

Chapter Four: A Method for Phenomenology

This is a chapter which requires careful reading, because it is here where so many readers have misunderstood what Dennett's program is all about, attributing to his theory a denial of consciousness, for instance.

1. FIRST PERSON PLURAL

Tradition has had it that we can rely upon our introspection to tell us what exactly is going on in our consciousness, but that tradition was incorrect, as we saw earlier. What has been going on? Dennett thinks he knows: "what we are fooling ourselves about is the idea that the activity of 'introspection' is ever a matter of just 'looking and seeing.' I suspect that when we claim to be just using our powers of inner observation, we are always actually engaging in a sort of impromptu theorizing—and we are remarkably gullible theorizers, precisely because there is so little to observe and so much to pontificate about without fear of contradiction."(pp 67-68)

The possibility that this could be what happens thus undercuts any hope that "pure" phenomenology truly gives us a neutral story about what goes on inside us.

2. THE THIRD PERSON PERSPECTIVE

Recognizing this problem, the behaviorism movement emphasized a change of focus in the science of mind—from relying on introspection, to only counting objectively measurable data subject to third person verification. Reports by subjects about what is going on in their minds cannot be so measured or verified, so must be approached with some skepticism.

Now here Dennett makes a very important point about what he is and isn't saying:

This methodological scruple, which is the ruling principle of all experimental psychology and neuroscience today (not just 'behaviorist' research), has too often been elevated into one or another ideological principle, such as:

Mental events don't exist. (Period!—this has been well called 'barefoot behaviorism.')

Mental events exist, but they have no effects whatever, so science can't study them (epiphenomenalism—see chapter 12, section 5).

Mental events exist, and have effects, but those effects can't be studied by science, which will have to content itself with theories of the 'peripheral' or 'lower' effects and processes in the brain. (This view is quite common among neuroscientists, especially those who are dubious of 'theorizers.' It is actually dualism; these researchers apparently agree with Descartes that the mind is not the brain, and they are prepared to settle for having a theory of the brain alone.

These views all jump to one unwarranted conclusion or another." (pp 70-71)

So much for the view that Dennett doesn't think we're conscious—right here he distances himself from that very position, and from "barefoot behaviorism" by name. It's also worth noting once again that Dennett is using the term "dualism" to refer not just to the traditional Cartesian substance, but to any view which makes something non-physical the focal point of consciousness.

The point of this chapter is that a scientific approach to the mind must be conducted from the third person perspective, and that Dennett is going to steer clear of certain ideological excesses—described above.

3. THE METHOD OF HETERO PHENOMENOLOGY

Dennett proposes heterophenomenology as a theoretically neutral way to collect data about people's inner experiences that sticks to the standard methodologies of science. What he is trying to be neutral about is, among other things, the zombie argument. That is, this method does not assume that its subjects must necessarily be conscious—if it did, it would beg the question against zombie-philosophers and outright eliminativists, and we don't want

to do that. But if any class of entities was ever thought to be consciousness, it would have to be normal, adult human beings, and so this is the class heterophenomenology is for.

The raw data that must be used by any science of consciousness is varied—we can measure physical events in a human body and record them with all kinds of devices, we can film the experiment, and collect an audio feed as well. One thing we might have to do with our subjects, something which doesn't often matter or would never occur in other domains of science, is interpret some of their physical actions as speech acts.

This involves two steps—a) some of the physical things a subject does can be translated as text in the subject's native language (which involves converting the raw data—mere sounds—into something abstract: words); and b) those texts are further interpreted as: “not mere pronunciations or recitations but assertions, questions, answers, promises, comments, requests for clarification, out-loud musings, self-admonitions. . .we must treat the noise-emitter as an agent, indeed a rational agent, who harbors beliefs and desires and other mental states that exhibit intentionality or ‘aboutness,’ and whose actions can be explained (or predicted) on the basis of the content of these states.”(p. 76). This is what it means to take the intentional stance.

Dennett has written at length about what is going on with the intentional stance, but the most important thing to stress here is that taking the intentional stance towards an entity is utterly neutral with regards to what is actually going on inside that entity. For instance, we can justifiably take the intentional stance towards a zombie or Deep Blue, because the only justification in either case is the pragmatic or predictive pay off for doing so.

4. FICTIONAL WORLDS AND HETEROPHENOMENOLOGICAL WORLDS

Isn't the very neutrality of the heterophenomenological method regarding zombies and intelligent computers a problem?

“The fact that there is a single, coherent interpretation of a sequence of behaviors doesn't establish that the interpretation is true; it might be only as if the ‘subject’ were conscious; we risk being taken in by a zombie with no inner life at all. . .We can't be sure that the speech acts we observe express real beliefs about actual experiences; perhaps they express only apparent beliefs about nonexistent experiences.”(p.78)

The way out of this problem is to compare the task of interpreting the text of subjects' speech acts to the task of interpreting a work of fiction. By so doing, we are “canceling or postponing difficult questions about sincerity, truth, and reference.”(p.79) The text is about a heterophenomenological world the way a novel is about a fictional world. The subject has the last word about everything that is in his or her heterophenomenological world, just as Doyle had the last word about the world of Sherlock Holmes, which has exactly the same metaphysical status.

“This permits theorists to agree in detail about just what a subject's heterophenomenological world is, while offering entirely different accounts of how heterophenomenological worlds map onto events in the brain (or soul, for that matter).”(p. 81)

Note how very different this is from the claim that mental states themselves are fictitious entities like Holmes and Watson. That is not what is being claimed. What Dennett is urging us to do here is adopt a view towards the text produced in experiments that is neutral about what the text refers to. We are temporarily suspending judgment, giving heterophenomenological worlds the status of fiction, while holding out for the both the possibility that elements in that world are about real things, or that elements in that world are about some kind of illusion or other. Holmes is an element of Doyle's created world that does not refer to anything real in our world; the London he inhabits does.

5. THE DISCRETE CHARM OF THE ANTHROPOLOGIST

This neutrality is compared in section five to the neutrality of an anthropologist researching the religious beliefs of a native tribe, who believe in a forest god called “Feenoman.” We take in everything the locals say about Feenoman and write it in our journals, and this text comprises a catalog of all the things people want to say about Feenoman. Some of what they say may be contradictory, and if so, we note this as well, remaining properly agnostic about what Feenoman “really” refers to.

True believers in Feenoman may object that really, Feenoman truly exists and possesses all the powers and abilities they attribute to him. This also gets noted, but we aren't ready to pass judgment about that—yet.

6. DISCOVERING WHAT SOMEONE IS REALLY TALKING ABOUT

We can imagine a course of investigation which might lead us to believe that a character in a purported work of fiction was actually based on someone in real life. Furthermore, we could satisfy ourselves that we had discovered that person even if the author protests that she is not who he really was writing about—we might believe the author was unconscious of the relation, real as it may be.

Similarly, we can also imagine discovering a real person responsible for the various myths and stories about Feenoman. He might not possess all the attributes the Feenomanists have claimed for him—for instance, he can't fly, and uses herbs rather than miraculous powers to heal the sick. But it's easy to imagine making discoveries which would convince us he is the source of the religion, even if he only had some of the fabled attributes, so long as he had enough of the right ones.

Some of the natives might be willing to concede, after such a discovery, that they had been wrong about Feenoman, at least in part. The “die-hards” might instead retain their orthodox beliefs by imagining a supernatural, “parallel” Feenoman residing in heaven (or the astral plane, perhaps) while his flesh and blood agent lived as a man on Earth. (pp 84-85).

Back to our heterophenomenological worlds: “if we were to find real goings-on in people's brains that had enough of the ‘defining’ properties of the items that populate their heterophenomenological worlds, we could reasonably propose that we had discovered what they were really talking about—even if they initially resisted the identifications. And if we discovered that the real goings-on bore only a minor resemblance to the heterophenomenological items, we could reasonably declare that people were just mistaken in the beliefs they expressed, in spite of their sincerity. It would always be open to someone to insist—like the diehard Feenomanist—that the real phenomenological items accompanied the goings-on without being identical to them, but whether or not this claim would carry conviction is another matter.”(p. 85)

Dualists or anyone who believes in some version of “mind-stuff”, however conceived, are in exactly the position of the die-hard Feenomanist. So are believers in Chalmers' zombies and most other versions of Mysterianism, who often make the claim that “Brain events seem too different from phenomenological items to be the real referents of the beliefs we express in our introspective reports.”(p. 85) To open the reader's mind to the possibility that such identifications might be possible, we turn to Shakey the robot.

7. SHAKEY'S MENTAL IMAGES

Shakey was a robot developed during the 60's. It was essentially a box on wheels with a television eye whose computer “brain” was in another room and connected via radio. It's environment consisted of sparse rooms with boxes, ramps, and platforms. The scientists and engineers who worked with Shakey communicated to it by typing instructions in a simple semi-English: “PUSH THE BOX OFF THE PLATFORM”. (P. 86)

Dennett goes into technical details about just how Shakey processed visual images, in effect describing how its perceptual system worked, as a kind of independent “black box” or homunculi the robot needed to do its thing. He then adds another system to Shakey, based on a simple question-answering system known as SHRDLU (this is a science fiction thought experiment, not something anyone actually did).

What could Shakey's responses to the question, “How do you tell the boxes from the pyramids?” be? Dennett gives us three possibilities:

- 1) I scan each 10,000-digit-long sequence of 0's and 1's from my camera, looking for certain patterns of sequences, such as. . .blahblahblah (a very long answer if we let Shakey go into the details).
- 2) I find the light-dark boundaries and draw white lines around them in my mind's eye; then I look at the vertices; if

I find a Y vertex, for instance, I know I have a box.

3) I don't know; some things just look boxy. It just comes to me. It's by intuition.(pp.92-93)

Each level involves giving Shakey a different level of access to its own internal goings-on. We ourselves are capable only of answering like 3) when asked about how it is we do most of our perceptual tricks. But for some tasks, we can give answers that are more like 2)—when asked about how one solves puzzles of rotating visual imagery, for instance.

Let's focus on level 2 type answers, then. A monitor can be attached to Shakey and we can watch the various tasks alluded to in answer 2) as they happen (the book actually depicts several pictures of just this thing). Of course, the monitor is for our benefit, not Shakey's—you can unplug it with no impairment to the robot's performance, for the robot isn't observing anything on the monitor. Shakey does have a camera, which is where the images on the monitor come from, but there are no visual images going on within Shakey, not anywhere. There are various bits in Shakey's CPU being switched on and off as the program which processes the signals from the camera does its thing, but that's "all" that is happening.

Of course, as part of the heterophenomenological method, we treat Shakey's statements about visual images as expressions of belief—beliefs we are suspending judgment about. A more sophisticated robot, one that achieved such a level of A.I. that it simulated a real human being, might "grant that what he was doing might be only allegorically describable as image processing—though he felt overwhelmingly inclined so to describe what was happening. In this case we would be able to tell him that his metaphorical way of putting it was entirely apt."(p.93).

But there is also a "diabolical" fourth way we could program Shakey, so that the answer to the question was utterly spurious, bearing no plausible relationship to what was going on inside:

"I use my TV input to drive an internal chisel, which hews a three-dimensional shape out of a block of mental clay. Then if my homunculus can sit on it, it's a box; if he falls off, it's a pyramid.' There would be no truth-preserving interpretation of this report; Shakey would just be confabulating—making up a story without 'realizing' it.

And this possibility, in us, shows why we have to go to the round-about trouble of treating heterophenomenology as analogous to the interpretation of fiction. . . [people] don't have any way of 'seeing' (with an inner eye, presumably) the processes that govern their assertions, but that doesn't stop them from having heartfelt opinions to express."(p. 94)

Shakey has beliefs about images, but the real referents of its beliefs are not images, just something it was talking about under the "guise" of an image—like the real Feenoman behind the sometimes errant beliefs about him. (p.94). The report of a conscious experience could be understood in the same way as a report about a brain event. Or alternatively, it could be understood as an errant "confabulation," perhaps the consequence of the quirks of brain evolution.

8. THE NEUTRALITY OF HETEROPHENOMENOLOGY

Heterophenomenology must be neutral with regards to those possibilities, and others we can think of as well. That's the way to avoid begging questions and falling into intuitionist traps about phenomenology, while collecting data about what things are like from the first person perspective amenable to empirical investigation. The empirical task comes later, where we try to see what, if anything, can be the candidates about what subjects were "really" talking about all along. And that's where we turn in chapter five, where Dennett introduces his model for the first time.