

Lucian Wischik

UNCONSCIOUS CARTESIAN DEMON

Abstract. A picture of the mind is proposed in which there are large unconscious depths, and in which we are being misled into thinking that the mind is more rational than is actually the case. Within this framework, introspection would lead to the following situation: there would be a multiplicity of theories, a systematic over-valuation of rationality, and a lack of insight into the functioning of the unconscious. Dennett's theory of vision is compared with a neuroscientific theory, and it is indicated how Dennett was led to suspect conclusions by his introspective methods. It is concluded that introspection alone is an unreliable method in the philosophy of the mind.

History and Philosophy of Science Second M.Phil. Essay:
General Principles in the Philosophy of Science
Queens' College, Cambridge
February 1997

4970 words approx.
Supervisor: Dr. P. Lipton

1. INTRODUCTION

Suppose that there are unconscious depths to the mind; and suppose that the mind is being consistently, systematically and subtly misled not by some malevolent dæmon but by its own unconscious. We present evidence that this picture of the mind is true, and that introspective approaches to the study of the mind therefore tend to yield false results.

The first part of this essay provides an overview of the argument that introspection has the potential to be flawed, given the assumption of the unconscious. The rest of the essay presents in detail a small part of the argument. We explain why we consider Dennett to be an introspectionist, and examine aspects of his theory of vision from Conscious Explained¹. We then present a parts of an alternative theory of vision, which integrates the picture of the unconscious with neuroscientific evidence. The two theories are compared. The fact that Dennett's introspective method has in some cases led him to rule out the alternative theory without sufficient grounds suggests that his method is flawed; we attempt to trace precisely how his flawed conclusions can be seen to follow from the shortcomings of his method.

It may seem that the picture of mind with unconscious has been designed specifically to make introspection as unreliable as possible, in much the same way as Descartes postulated his unbelievable dæmon to cast doubt upon sensory inputs. Just as Descartes concluded that it was epistemologically unsound to rely upon sense data because of the possibility of the dæmon, so we will conclude that it is philosophically unsound to rely upon introspection because of the possibility of this unconscious.

However, the neuroscientific evidence will effectively deal a second, stronger blow to introspection. As mentioned above, it will be used as the basis for a theory against which Dennett's theory can be compared and found wanting. But additionally it will suggest that not only is introspection potentially unsound because it may be misled, but also that introspection is actually unsound because it does get misled.

2. OVERVIEW OF ARGUMENT AGAINST INTROSPECTION

The anti-introspectionist project is large. We present in this section and the next an overview of the project, and indicate the particular small part of it that will be argued in this essay.

The sort of unconscious we suppose is illustrated in Figure 1 as a pyramid. Information comes in through the senses at the bottom and is progressively refined and re-expressed in higher-level representations towards the top. We suppose that only a fraction of the higher-level representations, at the very peak of the pyramid, are the

¹References consisting solely of a page number refer to the book Consciousness Explained [Dennett 1991].

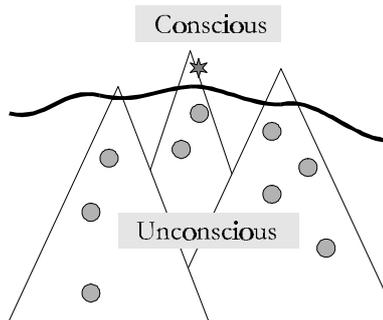


Figure 1. The unconscious mind as a pyramid of representations. Only those things above the line are taken to be the subjects of conscious experience.

subjects of conscious experience; the rest of the mind is unconscious² information-processing. There is no reason to suppose that the high-level representations should have an especially faithful relation to the sensory inputs except in as much as is useful, in an evolutionary sense; and there is no reason to suppose that high-level representations should themselves contain an accurate picture of the non-conscious parts of the mind.

We will further suppose that the human mind has an instinctive tendency to post-hoc rationalization: that is, after some sequence of events has been observed, the mind will typically look back and put a 'spin' on them that ascribes a rational, logical story where perhaps none was present. In particular, when some arational, unconscious mental process has taken place, the mind will instinctively look back over the process and falsely decide that the process was in fact rational and conscious.

In this essay, we term introspective' those investigations of the mind that take as their input data only the high-level representations in the mind. For instance when someone actively examines the state of their own mind they are being introspective. This is because, as stated above, the only parts of their mind that are taken to be subject to conscious experience are the high-level representations. Essentially, an introspective investigation is one that can be carried out from an armchair (or within Descartes' warm room).

We could expect three things from an introspective investigation of the mind. They will be detailed below but in summary are as follows: there will be many different introspectionist theories of the mind with little to set one above another; post-hoc rationalization will lead introspectionists to a distorted picture of the mind in which rationality plays too great a part; and introspectionists will either neglect the importance of the unconscious or will not be able to study it adequately.

It could no doubt be argued that introspectionists are not naïve, and that they could carefully avoid the pitfalls of their method and adequately piece together a picture of the unconscious from the scant clues available. But we suggest that even a careful introspectionist is likely to be misled. The justification of this claim, by example, forms the main part of this essay.

The example we have chosen is Dennett, in Consciousness Explained, for the following reasons. First, Dennett is a reductionist and a materialist: this means that his theories will be to some extent commensurable with the neuroscientific picture we will outline. Second, a significant part of the book is devoted to an explanation of his method—heterophenomenology—which, we will suggest in the following section, is essentially *third-*

²It should be noted that the term 'unconscious' in this essay refers simply to the automatic non-conscious processing of information in the mind. This is quite different from Freud's notion of the unconscious as a fully-functional manipulating homunculus.

person introspectionism. It is useful that he makes his method explicit, since this essay is an argument against the method.

Some may claim that Dennett 'denies consciousness', ascribing far less significance to consciousness than is justified. Given this claim it may be suggested that, in our attempt to stress the importance of the unconscious at the expense of the conscious, we have made a perverse choice of target. Such a suggestion, however, makes a category error. In denying consciousness Dennett is addressing the first of the two central questions in the philosophy of the mind: he claims that there is not something special about subjectivity or qualia, or at least that our intuitions about it are wrong. With regard to the second central question in the philosophy of mind—that of the functioning of

Dennett does accept the existence of conscious parts of the mind (whatever 'consciousness' may be). In this essay we will consider only this second question; and we will suggest that Dennett's methods have led him, in constructing a theory of the functioning of the mind, to focus too much on the conscious at the expense of the unconscious.

3. MISTAKES OF INTROSPECTION

In the previous section it was indicated that introspection might have three problems: a multiplicity of theories; a distorted evaluation of the significance of rationality; and an unclear picture of the unconscious. In this section we explain and expand upon these problems.

First, recall that only the high-level representations at the very tip of the pyramid of the mind were taken to be the subjects of conscious experience. Although these relate directly to the unconscious immediately beneath them, they relate only incidentally to the other high-level representations adjacent to them. The resulting material with which the introspectionist must work will be a confusing jumble of partially related phenomena, upon which it is difficult to impose order. Introspectionists will try to discover patterns in the data, but the relationships between different high-level representations will be extremely complex. One person may focus on a subset of the data and conclude that the mind works in terms of beliefs and desires; another may focus upon a different subset and conclude that the fundamental unit is perception; another may focus upon yet another subset of the data and conclude that propositional, logical beliefs and deduction are the best way to look at the mind. All theories will adequately explain part of the data, but none will explain all; and it will be impossible to find objective grounds for setting one theory above another. We term this problem the multiplicity of theories.

Second, we supposed that there is a natural human tendency for post-hoc rationalization, for perceiving logical stories or patterns even in cases where such patterns are not present. Since there are often regularities in nature, this tendency to spot patterns will generally be useful. But if there are mental processes that are arational and unconscious, then the mind will look at them and instinctively and falsely decide that the processes were in fact rational and logical; and, since we associate logic with conscious reason, the natural conclusion will be that the processes were conscious. Post-hoc rationalization will therefore systematically fool introspectionists into ascribing more rationality, logic and consciousness to the mind than is actually the case.

The third problem with introspection is its lack of insight into the unconscious. There is no reason to suppose that any high-level representation should contain an accurate picture of the functioning of the unconscious parts of the mind: the

unconscious is not visible to the conscious. Although it might be argued that our consciousness necessarily cannot be fooled about what goes on in the conscious parts of the mind (simply because we are by definition aware of them), the same certainty does not hold true for the unconscious.

Introspectionists who do not appreciate the significance of the unconscious will have a tendency to construct theories of the mind which focus too much upon the conscious at the expense of the unconscious. This is an easy mistake to fall into: "We are aware of the mind being composed of conscious parts; therefore the mind is composed of conscious parts." And, since we assumed that only the representations at the very tip of the pyramid are the subjects of conscious experience, introspectionists who do appreciate the unconscious will have a difficult job trying to piece together a picture of how it works from the scant clues available at the conscious level.

In the previous two sections we have given an overview of the argument against introspection and listed three shortcomings that may be expected of introspective theories of the mind. In the next sections we present Dennett as an introspectionist, and point out where he exhibits the above shortcomings.

4. HETEROPHENOMENOLOGY, AND MULTIPLE DRAFTS

Phenomenology is the study of mental experiences, especially conscious mental experiences. When we experience heat, cold, day-dreams or anger, these are phenomena. Relating phenomena to the pyramid picture of the mind outlined above, phenomena may be identified with those high-level representations in the mind that are the subject of conscious experience.

There are two central problems with phenomenology. First, it may be that phenomenology is not universal: when we question another person about their phenomena, we have no guarantee that their phenomena are in any way similar to our own— even that they experience phenomena at all. For instance, if a friend sees a red piece of paper, it might be that his subjective experience of red is actually more like our own subjective experience of green; or it may even be that he simply has no subjective experience of the colour. The second problem with phenomenology is the very *subjectivity* of subjective experiences. Phenomena are essentially connected with a single point of view. Since science requires its subject material to be observed consistently from many views, it would seem difficult to study phenomena in any objective, scientific manner.

Dennett presents a new name for a method, heterophenomenology, which circumvents these two problems. This method is distinguished by a platitude: we listen to a person's description of their phenomena and we accept what they say (if only for the purpose of discussion). Essentially, we discuss the phenomenological world of a particular subject in much the same way as we can discuss the imaginary world expressed in a work of fiction. The heterophenomenological method neither challenges nor accepts as entirely true the assertions of subjects, but rather maintains a constructive and sympathetic neutrality, in the hope of compiling a definitive description of the world according to the subjects" [p. 83]. We may think of heterophenomenology as third-person introspection, since it uses as data the content of the conscious experience of others.

Heterophenomenology circumvents the problem of universality because it simply does not matter if the friend's experience of a colour is in some way different from our own: we agree that the friend claims to have experienced red, and we then talk about

the conscious experience which the friend claims to have had. It circumvents the problem of objectivity because the written transcript of a subject's description is a concrete fact that can be objectively studied.

It might seem bizarre that, in attempting to come up with an explanation of how the mind actually works, Dennett should use a method that treats the accounts of subjects as fiction rather than fact. One response to this charge might be to argue that there is some way in which the subject's mind works and that, even though his account is a fiction, the fact of its generation does shed some clues as to his mind's actual functioning. This response is similar to the attitude of the non-naïve introspectionist, mentioned in section two, who pieces together clues as to the mind's functioning from the scant clues available at the conscious level.

Dennett deploys a different response to the charge. He tacitly assumes that everyone has essentially the same sort of content of thought and behaviour, so that the results from one subject can reasonably be applied to others. The fiction that he describes is not the content of the mental experience, but rather the subjectivity of the experience—the qualia. Although qualia are labeled as fiction, this is not to say that Dennett is actually denying that subjective experiences exist (at least, not in the book Consciousness Explained); rather, he maintains a constructive and sympathetic neutrality with regards to their existence. When he uses the word 'consciousness,' he is talking about the content of conscious experience, and he sets aside the issue of its subjectivity.

Heterophenomenology is the first of the two foundations of Consciousness Explained; the second is the rejection of the Cartesian Theater in favour of the Multiple Drafts model of the mind. It takes Dennett the entire book to explain why the Multiple Drafts model of the mind, as a theory, has greater explanatory power and elegance than the Cartesian Theater; we limit ourselves here to a brief outline.

The Cartesian Theater would be some central region in the brain which does the actual business of consciousness. Data would come in from the peripheral systems of the brain to be presented in a sound-and-light show to a solitary but powerful audience: the Ego, the Central Executive, the Master Judge, the audience of the Theater" [p. 227]. Decisions would be made here, and then passed down to be carried out by the rest of the brain. It may be that few people explicitly hold this idea, of an homunculus inside the mind who does the real work; but, according to Dennett, the image has a pervasive influence which misdirects most philosophers of the mind.

In the alternative Multiple Drafts model, the mind is seen as a pandemonium of competing and co-operating agents. There is no central consciousness, no central time [Minsky 1985, pp. 107-108]. Information moves around the brain and is processed by various mechanisms in various regions [p. 144]. What we experience is the product of many processes [p. 112]. Effectively, consciousness is done by the entire brain rather than just a small area. If a particular part of the brain makes some discrimination, then this local discrimination is all there is to it: the message does not have to be passed into a different part of the brain to actually become conscious [pp. 113, 292].

The two bases, of Heterophenomenology and Multiple Drafts, are the foundation from which Dennett builds his theories of the mind.

5. NEUROSCIENTIFIC VISION

In this section we present some neuroscientific evidence concerning vision. In the next sections Dennett's explanation of vision is presented and compared.

We propose that the vision has a hierarchical structure, and that different areas in the brain correspond to different levels. These areas of the brain are automatic feature-detectors: thus, the mental hierarchy is one of representations. A representation might informally be defined as a collection of neurons which discriminate some feature, such that the fact of this discrimination can be readily detected by some other collection of neurons. The only job of a representation, then, is to be registered by other parts of the brain.

From the anatomical perspective there is good evidence for hierarchy within the visual systems in the brain [Essen & Maunsell 1983]. In the macaque monkey, around a dozen brain areas have been identified that deal largely or exclusively with visual information. The connections between these areas can be found by staining neurons in one area with a dye, and determining to which other areas the dye spreads. Connections can also be assigned a direction according to precisely how the neurons are connected. A well-defined hierarchy can be inferred from these connections (See Figure 2). In other words, the areas can be unambiguously assigned to specific ranks in relation to one another. (This fact is perhaps surprising: it would be easy to conceive of other more random connection patterns that had no hierarchy).

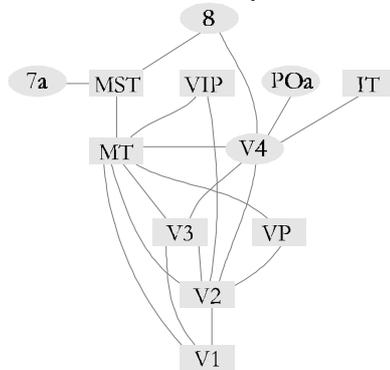


Figure 2. The hierarchy of visual areas in the macaque monkey. Each code indicates a particular anatomical area in the brain. From [Essen & Maunsell 1983].

From an information-theoretic perspective there is also evidence for a hierarchy, in terms of the processing and reorganization of information. A bit, or binary digit, is the fundamental unit of information. The optic nerve carries 10^8 bits (equivalent to the entire content of all twelve volumes of Burton's translation of the *Arabian Nights*) every second. However, it is estimated that the information intake for decision-making is of the order of tens of bits (just a few words) per second. This strongly suggests that visual information gets progressively simplified.

We have suggested that the hierarchy in the visual systems is a hierarchy of representations. One class of representations, close to the bottom of the hierarchy, are *retinotopic maps* where the representation has a similar sort of layout to the actual picture falling on the retina. Some retinotopic maps may mark out boundaries of intensity in the image, or re-express the image in different geometries (Figure 3 illustrates the transformation of an image from a log-polar coordinate system into a Cartesian coordinate system). Retinotopic maps are common in the brain: in the owl monkey, at least sixteen have so far been discovered [Allman 1971].

At higher levels in the hierarchy, representations become more abstract. They may be of three-dimensional geometrical shapes such as cylinders; and there may be levels at which shapes are represented in viewpoint-specific terms and higher levels at

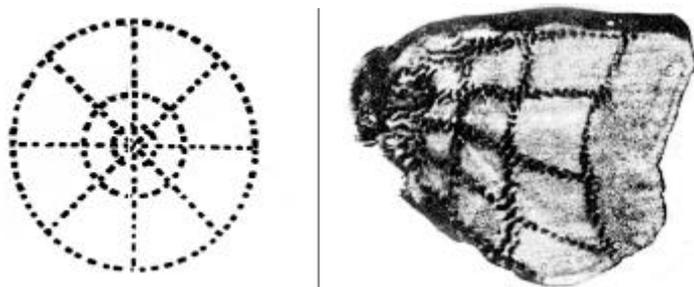


Figure 3. Slice of cat retina demonstrating a retinotopic map. A cat is injected with a radioactive glucose trace and its eyes exposed to a circular ring-and-ray pattern. The cells that fire in response to the image absorb glucose and so become stained. Shortly after, the cat brain is dissected and photographed. [Totell et al 1982]

which they are represented in viewpoint-independent terms [Marr 1982 pp. 307]. High level representations may also be of properties such as colour or texture, or of entities such as 'chair' or 'table'.

The remarkable property of the visual hierarchy is that more information flows from the higher levels back down towards the retina than flows in the expected upwards direction [Essen & Maunsell 1983]. From an overall perspective, this accords with the standard experience that knowledge is important in vision. (For instance, it is difficult to recognize an object away from its canonical perspective [Humphreys & Quinlan 1987 p. 79]. This may be readily tested by trying to recognize upside-down faces).

It seems likely that the visual system uses an ingenious feedback strategy. At each level in the visual hierarchy, the brain renders or projects how this high-level representation would appear at the next lower level; any differences between what is expected and what is experienced would be noted, and passed up as feedback so that the higher-level representation may be corrected (see Figure 4). Within this scheme, the figuring-out of an image is an interactive process between levels.

It would take time for this interactive process to settle upon a satisfactory high-level representation. This is born out by experiment. Subjects were shown various shapes: when the shape was shown for only 100ms, then the subjects' memory was only of its 2D projection; when the shape was shown for longer, the subjects remembered it as a 3D structure. The data suggests that successively more abstract descriptions can be generated during visual object processing, progressing from a retinally coded description, to a non-retinotopic viewpoint-specific description, to a yet more abstract description which can serve the recognition of an object across different viewing distances and angles. These different descriptions appear to take successively longer times to generate" [Humphreys & Quinlan 1987 p. 56].

We are clearly not aware of performing such tasks as complex log-polar transformations of what we see. This accords with the notion of the unconscious suggested in the introduction, where a large part of the information-processing was unconscious and only the very abstract representations at the top of the hierarchy are the

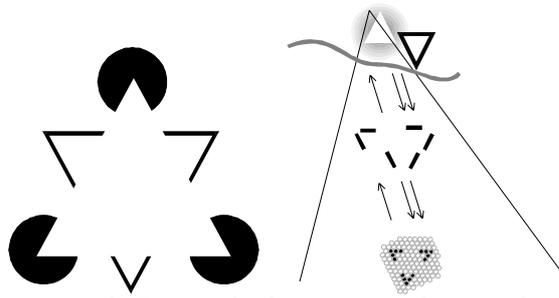


Figure 4. The Kanizsa triangle. It is very hard not to see a white triangle, even though we tell ourselves that none is really there! At high levels of representations, an abstract description consists of 'white triangle superimposed on black line triangle'. This high-level image is projected down a layer into lines; the middle layer reports that, yes, the high-level abstraction is indeed consistent with what actually falls on the retina. Effectively we do see a white triangle, because a white triangle is the subject of our conscious experience at the top of the pyramid.

subjects of conscious experience. Higher-level representations change slowly; and the high-level representation maintains a constant conceptual representation of objects and the environment, independent of inconsistencies such as eye-movement or changing perspectives. We are conscious only of these slowly changing high-level representations: this is why visual consciousness is continuous [Minsky 1985 pp. 256-257].

The key features of this neuroscientific theory of vision are as follows: there is a hierarchy in vision; low-level processing is not the subject of conscious experience; feedback is centrally important; high-level visual perception is slow; visual consciousness is continuous. We have firstly suggested that there is a significant unconscious part to visual processing: therefore, if introspection is suspect given the supposition of an unconscious, then introspection will be suspect with regards to vision. Second, we have presented a theory of vision against which we can measure Dennett's theory.

6. DENNETT'S VISION

In this section we outline the part of Dennett's explanation of vision that relates to filling

The retina presents an untidy image to the brain: there are blood vessels covering its surface; the optic nerve creates a blind spot, subtending approximately six degrees, in which no retinal image is possible; and acuity of vision and density of photoreceptors degrade away from the center of vision [Abramov et al 1991]. Heterophenomenologically, no subjects report experiencing a blind spot in the course of normal vision: how should we explain this?

A natural explanation might be that the brain somehow guesses what should be in the blind spot from its surroundings, fills the image in, and then projects this completed picture onto some further cinema-screen which is observed by the central consciousness. Dennett disagrees strongly: [this is a] dead giveaway of vestigial Cartesian materialism" [p. 344]; it would be a waste of paint" [p. 128].

Instead, Dennett suggests that subjects do not report any discontinuities at their blind spot because there is no part of their brain looking for discontinuities. There are no homunculi who are supposed to care about information arising from the part of the visual field covered by the blind spot, so when nothing arrives, there is no one to complain" [p. 357]. Essentially, things only have to be discriminated once, and once

discriminated they are potentially able to become conscious. Nothing gets discriminated about the blind spot, so nothing about it becomes conscious; lack of consciousness is quite different from consciousness of a blank spot.

Dennett concludes that consciousness is discontinuous, but that we are not aware of the fact. 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. And this leads even very sophisticated thinkers to make crashing mistakes, perfectly epitomized by Edelman: 'One of the most striking features of consciousness is its continuity.' This is utterly wrong. One of the most striking features of consciousness is its discontinuity. the discontinuity of consciousness is striking because of the apparent continuity of consciousness. as Minsky puts it, 'Nothing can seem jerky except what is represented as jerky.' [p. 356]

It is a little disingenuous of Dennett to quote Minsky [Minsky 1985 p. 257] at this point, for Minsky had been arguing that consciousness is continuous and had presented a structure of vision similar to the neuroscientific picture outlined in the previous section. The striking features of that picture were that the brain does provide something—a representation at a higher level—and that the consciousness is continuous, because the high-level conscious representations are abstract and change only slowly.

Given the evidence presented in the previous section, Dennett's rejection of filling-in is at best unwarranted and at worst incorrect. We presume that his rejection of the hierarchy of vision with constant images at the top, stems from a general rejection of anatomical hierarchies in the brain. This will be discussed in the next section.

7. INTROSPECTION VERSUS ANATOMICAL HIERARCHY

Much of Dennett's argument is an attempt to discredit the notion of the Cartesian Theater. He writes that the idea of 'special center in the brain is the most tenacious bad idea bedeviling our attempts to think about consciousness' [p. 108] and that, even though most philosophers superficially reject the notion of an intelligent homunculus watching cinema screens, their thoughts remain influenced by it. We argue in this section that, because of Dennett's use of the flawed introspective method, his (valid) rejection of an internal homunculus leads him wrongly to dismiss anatomical hierarchy.

As was suggested in the introduction, introspection suffers from post-hoc rationalization: in reporting their experiences, introspectionists will tend to ascribe far more to a conscious rational part of their mind than is actually justified. That is, introspectionists will be misled in their account of their own mental states; and they will be misled into positing a powerful, central, rational core. This 'Cartesian experience' becomes for Dennett a significant problem because as a heterophenomenologist he must accept—only for the sake of discussion—the strong introspective experience of a central rational consciousness [p. 107].

The natural way to solve the problem of the mind is to break it down into a hierarchy [p. 485]. There are two sorts that can be proposed: a virtual functional hierarchy, and physical anatomical hierarchy.

In a virtual hierarchy, it is only the parts at the very lowest level that physically exist: when we talk about higher levels, we are using a verbal shortcut for describing a collection of lower-level entities which have collective behaviours [pp. 39,262]. For instance, a 'year-group' is a higher-level virtual entity that does not physically exist, but is a convenient way of talking about the lower-level collection of students.

In a physical anatomical hierarchy, every different level has a separate physical existence. For instance, when we talk about a Head of Department, lecturers and

graduate students, we are referring to entities that exist individually: it is not the case that the Head of Department is merely a convenient way of talking about the graduate students.

As discussed, introspectionists tend ascribe far more significance to a central, rational, conscious core. When an introspectionist draws a hierarchy, it will have this significant consciousness at the very peak. In an anatomical hierarchy, this peak would have to be an area of the brain that does consciousness: a fully-fledged internal homunculus. In a virtual functional hierarchy, the peak of the pyramid is merely a convenient way of talking about the brain as a whole: so placing consciousness at the peak merely says that the brain as a whole exhibits consciousness. The introspectionist who rejects the notion of an internal homunculus is thus forced to reject anatomical hierarchies and espouse functional hierarchies.

This may be seen in Dennett's treatment [pp. 227-242]. He introduces a proposed anatomical hierarchy for speech generation (Figure 5). This anatomical hierarchy indeed exhibits a conscious peak: In the upper left-hand corner a functionary who looks suspiciously like the Central Meaner makes his appearance in the guise of the Conceptualizer, armed with lots of world knowledge, plans, and communicative intentions" [p. 233]. Dennett rejects this sort of bureaucratic hierarchy of agents and proposes instead a pandemonium model, in which many equal agents compete and cooperate. In the case of language generation, Dennett suggests that words, phrases and grammatical structures are the equal agents, competing to get themselves spoken. This structure would be a virtual functional hierarchy, in that the overall functioning of a group of such agents might be considered equivalent to the functioning of the speech subsystem.

It was presumably his rejection of anatomical hierarchy that led Dennett to his suspect conclusions about vision: that consciousness is continuous; and that no filling in occurs, even at a higher level. If instead we approach the problem of vision without being misled by introspection, without overvaluing conscious, then we can construct a theory that is in far better agreement with the neuroscientific evidence.

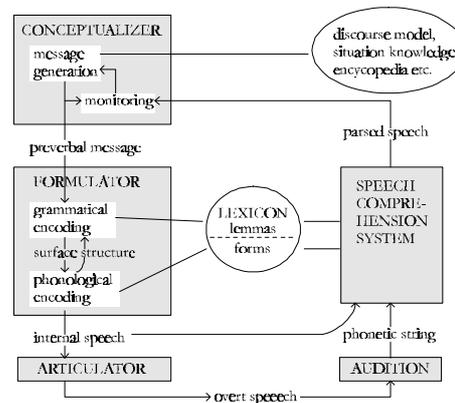


Figure 5. A proposed anatomical hierarchy for speech. Each area corresponds to a particular part of the brain. Details need not concern us. Dennett rejects the picture, because it would ascribe too much conscious rationality to the 'conceptualizer' at the top left.

8. SUMMARY

We started by supposing a picture of the mind where there is a large unconscious, and where we are systematically misled by post-hoc rationalization. Within this supposition it was suggested that the introspective method is likely to produce a multiplicity of theories with little to set one above the other, a systematic overvaluation of consciousness and rationality, and sketchy pictures of the unconscious.

We suggested that, because the introspective method overvalues consciousness, any proposed hierarchies of the mind would tend to have a large and significant consciousness at the peak. Two sorts of hierarchy were presented: an anatomical hierarchy in which different levels correspond to different brain regions, and a virtual functional hierarchy in which different levels are merely different ways of talking about the same regions.

In the case of an anatomical hierarchy, that conscious peak would correspond to an area of the brain that does consciousness: an internal homunculus. If we reject the internal homunculus, then we would be forced to reject anatomical hierarchies. This invalid rejection is thus a consequence of introspection.

As an example, we discussed Dennett's theory of vision: its conclusions (such as a denial of filling-in and of continuity in consciousness) were a consequence primarily of his rejection of anatomical hierarchy. Because Dennett's methods led him to these conclusions without good reason, we concluded that Dennett's methods were unreliable. This approach was similar in structure to that of Descartes: if there is a malevolent dæmon, then sensory input is unreliable; if there is an unconscious, then introspection is unreliable.

The neuroscientific evidence also led us to a stronger form of the conclusion: not merely that introspection has the potential to be unreliable because of the possibility of the unconscious, but also that introspection is actually unreliable because there is an unconscious.

Thus, introspection alone is unreliable. We suggest that neuroscience should be used in conjunction with conventional philosophical approaches to build a satisfactory and comprehensive theory of the mind.

BIBLIOGRAPHY

- Abramov I., Gordon J., Chan H. 1991. *Color appearance in the peripheral retina: effects of stimulus size*. J. Opt. Soc. Am. 8.2:404-415.
- Allman 1971. *Representation of the visual field in striate and adjoining cortex in the owl monkey*. Brain Research 35:89-106.
- Dennett D. 1991. *Consciousness Explained*. Penguin Books: London.
- Essen D., Maunsell J. 1983. *Hierarchical organization and functional streams in the visual cortex*. Trends in Neurosciences September.
- Humphreys G., Quinlan P. 1987. *Normal and pathological processes in visual object constancy*. In Visual Object Processing ed. Humphreys G., Riddoch M.
- Marr D. 1982. *Vision*. Freeman & Co.: San Francisco.
- Minsky M. 1985. *The Society of Mind*. Heinemann: London.
- Tootell R., Silverman M., Switkes E., DeValois R. 1982. *Deoxyglucose analysis of retinotopic organization in primate striate cortex*. Science 218:902-904.