very few cognitive scientists who would argue, on general principles, that properties of the world must be represented by like properties in the brain; fewer still who would be keen to champion the existence of peculiar phenomenal properties in an equally strange mental medium. Dennett believes, however, that even though we all know better, the remnants of the naïve view continue to inadvertently shape our theories of visual processing and phenomenology – for example, when we assume that the order of occurrence for external events must be represented by the order of occurrence of neural events, which in turn give rise to a perception of the external events as occurring in a particular objective sequence. This brings us to (2) below, another example of how this intuitive view permeates our theoretical understanding of conscious experience.

(2) Judgments versus Visual Representations: Filling in versus Finding out

Put in telegraphic form, Dennett's second injunction against "filling in" in visual processing is as follows. *Given a "conclusion" by one visual module, in whatever representational form it may arise, there need be no transformation of that "judgment" into a "purely visual mode" of representation in order for us to have a paradigmatically visual experience.*

To understand what lies behind this view, consider something that we will call "Marr's Paradox". This is a paradox that seems to arise if one attempts to understand how conscious vision is related to the neural computational events posited by a typical model of visual processing – say, Marr's theory of vision.[3] (Marr, of course, had nothing to say about the origins of consciousness or about those of visual phenomenology either; nor did he wish to claim that his computational theory of vision constituted anything like a complete explanation of human vision. Without misrepresenting Marr's own claims, however,

we can use Marr's theory to understand the perplexing nature of visual phenomenology.)

On Marr's theory, recall, the principal task of the system is *shape recognition*. The model contains four distinct representational stages. The purpose of each is to draw out or make explicit certain kinds of information contained in the original retinal image. On Marr's scheme, then, the four stages are as follows: (1) an image that represents light intensity at points in the retinal image, the representational primitives of which are intensity values at each point; (2) the Primal Sketch that makes explicit intensity changes and their spatial organization, the primitives of which are zero-crossings, blobs, terminations and discontinuities, edge segments, virtual lines, groups and boundaries; (3) the 2½-D sketch that makes explicit the orientation and depth of the visible surfaces, essentially a line drawing with dotted lines (discontinuities in object surface) and arrows (surface orientation); and (4) the 3-D Model that describes shapes and their spatial organization using hierarchically organized models, stick figures to which volumetric primitives (e.g., generalized cylinders) are attached. Each level of representation, on Marr's view, is to be entirely derivable from processes that act upon the lower level representations (as well as from fully general implicit assumptions). For example, the 2½-D sketch is derived from the Primal Sketch by incorporating the results of a number of processes. Among others, there are the processes that take, as their input, particular descriptions from the Primal Sketch and produce as their output the values for the representational primitives of the 2½-D sketch (relative distance, discontinuities of depth and surfaces, and surface orientation): stereopsis, structure from motion, shape from shading, optical flow and occlusion cues. Thus each new representational level (and the information that is thereby made explicit) is constructed from a series of inferential processes using lower level information.

Here we have the specific case of our general issue: exactly how do these representations and processes meet with our visual phenomenology? When we look at the garden outside the window, we do not see the blobs, boundaries and zero-crossings of the Primal Sketch. Nor do we see little dots and arrows affixed to cartoon outlines, the primitives of the 2½-D sketch; nor the stick figures with their affixed generalized cylinders, the primitives of the shape-recognition system. If any of Marr's representational stages "looks like" our conscious view of the world, it is surely the initial photographic display; cylinders and stick figures are a far cry from what we *see*. In other words, there is potential paradox at hand: starting with the photographic image, the more processing steps taken, each one serving to make explicit *more* information from the retinal image, the less like our visual phenomenology the representational primitives appear to be. So a mystery about visual processing and conscious phenomenology arises: how could the end results of visual processing actually *be* what we know as visual experience? What part in our rich and varied visual lives could the representational primitives of the higher processing levels (the stick figures, cylinders, dots and arrows) actually play?

Marr's theory of vision seems to present us with a paradox only because we start with two conflicting assumptions. First, we assume that it is the upper levels of visual processing – those levels at which the information contained implicitly in the retinal image is rendered explicit – that ought, *prima facie*, to give us some clue about the relation between neural representation and visual phenomenology. After all, if the raison d'être of representational stages, indeed of the processes of vision itself, is to weed out salient information from the retinal image, then surely the *results* of that process – that which is made explicit – support conscious experience. We also seem to assume, on intuitive grounds, that certain kinds of representations are more likely to support particular kinds of phenomenal experience – e.g., that the imagistic aspects of visual experience must arise from iconic representations, that felt judgments ("Aha! now I can make out the giraffe against the trees!") are the product of inferential processes, and so on. (Note that this assumption, while intuitively appealing to even sophisticated researchers, runs uncomfortably close to positing Cartesian figment.) However, given Marr's theory of visual processing, these assumptions cannot both be correct. If imagistic representations are required for imagistic experiences, then the upper representational levels of Marr's theory could not be what underlies visual experience; on the other hand, if it is the upper levels of visual processing that support conscious experience, then the relationship between experience and representational vehicles could not be anything like what our intuitions suggest. Iconic representations are not required for imagistic experience. This is the implicit contradiction that gives rise to our sense of paradox.

Marr's Paradox makes clear, then, that *neither of these assumptions is necessarily true* – both are merely contingent propositions. Moreover, because they are inconsistent, one of them must be false. Perhaps attempting to pair a single type of representational form selected from the latter stages of visual processing (e.g., one instance of a 3-D sketch) with a conscious visual experience (e.g., looking out at the garden) is *not* the way to understand the origins of visual experience. On the other hand, perhaps the relation between representations and experience does not require the similarity of representational and

phenomenal properties that we intuitively assume. Indeed, a more interesting suggestion – one that Dennett makes – is that perhaps *both* assumptions are false: *perhaps many of the representational levels and, indeed, the inferential processes that produce them are all part of the representational basis of visual phenomenology.* Were this the case, what underlies our conscious visual experience would be nothing like what our conservative intuitions suggest.

Not many researchers working in mammalian vision continue to believe that Marr's computational theory is likely to be correct – and those who do, tend to think that Marr's model of shape recognition is but a small part of visual processing. Whether or not Marr's theory is adopted in the end, however, it is not clear that any other of the computational theories of vision currently on offer gives us any clear guidance in addressing the kinds of issues raised by Marr's paradox. To put the point another way, looking at the 26 known visual areas of the brain, whatever computational functions they turn out to perform, it seems plausible that, like Marr's 3-D hierarchical models for shape recognition (of sticks and attached cylinders) much of the representational machinery may not fit neatly into our preconceptions about conscious experience. Whatever the true computational story of vision, it will probably include multiple levels of representation, a variety of representational types and an abundance of processes that mediate between those different representational spaces – all properties that are unlikely to render intuitively transparent the relationship between visual representational vehicles, their content and our visual phenomenology. Indeed, it is Dennett's consideration of multiprocessor models of vision, both computational and physiological ones, that has led him to similar conclusions about representation and phenomenology.

More specifically, Dennett begins with the assumption that there is no Cartesian Theatre – that conscious perception occurs without any "central observer" who "views," from a particular place or vantage point, the results of the multiple visual processes. He then reasons that:

(a) There need be no collation of the individual results of the modular visual processes, *no gathering together of these perceptual conclusions into a single informationally cohesive "package."*[3]
(b) Given (a), and in light of the fact that vision involves inferential and imagistic processes throughout, our conscious visual perceptions of the world are not the products of picture-like neural representations alone nor of any other unified form of representation characteristic of vision; *in this sense, there are no strictly visual perceptions, no perceptions that result from a single form of visual representation.*

(c) Given both (a) and (b), *there need be no re-rendering of the individual results of modular processes into some standard perceptual form*: "once a judgment is made it need not be made again."

Putting these three conclusions together, Dennett denies that there is a second kind of filling in: *given a conclusion by one visual module, in whatever representational form it may arise, there need be no transformation of that judgment into a purely visual mode of representation in order for us to experience it as a visual experience.* There need be no filling in (in standard visual representational form) as opposed to finding-out.

(3) Stockpiled Knowledge versus "Need-to-Know" Access

The world in which we live is, to a large and often predictable extent, stable and unchanging (or so it seems to us, here in central Illinois); occasionally a fly goes by, the leaves on a tree flutter in the wind, or a person approaches. Given this overall stability, Dennett suggests, there is little need for a visual system to process and store detailed representations of the visual world's multitudinous objects and properties. Visual perception will work equally well (indeed, *better*, for reasons of limited storage and ease of access) if the eyes are moved from point to point as particular aspects of the world become of interest. Even then, incoming information may be only partially processed ("nothing of interest there"), other parts may be given a more thorough treatment ("let's just have a better look at that"), while only some small subset of these visual conclusions are likely to be stored for future use. Our representation of the visual world, Dennett speculates, is probably more gappy than our introspective access would ever lead us to suspect.

Still, if our representations of the world are so sparse, why does it seem to us that our visual experience is relatively seamless and extremely complex? Here, Dennett provides two answers. First, the absence of a stimulus must be recorded *as* an absence before it can be experienced as such. For an absence to be experienced or, more relevantly, remedied, there must be a judgment of the form "information of type x is missing." But the visual system, operating as it does on a need-to-know basis, makes no such judgments; hence *no attempt at all – either by inference or otherwise – is made to fill in the missing information of our normal gappy visual representations.*

Second, Dennett suggests that perceptual representations often have what one might call vague or abstract content – propositions of the form 'lots of brightly clothed people crossing the bridge' or 'a circle filled evenly with twinkles of the very same type' or simply 'a line that changes colour somewhere in between the red end and green end'.

Just as the Belloti painting does not represent specifically the finery of the people crossing the bridge, the visual representations themselves lack information about specific aspects of the world; but the representations *depict* the world *as* having complex properties, just as the Belloti painting gives the impression of many individuals, sumptuously attired. In other words, if the content of the representation is suitably abstract, sparsely detailed representational content need not be manifested, phenomenologically, as a sparsity in the perceived properties of the world. *Hence, again, there is no felt loss of information nor any need to fill in or find out the missing details given the abstract content of the representation.* (If something seems amiss here, worry not. We will come back to this point later.)

Hence, Dennett's third moral about filling in: *even though, in normal visual processing, the visual content is both gappy (piecemeal) and abstract, there is no reason to suppose that this information must be supplemented in order to produce our rich visual phenomenology.*

II. Gliding on Past

All three denials of representational filling in are offered by Dennett as general principles of conscious visual processing. Thus his view on the blind spot is offered merely as one speculative case, one instance in which these principles might be borne out. (Note that Dennett's explanation of the blindspot need not apply to the phenomenology of artificial scotomas. Whatever strategies the visual system has developed for *normal functioning* with the blindspot need not be the same kinds of strategies that are used in cases of breakdown – here, in the case of artificial scotomas, when suddenly the system fails to receive information of the usual kind.) What, then, according to Dennett, happens when the visual space occupied by the blind spot fills in? Dennett postulates that, at some level of visual function, the system infers, on the basis of information from surrounding areas in the neural map, what the blind spot ought to have contained. The exact content of such inferences depends upon the stage or area of representation at which the inference is made – a content consonant with the information processing tasks of the particular neural site. If the inference is made further up in visual processing, then, the judgment will necessarily have a content that is propositionally abstract (e.g., "small twinkles exactly like the surrounding ones") *for that is all the inference from the surrounding areas warrants, given the more abstract content of the representational primitives (or type of "content determinings," in Dennett's terminology).* And once such inferences are made, they will be neither confirmed nor disconfirmed directly, for there are no neurons that are capable of

providing the requisite consistency check. As this is the normal case, the absence of confirmation is not judged *as* an absence. (Here it is clear why Dennett's explanation of the blind spot cannot be used to explain the phenomenology created by an artificial scotoma. With an artificial scotoma, there *are* neurons, below the level of pattern recognition, that are responsible for monitoring the blank area of visual space. Some story must be told then about why, when only noise is received from such lower level neurons – for visual neurons will continue to fire randomly in the absence of stimulation – the inference is not judged false.) The spot is labelled "just more of the same twinkling" and in virtue of that label alone, we experience the blind spot as filled with twinkles. There is no reason why the judgment must be reinterpreted into the correct form of visual representation (for there is none) nor painted into psychological space with some kind of mental figment (for there is none of that either), before the experience can occur. Dennett says, by way of summary,

> The brain doesn't have to "fill in" for the blind spot, *since the region in which the blind spot falls is already labeled* (italics ours) (e.g. "plaid" or "Marilyns" or "just more of the same"). If the brain received contradictory evidence from some region, it would abandon or adjust its generalization, but not getting evidence from the blind spot region is not the same as getting contradictory evidence. The absence of confirming evidence from the blind spot region is no problem for the brain; since the brain has no precedent of getting information from that gap in the retina, it has not developed epistemically hungry agencies demanding to be fed. Among all the homunculi of vision, not a single one has the role of co-ordinating information from that region of the eye, so when no information arrives from those sources, no one complains. The area is simply neglected. (Dennett 1991, 355)

The various ways in which Churchland and Ramachandran pass by Dennett's position should now be clear. First, Churchland and Ramachandran misconstrue Dennett's specific views about how the blind spot is handled during normal visual processing. On Churchland and Ramachandran's interpretation, Dennett simply denies that the brain provides or infers any information about the contents of the blind spot. For example, they cite him as saying, as evidence of this view: "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 something" (Dennett 1991, 356). On the contrary, Dennett assumes that the visual system *is* providing something, namely, an inference with an abstract propositional content; what it ignores is the lack of any direct confirmation or refutation of that inference.

Admittedly, it would greatly surprise Dennett if the phenomenon of filling in for the blind spot turned out to be explainable solely in terms of very low-level visual processes – if, say, the visual map in V1 contained individual "healthy" neurons that represented the area of visual space covered by the blind spot, and that thereafter absolutely no further special "finagling" (judgments or inferences) occurred.[4] After all, the experiments make clear that the principles that govern filling in utilize higher level properties of the visual scene – continuity, shape, relational properties between shapes, and so on – in order to determine the contents of the blind spot. So if filling in occurs at the level of V1 alone, somehow the patterns of neural activity created must be such that they can sustain the orderly computational processes of pattern recognition that occur further down stream, at the higher levels of processing. How is it, one might well wonder, that the pattern of activity qua pattern is recreated over the entire cortical area of the blind spot if information from above (from pattern recognizers) is not used? Note that, by Dennett's lights, such filling in cannot come about by means of the feedback of information from the higher processing levels (i.e., through the information discerned by the higher levels about the patterns of neural activity surrounding the blind spot), for such "back-filling" would violate Dennett's rule that "once a judgment is made, it need not be made again." If the patterns of the surrounding areas have already been discerned and the inference that the blind spot contains pattern x has already been drawn, why would it be necessary to re-write this inference into the representational space of V1? It is simply concluded that pattern x occurs. It seems natural for Dennett to suppose, then, that inferences about the blind spot occur at higher processing levels, those at which the various complex properties of the retinal image are discerned, and those at which the content of the representational primitives is of a vague or abstract kind.

Churchland and Ramachandran, on the one hand, and Dennett, on the other, seem to agree, in other words, about the interesting empirical questions yet to be answered. Both want to know: "what is the level at which inferences about the blind spot occur?" and "how rich is the informational content of those judgments?" It is not their *philosophical* intuituions that here diverge but their *empirical* intuitions – namely, each party's best guess about where in the visual processing hierarchy such inferences are most likely to occur.

More importantly, Churchland and Ramachandran seem to miss Dennett's proposal that there need be no single form of visual representation across visual cortical areas – that our phenomenology might be supported by a large variety of representational types and inferential processes. For this reason, they also pass by his contention that we cannot use the standard dichotomies in marking the distinction between the kinds of representations that underlie visual experience and those that are involved in non-visual conceptual thought. We cannot say that what separates visual perception from non-visual thoughts are the standard distinctions between imagistic and sentential processes, between perceptual states and conceptual ones, and so on. Churchland and Ramachandran, however, seem to equate "genuine" visual experience with *neural activity in any visual sensory area*, i.e., any neural area the cells of which can be driven by, at least in part, visual stimuli. Instead, if Churchland and Ramachandran want to make contact with Dennett's point, they must argue that one of these distinctions serves to mark the relevant divide or that, while these particular dichotomies are false, nonetheless our visual experiences are underwritten by an as yet unspecified form of visual representation. Then Churchland and Ramachandran would need to show that visual processes of the hypothesized type are both necessary and sufficient for our visual experience of filling in (e.g., our experience of the artificial scotoma filling in with twinkles). The neurophysiological evidence that Churchland and Ramachandran present, in other words, does not address the issue at the right level of abstraction – at the level of *neural representational structure and content* as opposed to at the level of single-neuron activity. The issue at hand, once again, is this: what kinds of representations result in human visual experience?

Finally, because Churchland and Ramachandran do not address the question of what makes a particular kind of representation a visual one, they miss what, by Dennett's lights, would or would not count as empirical evidence in favour of filling in. Thus Churchland and Ramachandran cite the declarations of the experimental subjects, those who attest to the very visual feel of filling in as evidence for the genuinely visual nature of the representations. Yet such declarations, however sincere, do not meet with Dennett's argument. If the form of a visual event need not match the resultant phenomenology, then Dennett would deny that any conclusions about representations can be drawn from our introspective reflections (or from the sincere declarations of other people about their own experiences). Again, to counter Dennett's view, the following question must be raised: is it really true that conscious visual episodes, even those that seem very vivid and complex, could be the product of inferential processes or judgments alone? It is a positive answer to this question, we think, that lies behind Dennett's denial of filling in.

In these three ways, Churchland and Ramachandran's empirical data and Dennett's theory of consciousness pass each other by.

III. Docking the Freighter

If Churchland and Ramachandran have failed to distil the essence of Dennett's view, this is a position with which many readers of *Consciousness Explained* will sympathize. To a large extent, the authors' interpretative problems result from Dennett's indirect narrative style: as we said earlier, Dennett's solutions (or dissolutions) to specific puzzles about consciousness presuppose a certain relationship between sub-personal neural representations and the phenomenology that we each, as persons, experience, but no account of it is ever explicitly given.

What, then, is Dennett's answer to this central puzzle? Dennett's view, we suspect, is simply this: *it is the content of neural judgments, not their form or structure, that is transparent in conscious experience.* The very feel of our phenomenology, in all its apparent complexity and variety, in some sense *is* the representational content of a narrative stream, a stream that has been precipitated by particular probes or questions. When you listen to the Brahms trio, say, when you delight in the unexpected progressions and wonder at their subtlety and innovation, what you experience is the content of a series of conclusions: here, presumably, conclusions drawn in answer to your own sub-personal expectations or your conscious personal-level queries about harmonic progressions and musical form. The character of your phenomenological experience depends upon your (often sub-personal) judgments about "how things lie" – the state of the external objective world and, in other cases, the state of your own neural processes. In other words, on Dennett's view, "what it is like" is a matter of how things are judged to be.

This view, that the nature of phenomenological events is a function of the content of the relevant neural states, is both radical and puzzling. (In particular, how are we to understand this notion of "transparency"? In what sense does the content of neural events shine through to conscious experience?) It is also a view that, prima facie, does not sit easily with all of Dennett's other philosophical commitments (as we will discuss later). In support of this interpretation of Dennett's theory of consciousness, then, consider the following four examples from *Consciousness Explained,* in which Dennett's explanations of various properties of conscious experience seem to presuppose just this position.

First, in the Prelude, which is given as a preview to the main themes of the book, Dennett presents a theory of hallucinations. The problem with which Dennett begins is this: how could we have detailed and vivid hallucinations (or dreams) in the absence of the detailed sensory information about the world that normally informs visual experience? Psychoactive drugs do not contain within them any "script" or "video"

of the hallucinations which they cause; white noise, by definition, does not contain the content of the hallucinations it often fosters. Dennett's solution is to suggest that, in the absence of orderly sensory input, the top-down hypothesis-generators of normal perception (primarily in vision and audition) begin to produce hypotheses which are then randomly confirmed or negated by noise on the input channels. While, under normal cognitive conditions, such random confirmations would be noticed and questioned by epistemic "demons" whose job it is to check for the consistency and plausibility of the confirmed hypotheses, in the case of hallucinations and dreams, such demons are uncharacteristically lax. They passively accept whatever random hypotheses are confirmed; they show a lack of "epistemic hunger." The conclusions are simply *seen* – or heard or smelled and so on, depending upon the sensory information of the hypotheses.

What is interesting here is that Dennett does not find it problematic that a hypothesis – a linguistic or quasi-linguistic entity – could give rise to *sensory* phenomenology simply in virtue of being *accepted.* How is it that an affirmative answer to the question "Am I hearing a Brahms trio?" could actually cause me to hallucinate, in its full auditory glory, a performance of a Brahms trio? Why would that produce a complex auditory hallucination?

As we saw above, Dennett does not believe that visual experiences must result from pictorial representations, nor, more generally, that each modality of sensory experience has a characteristic representational form. Moreover, in "Seeing Is Believing – Or Is It?", Dennett concludes that sentential or linguistic representations need not be impoverished in their content relative to more pictorial or imagistic forms of representation. As he puts it: *"There is no upper bound on the richness of content of a proposition"* (Dennett's italics) (This volume, p. 162). In other words, if linguistic hypotheses can cause visual sensations, and if there is no upper bound on the propositional content of sentence-like representations, there is no reason why mere hypotheses could not cause rich sensory events – for example, hearing a particular Brahms trio. No reason why not, that is, if one assumes that it is content alone that determines "what it is like."

Second, in Dennett's discussion of the representation of time in Chapters 5 and 6 of *Consciousness Explained,* he argues that no amount of evidence, either from the natural sciences (in the form of neurophysiological findings about the brain or psychophysical research) or from our own introspective reports, could serve to pinpoint the exact time of occurrence of a conscious event. In conscious awareness, says Dennett, we experience only the time as represented, not the time of the neural event as it actually occurs; we have access to the content of

the representation as opposed to the properties of the vehicle itself. According to Dennett, "introspection provides us – the subject as well as the 'outside' experimenter – only with the content of the representation, not with the features of the representational medium itself" (1991, 354). Admittedly, one way to interpret this statement is to give it a very weak reading: perhaps Dennett is saying only that we lack access to the properties of representational vehicles qua vehicles just because, as a matter of fact, we do not know the bridge laws that connect the properties of our phenomenology to any of the other non-semantic properties of the representations.[5] Still, if it were the (practical or theoretical) absence of bridge laws that blocked our access to the non-semantic properties of neural representations, then our access to representational content ought to be equally impaired. To date, that is, there are few commonly accepted bridge laws that link content and experience (and even fewer known to those of us who are not psychophysicists) so *mutatis mutandis*, with introspection as our only guide, the content of our conscious experiences ought to be just as opaque to us as all of the other properties of neural representations. In other words, if Dennett insists that there is a difference in our introspective access between the content of a representation and the other properties of the representational vehicle, then Dennett is making the stronger claim: that representational content is somehow transparent to experience in a way in which the other properties of representational vehicles are not.

Assume, then, that any complex conscious event will require any number of content determinings or conclusions, and that each of these will occur at slightly different times at a variety of neural sites. All are necessary for the occurrence of the conscious event, yet together they have no single point of onset. But if *subjective* experience does not point to any particular moment in time as "the" time of onset, then there is, in principle, no means of pinpointing the exact time, within a given window of opportunity, as "the" one at which a given conscious experience begins.

Criticisms of this argument aside, note that Dennett is not led by this train of thought into anomalous monism (not even on his darkest days). He does not conclude from his view that we have access to *only* the time of conscious events as they are represented, nor that we therefore have *no basis at all* for selecting neural representational events with which to pair the events of phenomenal experience (and hence no means by which to discover regularities between types of phenomenal experiences and types of neural representational events). In theory at least, it is possible to choose, among the multitude of neural events that occur at roughly the right time, those neural representations which are relevant, to select whatever neural events give rise to

a particular conscious perception or thought. And this we are able to do in virtue of matching the *content* of our phenomenological experiences with the content of neural events – by comparing what conscious events are about with what those neural events represent. Again, this is possible only on the presupposition that content both determines and is transparent in phenomenology.

Third, in Dennett's chapter in *Consciousness Explained* "Prosthetic Vision: What aside from Information Is Still Missing?" he describes experiments in which blind subjects were equipped with small video cameras mounted on eyeglass frames, plus a device that translated the low resolution images into an array of electrical "tingles" over the surface of the back. After several hours, the subjects could describe rudimentary shapes, read signs and even identify people's faces. But what of the feelings, asks Dennett, that the device produced? Did the electrical "tactors" succeed in producing prosthetic vision or was the electrical stimulation felt merely as "tingles," as tactile sensations of the usual kind? In reply, Dennett cites evidence for the transparency of the tingles, for the conclusion that the (visual) informational content of the video images informs the blind subjects' experience and thereby gives them visual experience: the subjects' felt point of view was directly ahead, the direction in which the camera was pointing, not directly behind; when the tactors were shifted from the back to the stomach, subjects were able to adapt almost immediately to the change; moreover, itches on the back were still felt *as* itches not "seen" as distal events. On the other hand, as Dennett admits, using the prosthetic device did not produce all of the phenomenological effects of normal vision in the experimental subjects – for example, *Playboy* pictures failed to be "interesting" to two blind male college students. Here, suggests Dennett, the problem might lie in the low spatial resolution of the images, the slow response rate of the "sluggish" tactile receptors and in the lack of neural connections from sensory cortex to areas responsible for emotional affect. He says:

> It is not clear how much would change if we somehow managed to improve the "baud" rate of prosthetic vision to match normal vision. It might be that simply increasing the amount and rate of information, somehow providing higher-resolution bit-maps to the brain, would suffice to produce the delight that is missing. Or some of it.... It might also be that some of the pleasure we take in visual experience is the by-product of ancient fossil traces of an earlier economy in our nervous systems...." (Dennett 1991, 342)

In other words, Dennett believes that insofar as the devices produce visual sensations, this is the result of the information common to both systems; and insofar as the subjects' phenomenology fails to mimic

that of normal vision, informational differences or a difference in neural connections to "earlier economies" account for the discrepancy – and this regardless of how the sense of touch packages its informational content, regardless of the form of the representational vehicles.[6]

As a fourth and final example of the proposed view, consider Dennett's solution to one of the classical problems of consciousness and representation, namely, the nature of secondary qualities, of the "purely qualitative" sensations of color, warmth, hardness and so on. Consider the problem of colours. We know that any number of different conditions can cause the sensation of red: our perceptions of red objects are not caused, on each and every occasion of perception, by a specific wavelength (or combination of wavelengths) of reflected light. We also know that our neural representations of red objects are not themselves coloured red. But if neither external objects in the world nor the inner neural states which represent them are red, where is redness to be found? Given that red is neither "out there" nor "in the head," must it not exist, as Descartes thought, in yet another realm, in the non-material mind? (Here again we see the naïve theory of consciousness at work.)

Roughly put, Dennett's solution to this puzzle is to propose that colours really do exist as properties of objects in the external world – but as extremely complex, heterogeneous properties, those which the human visual system has evolved to "detect." Given the facts about the light-reflective surfaces of the world, the human visual system and its color categories have evolved to make use of those contingencies; indeed, to a large extent, one could say that our color categories have *co-evolved* with the surface reflectance properties of various other biological entities. The color red, for example, is just whatever set of light conditions is now capable of producing in human observers the sensation of red, and this is true no matter how large, unruly or disjoint such a set might be. So colours are properties in the world and our sensations of color are simply representations of those (admittedly) peculiar properties.[7]

Setting aside the issue of primary and secondary qualities, what we want to highlight is why Dennett considers this an adequate answer to the question at hand – to the problem of the *consciousness* of secondary properties. To the average reader, we suspect, his answer will seem to side-step the issue: Dennett does not explain here why the detection of the peculiar property of red gives rise to its particular phenomenology, nor does he say where the *psychological* property of redness resides. Rather, the pedagogical move is to assure the reader that our perceptions of redness are representationally – or rather, informationally – respectable. Note, however, that if informational content is

"transparent" to consciousness under certain circumstances,[8] then Dennett's answer goes further towards providing an explanation than first appears. To understand the informational content of a representational state is to understand why a type of conscious perceptual event has its own particular feel. (This also explains Dennett's answer to Otto's objection that "[y]ou haven't said why pink should look like this!... Like the particularly ineffable, wonderful, intrinsic pinkness that I am right now enjoying" (Dennett 1991, 383). Dennett answers with an explanation of the affective qualities of conscious experiences (e.g., why we enjoy some of them and hate others) and by arguing against any "intrinsic" or "purely qualitative" aspects of colour sensations. In other words, the complete answer he gives to the problem of conscious colour states is in terms of their functional/informational properties – their informational ties to the world and their links to other cognitive/affective states.) Perhaps, then, this is the unspoken view that lies behind the text.

Despite the cautions of Churchland and Ramachandran (this volume), then, we think that Dennett's magic elixir for consciousness is not "Gilbert Ryle's 'Ghost-Be-Gone'." Instead, Dennett is a philosopher trying to float two separate projects that drift in different directions.[9] First, it is true that Dennett has a long-standing philosophical allegiance to both Ryle and Quine, and, just as Churchland and Ramachandran say, these strands in his philosophy seem to restrict the science of the mind's legitimate domain. From Ryle, Dennett inherits the view that our everyday ascriptions of psychological states cannot be construed as descriptions of internal processes.[10] According to Dennett's Intentional System Theory, ascriptions of beliefs and desires, made under the assumption of rationality, serve to capture only large-scale patterns of human behaviour. They do not describe internal cognitive events, nor do they reflect any underlying "true" ontology, the ontology of neurophysiological states. Rather, the ascription of such states provides us with *predictive powers* about human behaviour, powers that are not available to us using other kinds of generalizations (and this includes the generalizations of neurophysiology). Thus, the job of explaining beliefs and desires is viewed as outside the rightful domain of the neurosciences.

Similarly, one can see Quine's influence in the host of roughly Quinean indeterminacy claims that inhabit Dennett's writings – the indeterminacy of meaning, translation and intentional ascription, plus the indeterminacy of biological function, of the time and content of conscious events, and so on. In effect, such claims serve to limit the scope of scientific explanation. If there are certain questions for which, in principle, there can be no answers, then there is nothing that science

– or any other discipline for that matter – can contribute to their elucidation. (E.g., if there is no sense to be made of the time of onset of a conscious state, then there are no experiments to be done that might pinpoint time of occurrence.) Again, the science of the mind may seem denuded.

Despite these seemingly anti-scientific philosophical views, Dennett (like both Churchland and Ramachandran) is a true believer. We are creatures with cognitive powers and phenomenological experiences and science can – will – elucidate and explain these capacities. Moreover, science will do so whether the subject is our own (human) cognition and phenomenology, or whether it is that of the echo-locating bat (should the bat turn out to be conscious).[11] This is the domain of science, a domain it has even if – despite the fact that – at the outset of the investigation, introspection, folk psychology and our commonsense intuitions are not entirely helpful in delineating what body of facts needs to be explained. Even if, at the beginning, we are somewhat confused about exactly how we are.

This faith in scientific explanation motivates Dennett's views about the mind/brain described *at the "sub-personal" level*. For this – the neurofunctional or neurocomputational level – is the rightful domain of science. Here science can ask: what is the essential nature of our cognitive processes? How, in particular, do these processes work? Why is there consciousness at all? And what is the specific relationship of cognitive processes to our conscious phenomenology? Moreover, Dennett's answers to these questions, as given in *Consciousness Explained*, are actually quite theoretically staid. All tread along (or at least next to) the beaten path in cognitive science: they fall within the scope of the computational model of mind. He holds, for example, that cognitive processes are essentially computational processes – the brain is a "syntactic engine" which nonetheless manages "to approximate the impossible task, to mimic the behavior of the impossible object, the semantic engine" (Dennett 1987, 61); that consciousness arises out of the imposition of a culturally transmitted "virtual machine," a serial machine that "sits atop" the brain's more primitive parallel processes (Dennett 1991, 209-26); and finally, that (if our suggestion is correct) we experience what we do because the very shape and feel of a conscious experience is a function of the content of the virtual machine's relevant computational states.

In *Consciousness Explained*, then, we see Dennett attempting to honour his Quinean and Rylean views about persons while holding onto his faith in science as the means of investigating sub-personal neuro-computational processes. Whether Dennett can stay afloat on this divided structure, of course, remains to be seen. Often, we think,

the two separate projects – Dennett's philosophical doctrines about persons (or entire intentional systems) and his scientific/philosophical explanations of sub-personal processing – tend to collide or split apart. This results, prima facie, in a number of inconsistencies, incoherencies or (often) large areas of theoretical unclarity. We will close then by mentioning three such problems that Dennett needs to resolve.

First, Dennett must reconcile the transparency of content with a subject's ability to misdescribe or make mistakes about the nature of his or her own phenomenological experience. On the one hand, Dennett claims that there is no appearance/reality distinction for phenomenological experience. Because any episode of conscious awareness emerges only as the result of a (personal or sub-personal) probe, there is no difference between how things seem to be and how they really are. As we said earlier, on Dennett's view, "what it is like" for the subject just *is* a matter of how things are judged to be. On the other hand, when we as subjects of experience make reports about our conscious lives, it is entirely possible that we are in some way mistaken about the nature of those internal events. By treating such reports as possible fictions in need of confirmation, Dennett's method of "heterophenomenology" presupposes a personal-level capacity for error (Dennett 1991, chap. 4). But in what ways could the subject be wrong about his phenomenology if the content of the experience is transparent to the subject and indeed, may have come about as the result of a conscious probe initiated at the personal level? What needs to made clear, then, is the distinction between personal-level and sub-personal judgments (what these are judgments about) and the relation of each kind of judgment to our capacity for – and immunity to – error.

Second, if content is transparent to consciousness, Dennett must explain why there seems to be no real difference in our phenomenological experience of neural events with abstract content (our representations of non-foveal visual space or our representations of the blind spot) and our experience of neural events that contain detailed information about properties of the world (this is a criticism very close to Churchland and Ramachandran's). That is, when the content of a representation is suitably abstract – "just more Marilyns" or "lots of people in brightly coloured costumes crossing a bridge" – why does it not *seem* abstract to the subject, or more to the point, seem *any different than* our visual experiences that are supported by more concrete representational contents? We don't say, for example "well I see one Marilyn on the wall and the others I see there, um, in a kind of abstract way"; nor do we report that while the very small portion of one Marilyn centered on the fovea seems crystal clear, all of the other Marilyns (and the other parts of the foveated Marilyn) seem completely "fuzzy"

(although we may say that, at the very edges of peripheral vision, the world seems dim or blurred). On Dennett's theory of consciousness, one would think that seeing a wall of Marilyns would be a little like imagining a speckled brown hen – lots of spots, just not any particular spots anywhere in particular. Similarly, when we look at an (actual) bridge in the distance, and see lots of different people, each in bright and individual attire, the foveated portion of the visual scene provides the brain with specific (albeit partial) information about the individuals on that part of the bridge. For non-foveated vision and the blindspot, the visual system draws the conclusion "more brightly and distinctly attired people." But again it is hard to see why there is no apparent difference in our perceptions – no difference between how we see the people for whom we do have specific visual information and how we see those individuals about whom we have made only a vague generalization. (Indeed, here, the generalization cannot be literally "just more of the same" because our impression of is of *distinctly* attired individuals.) In short, Dennett needs to say more about the various forms of abstract content and about the relation of abstract content to conscious experience.

Finally – and most importantly – Dennett requires a notion of computational content that makes sense of his claim that content is transparent to consciousness, or more strongly, that content alone determines phenomenological experience. As we said above, Dennett adheres to a standard computational line about the semantics of sub-personal states. According to this familiar story, brains, as syntactic engines, have causal access to only the physical (or formal or syntactic) properties of mental representations. So computational systems manipulate symbols *in accordance with* their meaning or semantic properties but *by virtue of* their syntactic properties, through causal processes. Syntactic properties are what make the impossible object – the semantic engine – run. In turn, on Dennett's scheme, the content of symbols is determined by biological function: to understand what a particular state of a mechanism means is to step back and assess its natural function, both past and present – its present role in relation to inputs of the system, the mechanism's informational links with various other states, its connections to outputs of the system, and last, the evolutionary history of that mechanism or state. Thus the dual aspects of computational states, on Dennett's view of sub-personal events, their semantic and syntactic properties.

What kinds of properties are available to consciousness, then, according to Dennett? What is puzzling here is that according to the computational picture, content itself could not be transparent in a person's experience, for content is a *relational* property of states, one that often depends upon facts beyond the cognitive reach of the subject (say, upon the evolutionary history of a sensory mechanism). In a real sense, content could not be a legitimate or felt participant at all. On the other hand, Dennett claims that *none of the non-semantic properties of the representational vehicle are available to the subject in conscious experience.* Recall Dennett's statement, cited earlier, that "introspection provides us – the subject as well as the 'outside' experimenter – only with the content of the representation, not with the features of the representational medium itself" (1991, 354). Given Dennett's computational explanation of sub-personal processing, there would seem to be *no* properties of neural representations that could be felt in conscious experience.

A natural response to this dilemma is for Dennett to draw a distinction between the syntactic properties of representations and any other (non-semantic) properties that a representational vehicle might have. Let the syntax of the representational vehicle "go proxy" for content in the function of the syntactic engine. Let all of the other properties of representational vehicles, the "purely physical" or non-syntactic characteristics of representational states, remain hidden from view. In other words, the line drawn between that which finds its way into consciousness and that which does not is a distinction between two classes of non-semantic properties, between syntactic properties on the one hand and purely physical properties on the other. If Dennett wishes to maintain the computational party line, drawing this distinction might make for a plausible revision of his view.

In fact, Dennett's writings do not provide a stable answer to this problem nor does he fill in any of the details that would be necessary to flesh out this computational strategy.[12] To do so, he must provide an explanation of the difference between syntactic and non-syntactic (yet non-semantic) properties of representations. He must also give an explanation of why the difference has such profound consequences for human experience – why syntactic properties alone are transparent to consciousness. It is here, then, that we come to the most serious hurdle for Dennett's theory of consciousness: in order for Dennett to explain why syntactic properties alone are transparent to consciousness he must give a general explanation of the mysterious notion of transparency itself: what transparency amounts to given the prima facie profound difference between the syntactic properties of neural representations and the phenomenal properties of individual experiences, and why in general transparency occurs. In other words, if we are correct about Dennett's implicit views on content and consciousness, the heart

of the explanation of phenomenal experience has yet to be given. So, we will close by simply noting that Dennett needs to sort through these issues – and note that this, we suspect, will be a rather lengthy project.

Acknowledgments

This paper was presented in an earlier draft in November 1992 at the McDonnell-Pugh Seminar Series in cognitive science at Oxford. For their comments and helpful criticisms, we would like to thank Patricia Churchland, Daniel Dennett, Martin Davies, Denis Robinson and David Rosenthal.

Notes

1 It is very hard to speak of the relationship between neural processes and conscious events without speaking of the neural states as "giving rise to," "underlying" or "supporting" conscious events – that is, without the appearance of dualist sentiments. Despite appearances, this is not a philosophical commitment we intend to express.

2 This example, "Marr's Paradox", was first given by Kathleen Akins in a review of Michael Tye's book *The Imagery Debate* forthcoming in *The Philosophical Review.*

3 Note here that Dennett does not consider any other reasons, except the demands of the fictitious Cartesian Theatre, why the regular integration of visual information might be necessary. Nor does he explain his notion of a "probe," those judgments which constitute the individual acts of consciousness. As probes are judgments about the contents of individual modules at particular times, probes are surely agents of integration. But Dennett does not say how such "gathering together" occurs.

4 It is helpful to note here that, at the retina, there are no transduction cells in the area of the blind spot and hence no ganglion cells to carry the outgoing signal for this area of visual space either. So, at the first level of cortical processing, the lateral geniculate nucleus (LGN), there is a map of visual space but one from which the blind spot is simply omitted. There is no hole in the map, an area which contains only silent neurons or none at all; rather, those neurons that are activated by stimulation *around* the blind spot sit *side-by-side* in the neural map of the LGN. At the next cortical level, however, at V1, there *are* neurons in the neural map that cover the blind spot (from Churchland and Ramachandran, this volume; see also Fiorani et al. 1990; Gattas et al. 1992; Fiorani et al. [forthcoming]) – although, as Churchland points out, it is not clear what kind of information their signals contain.

5 We owe this objection to Martin Davies.

6 We cannot resist pointing out that somehow Dennett has missed the following point: that if the "very same" information about a woman's shape were obtained through more *normal* tactile means (yes, even through the back), the desired effects might well occur in the blind subjects – and this despite the sluggish response and low spatial resolution of the tactile transducers. So it seems highly unlikely that the human sense of touch lacks the sort of "ancient fossil traces" of links to "earlier economies" that are required for the experience of pleasure.

7 Although Dennett attributes this view to Akins (1989), Dennett and I draw somewhat different metaphysical conclusions about secondary properties. Although this is not the place to attempt to develop a complete account of properties, suffice it to say that, on my view, transducers with standard nomological causal properties need not "capture" (be responding to) a property in the world, not even a heterogeneous property. *Transducer mechanisms need not be used as property detectors* – the mere existence of such a transducer does not "create" a property in the world, nor are the outputs of such systems necessarily used either as representations of properties per se or as "content determinings" in Dennett's terminology. On my view, in order to claim that a sensory state is used as a representation of a secondary property (or that such a secondary property exists), a lengthy ontological story must be told about how organisms are able to use certain sensory states as representations of objects and their properties.

8 Namely, on Dennett's view, when it is incorporated into the working of the serial virtual machine.

9 Call this the "Multiple Air Mattress Theory of Consciousness."

10 Dennett also adopts Ryle's deep suspicion of "introspection" as a legitimate means of access to whatever neural processes actually do occur. In fact, Ryle held, in *The Concept of Mind,* that we do have conscious experience of our own thoughts, but that this process cannot be explained as "introspection" or "inner perception" or as any other kind of monitoring of one's current mental state. Rather, what we enjoy is "retrospection," a process that occurs only when we chance to ask ourselves what we have been thinking. As such, it is open to all the normal failures of any other kind of recollection or memory. Ryle says that "aside from the fact that even prompt recollection is subject both to evaporations and to dilutions, however accurately I may recollect an action or feeling, I may still fail to recognize its nature.... Chronicles are not explanatory of what they record." This is a theory that sounds for all the world like Dennett's own Multiple Drafts model – just as Churchland and Ramachandran would have it.

11 Dennett's clear support for the neurophysiological investigation of the "alien" consciousness of other animals might seem to contradict our interpretation of his view. After all, if it is only the representational *content* that is

transparent to conscious experience, and if this content is not affected by the way in which it is neurally packaged, how could the study of the vehicles of content – the packaging – lead to an increased understanding of another subject's phenomenology? What could neurophysiology possibly add to "outer" investigation, to purely behavioural discriminatory tests (e.g., to tests of whether the bat can discriminate between a freely fluttering insect and an inert meal worm tossed into the air by an experimenter)? Moreover, our understanding of computational states suggests that if any feature of neural events is a reasonable candidate for offering insight into phenomenology, it is surely the *structure* of the subject's present representations and the form of the representational system as a whole, its representational relations. Even Thomas Nagel, for all his reasoned objections to the possibility of third-person access to consciousness, sometimes speaks as if neural structure might give us insight into certain "structural" aspects of consciousness. And Dennett himself says: "we learn about the bat's world by learning the systematic structure of bat perception and behavior, not by imagining bat worlds or imagining our minds turned into bat minds." (1991, 442).

The answer to this apparent contradiction lies in Dennett's understanding of neural function, of what neural mechanisms in the brain are designed, by and large, to do. Throughout *Consciousness Explained*, Dennett speaks of individual processors as performing the common task of "content discrimination." It is the product of such discriminations or "determinings" that are subject to continual revision, that are discerned intermittently by probes, that compose the serial narratives constructed out of the probe's findings, and so on. (Note how well this conception of things sits with the low-level task of Marr's theory of vision, namely the task of making explicit information from the original visual image.)

Two things follow from this view. First, if "content determining" is the central task of (most) neural processors, then, at least theoretically, behavioural discriminatory tests are capable of finding out the processing results of (most) brain functions. We watch the bat to see its discriminatory capabilities. Neurophysiological investigation, on the other hand, merely offers a more direct route to discerning which content determinings are going on. "Looking in the head" while monitoring the informational input to the brain is often a speedier and more reliable scientific methodology. Second, if one views brain function in this way, one might also reasonably assume that the logical relations between the content determined will be mimicked by the structural relations between the vehicles of representation – that the systematicity of the content qua content will be mirrored by the structural systematicity of neural representational vehicles. So on this view, to look into the brain, by neurophysiological means, and see the relevant structures (the structure of the processing results as opposed to the structure of the processes themselves) *is* to look into the brain and see the appropriate content

determining – the two ride together, like a numeral and its number. On this view of neural functioning, then, it is natural to equate structure with content. And we suspect that this is what Dennett does in this and like passages.

Note that if the raison d'être of brain processes is not content discrimination (if one cannot always classify the processing answers in this way) or if the processes themselves are part of conscious phenomenology, then no such convenient equation between content and structure can be made.

12 For example, in the case of prosthetic vision given above, Dennett suggests that perhaps the absence of cortical links, from somatosensory cortex to "earlier economies" accounts for the lack of pleasure taken in the nude photographs. In other words, he seems to believe that the *functional facts on the basis of which we would ascribe certain content to normal visual "images"* (here, the cortical connections to earlier economies) are themselves responsible for a certain phenomenology – for feelings of pleasure. Then again, perhaps not. Perhaps these informational links from visual cortex to sub-cortical sites result in sub-cortical representations, the syntax of which finds its way into experience as felt pleasure. We think, though, that there is little point to this kind of speculation, for it is not clear that Dennett has a settled or considered view on the matter.

References

Akins, K. (1989) On Piranhas, Narcissism and Mental Representation: An Essay on Intentionality. Ph.D. dissertation. AnnArbor, MI: University of Michigan

———. (forthcoming) Review of *The Imaginary Debate* by Michael Tye. *The Philosophical Review*

Dennett, D.C.D. (1987) Three kinds of intentional psychology. In D. Dennett, *The Intentional Stance*. Cambridge, MA: MIT Press

———. (1991) *Consciousness Explained*. Boston, MA: Little, Brown

Fiorani, M., M.G.P. Rosa, R. Gattass and C.E. Rocha-Miranda (1990). *Society for Neuroscience Abstracts* 16: 1219

Fiorani, M., M.G.P. Rosa, R. Gattass and C.E. Rocha-Miranda (forthcoming). Visual responses outside the "classical" receptive field in primate striate cortex: A possible correlate of perceptual completion. *Proceedings of the national Academy of Sciences*

Gattas, R., M. Fiorani, M.G.P. Rosa, M.C.F. Pinon, A.P.B. Sousa and J.G.M. Soares (1992). Changes in receptive field size in V1 and its relation to perceptual completion. In R. Lent (ed.), *The Visual System from Genesis to Maturity*. Boston: Birkhauser

Ryle, G. (1949) *The Concept of Mind*. London: Hutchison