The Content of Perceptual Experience

John McDowell

The Philosophical Quarterly, Vol. 44, No. 175 (Apr., 1994), 190-205.

Stable URL:
http://links.jstor.org/sici?sici=0031-8094%28199404%2944%3A175%3C190%3ATCOPE%3E2.0.CO%3B2-I

The Philosophical Quarterly is currently published by The Philosophical Quarterly.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at http://www.jstor.org/about/terms.html. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at http://www.jstor.org/journals/philquar.html.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

http://www.jstor.org/
Tue Aug 23 05:55:12 2005

The Philosophical Quarterly Vol. 44, No. 175 *April 1994*
ISSN 0031-8094

THE CONTENT OF PERCEPTUAL EXPERIENCE

BY JOHN MCDOWELL

1. Daniel Dennett's aim, in his richly suggestive paper 'Toward a Cognitive Theory of Consciousness',[1] is to represent the content with which we persons have conscious dealings as a selection from the content that would figure in a sub-personal, cognitive-scientific account of the operations of our internal machinery. What effects the selection, according to Dennett's suggestion, is the fact that some of that sub-personal content is available to an internal public-relations organization that accounts for our linguistic output – not that we actually state everything that is in our consciousness, but the idea is that we could do so.

What is in question is precisely access to content, rather than something we could without qualification conceive as access to content-bearers: for Dennett remarks (p. 159) that 'we have no direct personal access to the *structure* of *content*ful events within us'. The idea is that 'events within us' are contentful *by virtue of* their structure: for the content possessed by an internal event or state is a function of its function in the organism, and this function 'is – in the end, must be – a function of the *structure* of the state or event and the systems of which it is a part' (p. 163). But only some of the events and states within us that possess content in this way make available to *us* the content that their structurally determined role in the system confers on them (although they transmit content freely among themselves, unnoticed by us). And when they make their content available to us, the structure in virtue of which they have it does not figure in an accurate phenomenology of our consciousness; it remains a topic for theory, not mere introspective noticing.

I think it is phenomenologically acute of Dennett to deny that we have 'direct personal access' to structure; the background thesis, that the content in question is possessed in virtue of structure, generates familiar temptations to suppose otherwise. Consider visual experience. The relevant internal

[1] In his *Brainstorms* (Hassocks: Harvester, 1978), pp. 149–73. Page references to Dennett, unless otherwise specified, are to this paper. Dennett has discussed issues about consciousness more recently, in *Consciousness Explained* (Boston, Toronto, London: Little, Brown and Co., 1991). But I think the earlier work is still worth attention. Dennett (in conversation) has agreed that the book does not supersede what I take issue with in the paper.

© The Editors of *The Philosophical Quarterly*, 1994. Published by Blackwell Publishers, 108 Cowley Road, Oxford OX4 1JF, UK and 238 Main Street, Cambridge, MA 02142, USA.

event will be described, in the cognitive-scientific framework Dennett is working in, as something on the lines of a computation of a representation of part of the environment from a pair of arrays of intensities and wavelengths. The way information is contained in this base for computation is naturally described as imagistic. If we assume that the contentful consciousness involved in a visual experience is a matter of access to such an event, it can be very tempting to equate the plain fact that our visual consciousness is of how things *look* to us with the theoretical idea (not a plain fact of consciousness at all) that what we have access to in this sort of consciousness is more than the contents of the computationally derived representations (at various levels: see pp. 157–8), and includes specifically the image-like character of the base of the computations – a matter of structure in Dennett's sense. But whatever may be true about the information-processing that takes place in the visual system, there are no images (two-dimensional arrays) in the phenomenology of vision: it is the relevant tract of the environment that is present to consciousness, not an image of it. It is to Dennett's credit that he resists this falsification of what visual consciousness is like.

2. Another thought of Dennett's seems less acute as phenomenology. Having introduced the notion of presentiments or premonitions by way of cases in which, for instance, one is struck, without knowing why, by the thought that someone is looking over one's shoulder as one writes (pp. 165–6), he goes on to apply the notion to ordinary visual experience (p. 166):

> Right now it occurs to me that there are pages in front of me, a presentiment whose aetiology is not known directly by me, but which is, of course, perfectly obvious. It is my visual system that gives me this presentiment, along with a host of others.

The suggestion is that these perceptual 'presentiments' are unlike, for instance, the presentiment that someone is looking over one's shoulder only in their connectedness with what precedes and follows them (which is presumably what makes the aetiology so obvious). They are like that sort of presentiment in that the aetiology, although obvious in this case, unlike that one, is, as we might say, phenomenologically extrinsic ('not known directly by me').

This suggestion seems phenomenologically off key, perhaps especially about visual experience. What it seems to threaten is the presentness to one of the seen environment. On Dennett's suggestion, that a seen object is there before one is a mere premonition, something one finds oneself inclined to suppose, unaccountably so far as anything contained in the

© The Editors of *The Philosophical Quarterly*, 1994.

experience itself goes. Or perhaps we can change the aspect, and say that the presence of the object is a hypothesis, the obviously best explanation of the premonition; here the claim of obviousness cannot undo the damage done by the idea of a hypothesis – it cannot give us back the idea that the object itself is presented to one's awareness.

Consider a basic (demonstratively expressible) singular empirical judgement, say a judgement one might express, in a suitable perceptual situation, by saying 'That cat is asleep'. The content of such a judgement depends on the perceived presence of the cat itself. A premonition would at best yield content to the effect that a cat that is in a certain region (and has, no doubt, all manner of visible properties, but registering the richness of the available content does not help) is asleep. This lacks the particularity of the original judgement – its relating to a particular cat (*that* one, as one will be able to say if one is in the right perceptual situation), not just to some cat that satisfies a general specification, however rich. We might try to recapture particularity by making out that the thought is carried to the right object by the 'obvious' hypothesis about the premonition's aetiology. But that is not how demonstratively expressible thought makes contact with its object. Such thought does not need to be 'carried to' its object by a hypothesis, because the object is directly there for the thinker.[2]

3. It may seem captious to complain about phenomenological niceties. But I think the off-key phenomenology reflects a serious epistemological difficulty.

Consider Kant's advance over Hume. Hume inherits from his predecessors a conception according to which no experience is in its very nature, intrinsically, an encounter with objects. What Kant takes from Hume is that there is no rationally satisfactory route from such a predicament to the epistemic position that we are in (obviously in, we might say). Transcendental synthesis (or whatever) is not supposed to be such a route: the whole point of its being transcendental, in this context, is that it is not supposed to be something that we, our familiar empirical selves, go in for. It would be a mistake to think we can domesticate Kant's point by detranscendentalizing the idea of synthesis, so as to suggest that the idea of encountering objects is put in place by interpretation of data, perhaps by inference to the best explanation, with the interpretation being something we do, or at least something that might figure in a 'rational reconstruction' of our being in the epistemic

[2] See Gareth Evans, *The Varieties of Reference* (Oxford: Clarendon Press, 1982), especially ch. 6.

© The Editors of *The Philosophical Quarterly*, 1994.

position we are in. That would just be missing *Hume's* point. Kant does not miss Hume's point. He builds on it: since there is no rationally satisfactory route from experiences, conceived as, in general, less than encounters with objects, glimpses of objective reality, to the epistemic position we are manifestly in, experiences must be intrinsically encounters with objects. But how could they be that if their aetiology were phenomenologically extrinsic?

Dennett's idea of experiences as presentiments requires a pre-Humean epistemological optimism. Not that Dennett would dream of urging epistemologists to revert to such a stance: his attention is simply not fixed on epistemological considerations. I think this sort of situation is not uncom-mon: we have a new-fangled move in the philosophy of mind, enmeshed in a quite old-fashioned philosophical difficulty, in a way that has no connection with the intentions of its proponents (who imagine that their questions and answers are simply insulated from that kind of thing).

Of course I am not suggesting that Dennett's conception of experiences as presentiments is the same as Hume's conception of experiences as impressions. But both conceive experiences as less than encounters with objects. Hume's good point, the one Kant builds on, generalizes beyond the specifics of his own picture, to warrant an epistemological pessimism about any such conception of experience.

4. I do not believe that the off-key phenomenology is just a gratuitous slip on Dennett's part.

Consider the framework role played in Dennett's thinking by the idea of what our visual systems tell us. His main aim is to capture the thought that *we* (and not just some sub-personal parts of us) are on the receiving end of this telling. I shall come back to that; but for present purposes, we need to focus rather on the content of the telling. What could our visual systems, conceived in the information-processing vein that Dennett is defending, tell us? (Or tell our brains, if one prefers something modelled on the original formula 'what the frog's eye tells the frog's brain'?)

Here is a candidate for being a possible message to me, on occasion, from my visual system: that there is a cat, with such and such properties, at such and such a position in my egocentric space. The considerations in §3 above actually suggest that the framework undermines even this, since it undermines the possibility that experiences, as it conceives them, might possess objective content at all; but we can let that pass for now, in order to get the 'presentiment' idea going. The relevant question at this point is rather this: could my visual system tell me, in addition, that a cat figures in the aetiology of the original message? How could my

© The Editors of *The Philosophical Quarterly*, 1994.

visual system, conceived as a sub-personal computing device, be in a position to tell me that? We are letting the visual system pass muster as capable of telling me about such things as the presence of cats with such and such properties, on the basis that it discriminates such circumstances more or less reliably. (Actually the truth is that it enables *me* to do that rather than that *it* does that, and this is crucial; but I am playing along with a different way of talking.) But my visual system is surely not a reliable discriminator of cases in which the input from which it starts has one kind of aetiology as against cases in which it has another. If, as a matter of routine, the visual system added a suitable message about aetiology to whatever it told me about the environment ('This message was caused by the fact it reports'), and if the visual system *is* a more or less reliable detector of features of the environment, the added message would be more or less reliably correct. But it would not be, as we are supposing the original message about the environment might be, a case of *informing* me of something. From the perspective, as it were, of the information-processing device, it would be more like an expression of blind faith: not the sort of thing that belongs in a sensible theory of the functioning of an information-processing device. (A routine additional message that might make a certain sense would be 'This is your visual system speaking'. This leaves Dennett's 'presentiment' idea untouched: it merely registers that the presentiments can be distinctively visual.)

5. If I am right so far, Dennett's basic framework necessitates the 'presentiment' idea: perceptual experiences as he conceives them would have to be presentiments. And the idea is deeply unsatisfactory, in a way that is not just a matter of phenomenological nuance. What has gone wrong?

To begin on a diagnosis, I want to turn away, for the moment, from Dennett's ultimate aim of accounting for personal-level psychological truth, and consider the perceptual lives of frogs. Dennett alludes to a famous paper called 'What the Frog's Eye Tells the Frog's Brain',[3] but he commends a suggestion, by Michael Arbib, that one might prefer the formula 'What the frog's eye tells the frog' (p. 163). His point is that 'sub-personal' content-ascription in the theory of frog perception is controlled by the requirements of a biological account whose topic is the life of *frogs* rather than the doings of their parts.

Still, the original formula 'what the frog's eye (or visual system) tells the frog's brain' was not wrong. What it enables us to stress is that we understand the 'sub-personal' metaphor of telling in terms of informational transactions between one part of an organism and *another part*. So

[3] J. Y. Lettvin *et al.*, *Journal of the Institute of Radio Engineers* (1959), pp. 1940–51.

© The Editors of *The Philosophical Quarterly*, 1994.

if we want to talk of informational transactions between part of a creature and the *creature*, we cannot simply carry over the metaphor; we need to work for the extension. In the metaphor, our parts talk to one another; they do not, at least in general, talk to us. Dennett does a great deal of work at the necessary extending of the metaphor for the case of persons, making room for the idea of what our visual systems tell *us* (this is the main aim of his paper). But the point should hold for frogs too. In a 'sub-personal' account of frog perception in which the frog's eye (or visual system) does some telling, say announcing that there is a bug-like object at such and such a position in the frog's motor space, what gets told of this will be another *part* of the frog, say one that the theorist labels 'motor control'. What entitles us to say, not just that the frog's visual system informs the frog's motor control of the presence of a bug-like object, but that it gets the message to the *frog*? It is part of Dennett's own point that there is no extra twist in the 'sub-personal' account of what happens in frogs, analogous to the extra twist in the sub-personal account of our inner workings that is supposed to make room for *us* to have access to the content of our inner states and events. How, then, does the *frog* get into the act?

I suspect that this question – which is, I insist, a serious one – tends to be suppressed because of an unfortunate feature of the otherwise excellent distinction between the personal and the sub-personal. Theories of internal information-processing in frogs are at best 'sub-personal' (I have needed the scare quotes at several points), not sub-personal, because there are no persons around in contrast with whom we can mark the standard distinction. It would be easy to think on these lines: the frog's being informed, by its visual system, of the presence of a bug-like object would certainly not be a personal-level involvement with content; so it is sub-personal; so why not simply *identify* it with the 'sub-personal' content-involving transaction that we already have in our theoretical sights, the frog's motor control's being informed of that by its visual system? But this would be confused. The point of saying that the theory of internal information-processing in frogs is 'sub-personal' is not that no persons are involved, something that is indeed equally true of talk whose subject is frogs themselves, but that the fundamental idea of such a theory is the idea of informational transactions between *parts* of frogs. If we speak in all seriousness (and why should we not?) of *frogs* learning about their environment through vision, what is in question is, by all means, not personal involvement with content. But it is *froggy* involvement with content, and it ought to be just as pressing to ask how this connects with sub-froggy informational transactions – how the frog gets into the

© The Editors of *The Philosophical Quarterly*, 1994.

act – as it is to ask how our personal involvement with content connects with sub-personal informational transactions within us.

What is more, if it were right to suppose that the personal-level involvement with content which we enjoy when we learn about the environment in (conscious) experience is a matter of access on our part to some of the sub-personally generated content that is being passed from part to part within us, it should be just as plausible to suppose that the involvement with content that is enjoyed by a frog when it learns about the environment in vision is to be understood in terms of access on the frog's part to its own interior. But this seems merely ludicrous. The frog's access in perception is, like ours, to the environment.

6. When we apply the idea that frogs learn, through vision, about features of their environment, we are subject to two controls. First, there is frog life, which, like all animal life, is a matter of more or less competently inhabiting an environment. In this context, we ask questions like the following: what features of the environment would a creature need to become informed of, in order to live in it with precisely the competence that frogs display? Second, there are the facts about frog perceptual equipment. Here we have questions like the following: is *this* the kind of thing that we can make sense of a creature's becoming informed of by the use of, say, eyes? Or, more specifically: can we understand how possession of a visual system that works like *this* makes it possible for a creature to become informed of just *these* features of its environment?

The specific questions that arise under the second head can be answered, in principle, by constructing characterizations of the relevant perceptual equipment as information-processing devices, which transmit their results to other (functionally specified) parts of the organism (for instance, 'motor control'). There is an obvious interplay between the results of investigation here and the answers we give to questions under the first head. Casual observation of frog life might induce the provisional thought that frogs become informed, through vision, of the presence of bugs. Then it turns out that a good theory of the relevant perceptual equipment fails to support the view that the equipment processes information about arrays of light into information about the presence of bugs. The equipment hardly processes information at all (it is a limiting case of an information-processing device), but rather simply reacts to any small moving speck. It is better to view the informational transaction as the transmission, to 'motor control', of information to the effect that a small moving speck is at such and such a point in motor space. So we recast our conception of what frogs become informed of: at best the presence of a bug-like object at a certain place. (Given the usual sort of environment that frogs inhabit,

© The Editors of *The Philosophical Quarterly*, 1994.

this is good enough for their somewhat low-grade competence.) Some may think even this goes too far; but it is hard to see how we could credit frogs with being less informed about their environment than this without representing them as not in touch with it at all – a position that has all the appearance of a philosophers' prejudice.

The fact that there is this perfectly intelligible interplay between what we decide we can correctly say, in content-involving terms, about frogs, on the one hand, and the detail of a content-involving (information-processing) account of the inner workings of the parts of frogs, on the other, is no reason to mix the two stories together. In the account of inner workings, one sub-froggy part of a frog transmits information to another: the frog's eye talks to the frog's brain, not to the frog. In the sense in which the frog's eye tells the frog's brain things, nothing tells the frog anything. We may still want to say that the frog gets told things. But what does *this* telling is not something in the frog's interior; that is what generated the idea that we could attribute dealings with content to the *frog* only if we credited it with something like introspection. Rather, what tells the frog things is the environment, making features of itself apparent to the frog, equipped as it is with frog perceptual apparatus. This is a different metaphor of telling, not in competition with the 'sub-personal' one. It is essential not to be misled by the enormous capacity for illumination that the 'sub-personal' account has (together, perhaps, with the true but irrelevant point that frogs are not persons) into thinking that the 'sub-personal' account exhausts the content-involving truth in this area of biology. The second metaphor encapsulates a whole extra field of truths. What is more, the involvement of content here, and only here, is literal: underneath the metaphor of the environment telling the frog things, we have the literal truth that the frog becomes informed of things. Whereas the content-involving truth at the 'sub-personal' level is irreducibly metaphorical.

The 'sub-personal' account of a sensory system, which treats it as an information-processing device that transmits its informational results to something else inside an animal, cannot adequately characterize what its sensory systems are for the animal (as opposed to what they are, metaphorically speaking, for the internal parts that receive the results of the information-processing): namely, modes of sensitivity or openness to features of the environment – not processors of information, but collectors of it.

It would be a confusion to think that the distinction I am making here is blurred by the minimal extent of information-processing in frog vision in particular. What the frog's eyes do for the frog is to put it in touch

© The Editors of *The Philosophical Quarterly*, 1994.

with moving specks in its spatial environment: things that are in fact bugs, in the sort of case that is sufficiently normal in frogs' lives. From the frog's point of view, its eyes enable it simply to pick up the fact that there is a moving speck (with luck, a bug) out there. From the point of view of the frog's 'motor control' (to speak in the terms of the 'sub-personal' metaphor), the presence out there of a moving speck is rather (at most) the best hypothesis the eyes (or, probably better in view of how little the eyes do, the whole system) can come up with in order to account for the input of light (what is in fact light, though the system does not even know this much) to the eyes. If all goes well, the frog is in direct touch with a feature of its external environment; the internal information-processing system is in direct touch only with structural properties of the immediate inputs to it – which, in the metaphor, it interprets as clues to the nature of the external environment. (Of course the frog does no such thing.)

What could an internal information-processing device really tell an animal? To give a positive answer, we should need to deal satisfactorily with the question I am suppressing, about how to make sense of the frog's being on the receiving end of 'sub-personal' telling; but my point now is not that we have no inkling how that might be done. What could an information-processing device *really* tell *anything* (including another component in a sub-personal or 'sub-personal' informational system)? It is essential to realize that the answer to this question can be, in fact is, 'Nothing', without the slightest threat being posed to the utility, or even the theoretical indispensability, of cognitive science.

A sub-personal or 'sub-personal' informational system is a physical mechanism, connected to its surroundings by transducers that convert physical impacts from outside into events of the sort that the system can work on, and perhaps by transducers that convert the system's end-products into physical interventions in the exterior. The system knows nothing even about the character of the immediate physical impacts on the input transducers, or the immediate physical interventions in the exterior that result from its operations by way of the output transducers, let alone about the nature and layout of the distal environment. The operations of the system are determined by structures exemplified in the initial contributions of the transducers, and in intermediate events and states in the system, which have no meaning for the system. In short, in Dennett's own memorable and exactly right phrase, the system is a syntactic engine, not a semantic engine.[4] The same goes for its parts.

[4] The idea is implicit in his remarks about the role of structure in sub-personal content-attribution (p. 163). For the phrase itself, see 'Beyond Belief', in Andrew Woodfield (ed.), *Thought and Object* (Oxford: Clarendon Press, 1982), pp. 1–95, at p. 26.

© The Editors of *The Philosophical Quarterly*, 1994.

Animals, by contrast, are semantic engines. To stick with the present example, they become informed that ... (say, that there is a bug-like object at such and such a position). The background against which this makes sense is their competent inhabiting of their environment. Now this competence would be quite mysterious if there were no interestingly structured machinery inside them, controlling their behaviour in a way that is responsive to impacts from the environment. We could not make sense of the competence that enables us to make sensible use of the claim that *animals* have dealings with content if we could find nothing inside them but, say, a completely homogeneous jelly. And nobody knows how to make sense of an animal's internal control mechanism, and connect it conceptually to the competence it is supposed to explain, except by describing it *as if* it were, what we know it is not really, a semantic engine, interpreting inputs as signs of environmental facts and, as output, directing behaviour so as to be suitable to those facts in the light of the animal's needs or goals. To insist that the attribution of content at this sub-personal or 'sub-personal' level is 'as if' talk is in no way to debunk it. The content-attribution is not, as it were, irresponsible: it is constrained by the physiological facts, in a way that is exemplified, on a small scale, by the discovery of how little interpretation ('as if' interpretation, we must say) can be credited to the visual systems of frogs. And it is surely clear, at least in a general way, how content-attribution that is only 'as if' can even so pull its weight in addressing a genuine explanatory need: the question is what enables us animals to be the semantic engines we are.

It is crucial to see that the question about real content with which we are helped by the 'as if' attribution of content to states and events in internal mechanisms is this causal or enabling question. One can easily fall into a temptation to suppose that the question is rather a constitutive one. If we could see dealings with content on the part of animals as somehow constitutively explained in terms of information-processing in their interiors, that might seem a protection against a metaphysical embarrassment. (After all, we might say to ourselves, cognitive science is science: maybe it is not quite clear that ecology and ethology are science.) But this temptation is disastrous: if we offer a constitutive explanation of genuine content in terms of a merely 'as if' attribution of content, we make genuine content fragile and problematic.

Dennett's basic picture is that *our* dealings with content are nothing but our access to some of the content manipulated by our internal information-processing systems, and this seems to be a case of succumbing to the temptation. It flies in the face of the insight that the

© The Editors of *The Philosophical Quarterly*, 1994.

internal systems are only syntactic engines: access to our interiors cannot be what constitutes our dealings with content, since there is no content in there, although it is enormously useful to talk as if there were. Dennett writes ('Beyond Belief' pp. 26–7):

> Somehow, the syntactical virtuosity of our brains permits us to be interpreted at another level as semantic engines – systems that (indirectly) discriminate the significance of the impingements on them.

The idea that our discrimination of the significance of the impingements on us is *indirect* reflects the idea that our becoming informed of environmental facts just is the upshot of the sort of computational process that we attribute to our perceptual systems – as if we were in the predicament of our nervous systems, blocked off from the environment by transducers rather than inhabiting it. No wonder our status as semantic engines becomes a mystery ('somehow'), and no wonder it is a comfort to make room for the suggestion that it is just a way of talking, not a fact (the syntactical virtuosity of our brains 'permits us to be interpreted as' semantic engines, rather than just making it intelligible that we are such). If we drop the attempt to read the envisaged explanations constitutively, 'permits' can take its proper significance: unmysteriously, the syntactic virtuosity of our brains enables us to relate to the environment in the direct way that is constitutive of our being the semantic engines we are.

7. Sidney Morgenbesser is said to have accused a cognitive-scientifically minded colleague of believing that our intelligence is Artificial Intelligence. We can now see a sense in which that ridiculous belief is almost correct: we ourselves have genuine intelligence, but there is Artificial Intelligence inside us – not in the sense of an artefact, but in the sense of an imitation or fake. That is to say: we have inside us something that is not intelligent at all (it knows nothing and understands nothing); even so, we can be enormously helped in finding it comprehensible how we can be intelligent, even though we are made of nothing but the stuff of which everything is made, including mere 'stupid' *things* like sticks and stones, by seeing how this completely unintelligent internal equipment of ours can have imposed, on top of the truth about its mindless manipulations of structures that are meaningless to it, an 'as if' description in terms of dealings with content. That makes it possible to understand how this mindless internal control system enables us to do what it takes to display genuine mindedness, namely to live competently in an environment.

'Display' may mislead here, encouraging the thought: maybe that is

© The Editors of *The Philosophical Quarterly*, 1994.

what it takes to *show* mindedness, but it can be there anyway, perhaps in a brain in a vat. I think we are now in a position to begin to see through this. I shall approach the point by way of another striking philosophical remark that we can now see in the proper light, namely John Searle's claim that we are brains in vats (vats of bone, not glass, with input and output linkages to the environment different from those in the standard mad-scientist fantasy; but vats none the less).[5] The truth is that our brains are indeed brains in vats, and that is exactly why we must not identify ourselves with them. To a brain it is all one whether its vat is glass or bone, and what, if anything, is outside its input and output transducers. To repeat, a brain knows nothing and understands nothing: all it does is to manipulate structures that have no meaning for it. That is not the truth about us. Without any threat at all to the enormous power of cognitive science to enable us to explain our mindedness (in one sense of 'explain our mindedness'; not constitutive explanation), we ought to be able to see that the sheer fact that a brain is going through the motions that an embodied brain goes through when a person thinks or experiences is by itself no ground at all for supposing that there is a mind in there. (It may be a different matter if the mad scientist's vat contains what used to be an embodied brain; perhaps memory can give subjectivity a tenuous foothold there.)

There is a persisting inclination to suppose that this cannot be right: if the brain goes through the right paces, it must at least *seem* to it that it thinks or enjoys experience, and then even if we can make out a difference between having it seem to one that those things are so and their being so, the seeming is enough for subjectivity. Here it really is to the point to respond: you might as well suppose it seems to an electronic calculator that things are thus and so.

8. I have been urging, in effect, that we take Dennett's distinction between attributing content at the personal level and attributing it at the sub-personal level as a special case of a more general distinction, between content-attribution at the level of the animal and content-attribution at the level of its internal machinery. At the level of internal machinery, it is useful to talk of sensory systems as information-processing devices; but for the animal its sensory systems are modes of openness to features of its environment. Information-processing characterizations of the internal machinery figure in explanations of how it can be that animals are in touch with their environments. The 'as if' content that is usefully deployed at the lower level helps make intelligible the genuine content

[5] See *Intentionality* (Cambridge UP, 1983), p. 230.

© The Editors of *The Philosophical Quarterly*, 1994.

that appears at the higher level by way of 'enabling' explanations, not as somehow constituting that content. Since there is no getting around the fact that the internal machinery is really only a syntactic engine, the attempt to see a constitutive relation between the lower and the upper levels undermines our hold on the fact that animals are semantic engines.

It will not have escaped notice that my descriptions of what sensory systems are for an animal that possesses them are Gibsonian.[6] A proper understanding of the relation between the two levels should help us to see through some cross purposes in a familiar style of cognitive-scientific response to Gibson's claims.

Gibson himself sometimes seems to deny that the idea of processing information has any role in characterizations of the operations of sensory systems. This is fully intelligible, given the fundamental point that he wants to make: that perceiving (something that animals do) is not processing information, but simply taking it in. In fact that claim is, as should by now be clear, fully compatible with recognizing that it can be useful to characterize sensory systems, *not* from the animal's point of view, as information-processing devices. The claim gives the framework within which such characterizations have their explanatory point.

Some cognitive-scientific critics of Gibson, not equipped with the distinction of levels, read Gibsonian descriptions of the sensory systems as if they were meant to serve the intellectual function that their own theories serve. It is not surprising if that makes the Gibsonian descriptions seem idiotic. David Marr, for instance, in his brilliant and path-breaking information-processing approach to vision, cites Gibson only as a half-baked precursor. When Gibson resoundingly, and rightly, denies that perceiving is processing information, Marr can hear only a reflection of failure to understand what information-processing is. In this framework, Gibson's positive remarks about the sensory systems can indicate at best a massive under-estimation of how difficult the information-processing task, the task of extracting information about environmental invariants from 'the flowing array of ambient energy', is.[7] And really this is an understatement: if we read Gibson as attempting to say something at

[6] See J. J. Gibson, *The Senses Considered as Perceptual Systems* (London: Allen & Unwin, 1968).

[7] See Marr's *Vision* (New York: W. H. Freeman, 1982), pp. 29–30. At p. 3 of *The Senses Considered as Perceptual Systems*, Gibson writes: 'The unanswered question of sense perception is how an observer, animal or human, can obtain constant perceptions in everyday life on the basis of ... continually changing sensations.' Marr misreads this. He takes Gibson to be acknowledging the urgency of that question, whereas Gibson is describing how the problems look in the approach that he is going to reject. See p. 2 ('The seemingly paradoxical assertion will be made that perception is not based on sensation'); p. 320 ('The puzzle of constant perception despite varying sensations disappears').

© The Editors of *The Philosophical Quarterly*, 1994.

the lower of the two levels, then given what is obvious about the physical impingements to which, say, visual systems are restricted, the idea of directly collecting information about 'environmental invariants' through vision looks like an appeal to magic. As Dennett says, quite correctly given this reading of Gibson, and completely missing Gibson's point, Gibson represents the visual system as 'a hunk of wonder tissue'.[8]

9. The distinction of levels equips us to see that there is nothing unscientific, no mere know-nothing refusal to acknowledge the rich promise of cognitive science, in denying that dealings with content on the part of perceiving animals should be equated with computationally described goings-on in their interiors. Moving to personal dealings with content, such as the conscious perceptual experiences of adult human beings, makes all kinds of differences. But there is no reason to suppose that it makes any difference on this point: our dealings with content, in our consciously enjoyed perceptual experience, are no more a matter of access to our own interiors than a frog's dealings with content are.

Dennett offers a picture of our internal functional organization, in which perceptual systems process bare data into comparatively rich information about the environment, and their products (at various levels of processing) are stored in a special short-term buffer memory ('M'), which feeds into a system controlling speech. (There is more complexity besides: see p. 155.) About this picture, he floats (p. 165) 'the bold hypothesis that you are a realization of this flow chart, and that it is in virtue of this fact that it seems – to us and to you – that there is something it is like to be you' (that is, that you are conscious). As far as anything I have said goes, this may well be right. The important point is that if it is right, it is right as a piece of cognitive science, with 'in virtue of' receiving a causal reading, not a constitutive one. The suggestion has the same shape as one to the effect that it is in virtue of possessing perceptual equipment that admits of such and such an information-processing characterization that an animal can be in touch with such and such features of its environment.

One might put this by saying that consciousness itself escapes Dennett's cognitivistic net: he offers what may be an enabling explanation of consciousness, but not a constitutive one. In one sense, this leaves us without an account of consciousness. We lack an account of what it is, even if we have an account of what enables it to be present. It would be a mistake to think this makes consciousness a metaphysical embarrassment: as if,

[8] 'Cognitive Wheels', in C. Hookway (ed.), *Minds, Machines and Evolution* (Cambridge UP, 1984), pp. 129–51, at pp. 149–50 (n. 21).

© The Editors of *The Philosophical Quarterly*, 1994.

in denying that consciousness is a matter of configurations in the satisfyingly material medium of the nervous system, we were committed to regarding it as a matter of configurations in an immaterial medium instead. As I have insisted, there is nothing inherently mysterious in a frog's being in touch with its environment (of course not in a 'personal' way); we take that idea in our stride, in the context of thinking about how the frog's life fits into the environment. There is no difficulty in principle, although all kinds of differences must be acknowledged, about extending this comprehension to our own case: our personal-level dealings with content are intelligible in the context of our distinctively human life.

One striking advantage to be derived from rejecting the idea that conscious perceptual experience is a special kind of access to content that is in the first instance sub-personal, i.e., to the content of events or states in our interiors, is that it enables us to repossess the phenomenology of perception. I have already said something about a phenomenological misrepresentation that Dennett is led to (see §2 above). Let me give another example: discussing the richness of experience, Dennett writes (p. 170):

> One experiences more at any time than one wants to say then. What fills the 'periphery', adds detail to one's 'percepts', inhabits 'fringe consciousness', is, as phenomenologists have insisted, *there*. Where? In M [the special buffer memory]. No more mysterious process of presentation or apprehension of inhabitants of phenomenal space is needed.

Certainly no such thing is needed, but Dennett's alternative answer, 'In M', is surely quite wrong as phenomenology. The phenomenologically right answer – which, once we have the status of sub-personal theorizing straight, we can recognize as the right answer, period – is: in the part of the world (ordinary objective space, not 'phenomenal space') that lies open to view.

10. At one point (pp. 160–1) Dennett remarks, in passing, that his construction has a Kantian flavour: his flow chart diagrams how intuitions (the sub-personal, and hence certainly not conceptual, content of states or events that result from perceptual data-processing) are 'knitted together' with concepts (which figure in the articulable shape that contents acquire when they are made available to the system that controls speech). I think the real lesson of Dennett's paper is this: a dualism of intuitions and concepts cannot be made safe by simply removing it from the sphere of the transcendental – by assigning the task of fitting intuitions and concepts together to something empirical, whether it is the empirical self (as in §3 above) or, as here, an empirically

© The Editors of *The Philosophical Quarterly*, 1994.

postulated internal apparatus. In either case, just because the 'synthesis' is not transcendental (nor therefore something that simply disappears if we attempt a domesticated formulation of what we can learn from Kant), there is an epistemological come-uppance. The great beauty of Dennett's paper is how rigorously he lets his theory control his phenomenological claims; their failure of fit can now emerge as a fault line along which the whole structure must break apart.

We must see our way to not needing to give an account of how concepts and intuitions are brought into alignment. Another way of saying the same thing would be that we cannot make use of the notion of an interface between mind (which inhabits the space of concepts) and world, where the world presents the mind with non-conceptual items for it to work into conceptual shape. Or (yet another formulation) we cannot make sense of the mind as a 'black box' in the world. (Obviously we can make sense of an interface between nervous systems and the world outside them.)[9] This paper has been about perceptual connections between mind and world, but the point has obvious reverberations for how we think of action too: if seen objects (say) are not on the far side of an 'input' interface between mind and world, there is, to say the least, no point in trying to represent objects acted on – which may, of course, be the same objects – as lying on the far side of an 'output' interface.

University of Pittsburgh

[9] Searle, in *Intentionality*, has the 'black box' picture of the mind firmly in place (we *are* brains in vats: see §7 above). This means that he cannot capitalize on what appears as an isolated insight, that perceptual experience is presentation, not representation (see p. 46). He tries to do better than Dennett's 'premonition' idea by insisting that the aetiology of an experience enters into its conditions of satisfaction. But with the interface in place, it is merely mysterious how it can do so.

© The Editors of *The Philosophical Quarterly*, 1994.