

Chapter Ten of *Consciousness Explained* doesn't have much of a unifying theme, unlike many of the others. It is primarily concerned with more stage setting for Dennett's attack on qualia and traditional ways of understanding the mind. Consequently, the ideas from section to section tend to be disjointed and do not always flow from one to another. I see this chapter as a bit of minor house cleaning before the main assault.

1: Rotating Images in the Mind's Eye

We can manipulate mental images, and those of us who are good at it can exploit the salient features of said images and put them to good, sometimes surprising use. One fantastic example comes from Richard Feynmann's *Surely You're Joking, Mr. Feynmann!*, and this passage (from pp. 85 – 86), quoted in *Consciousness Explained*, where Feynmann describes how he used mental images to analyze arcane theorems of topology, a subject he was not formally trained in:

I had a scheme, which I still use today when somebody is explaining something that I'm trying to understand: I keep making up examples. For instance, the mathematicians would come in with a terrific theorem, and they're all excited. As they're telling me the conditions of the theorem, I construct something which fits all the conditions. You know, you have a set (one ball) – disjoint (two balls). Then the balls turn colors, grow hairs, or whatever, in my head as they put more conditions on. Finally they state the theorem, which is some dumb thing about the ball which isn't true for my hairy green ball thing, so I say, "False!"

If it's true, they get all excited, and I let them go on for a while. Then I point out my counterexample.

"Oh. We forgot to tell you that it's Class 2 Hausdorff homomorphic."

"Well, then," I say, "It's trivial! It's trivial!" By that time I know which way it goes, even though I don't know what Hausdorff homomorphic means."

What mental images are, and whether it even makes sense to say they exist, has been a point of contention among cognitive scientists even to this day. Dennett focuses on a series of experiments that started with the work of Roger Shepard on the "mental rotation" of three dimensional Tetris-style block figures. Subjects would be shown drawings of two different figures and then asked whether they were the same shape seen from different perspectives. In cases where the drawings were indeed the same shape, the experimenters varied the degree of angular rotation between the pairs and also measured the time subjects took to reach a decision.

Typically, subjects would say that they rotated one of the shapes in their mind's eye until they could perceive a match or mismatch. One would think that if an image were truly being rotated, it ought to take subjects (for instance) twice as long to rotate an object through 90 degrees than 45 degrees. This is, in fact, what experiments have confirmed. It goes even beyond this, for there is evidence that the rotated images experienced both "inertia" and "momentum"!

On the surface, at least, this might appear to be evidence from science that there is, contrary to Dennett's central claim in *Consciousness Explained*, a Cartesian Theater. One researcher into mental images of these sorts—Stephen Kosslyn—has even offered a metaphor in terms of images being assembled and rotated on an inner display, rather like images on a CRT (cathode ray tube).

Dennett takes the CRT metaphor at face value, with the idea of beginning with something that couldn't possibly be true of human brains, then subtracting the "undesirable properties" to see what we are left with.

We're first asked to imagine the workings of a typical CAD system, which enables architects and engineers to display and rotate virtual 3-D objects on a computer screen. There is nothing image-like inside a computer which is responsible for the CRT picture—of course. Instead, a description of its planes and edges is defined by coded xyz coordinates and then defined points are grouped. Further codes are attached to groups to represent surfaces that are opaque, transparent, textured, colored, and so on. In this way, the utterly un-image like structures within the machine can produce a genuine image on the screen. In effect, a CAD system acts as a sort of "imagination-prosthesis", enabling people to do with much more precision what they can do with limited ability in their own minds.

Next, Dennett asks us to imagine an even more powerful imagination-prosthesis: a CAD system for blind engineers. Such a machine would take in as input questions in ordinary language, perhaps via Braille or a voice-recognition system. It would then output the answers in a similar fashion.

The first version he describes is the “Mark I CADBLIND,” which consists of an ordinary CAD system in front of which is a “Vorsetzer”—a computer vision system with a TV camera aimed at the CRT and robot fingers to operate the keyboard of the CAD. Recall Shakey from Chapter Four. Shakey was a robot with a visual processing system that could analyze simple shapes before it (blocks and pyramids, for example) and use that processing to carry out instructions to manipulate them. The Vorsetzer is similar, except its camera is aimed at a computer monitor on which the CAD system displays its output. It can look at the CAD images, manipulate them with its robot fingers, and answer the questions of the blind engineer.

But the Mark II CADBLIND shows just how irrelevant the CRT images really are—for it does away with the CRT and camera, replacing them with a cable, through which streams the bit map of the images normally produced on the CRT. This is still an absurdly inefficient way of solving problems for our blind engineer, as Dennett describes:

. . . if the information that a particular portion of the bit-map should represent a shadow is already “known” to the CAD system (if this is part of the coded description of the object from which the CAD system generates its images), and if that fact is part of what the Vorsetzer must determine in order to make its depth judgment, why doesn’t the CAD system just tell the Vorsetzer? Why bother rendering the shadow for the benefit of the pattern-analyzers in the Vorsetzer, when the task of pattern-rendering and the task of pattern-analyzing cancel each other out? (p. 292)

So next we arrive at the Mark III CADBLIND. The Mark III takes what it already “knows” about the virtual objects and their properties and passes this information on to the Vorsetzer subsystem by attaching various “labels” to different portions of the bit-map. Some of the spatial properties would be represented directly within the structure of the bit map—“shown” as it were. The labels would suffice to “tell” the subsystem about the rest (hence the chapter’s title).

We should expect that during the course of evolution, which is often sloppy and inefficient, human brains may have come to have various systems that bear resemblance to different versions of the CADBLIND. The phenomenology of the sketchiness of mental images would be evidence in favor of the brain’s use of shortcuts—“occasions on which the brain tells-without-showing” (p. 295). But cases where using mental imagery allows us to spot patterns which we wouldn’t recognize otherwise “points to the exploitation of visual pattern-recognition machinery that could only occur if one part of the brain went to all the trouble of preparing special-format versions of the information for the benefit of those visual systems.” (p. 295) The Feynmann story is such an example.

However, it’s important to realize how metaphorical the processes of the brains “showing” really are. Even a detailed bitmap of the Mark II CADBLIND will render information on color and shading via internal codes rather than internal colors and shades, so in a sense it is all “tell” and no “show.” Dennett writes:

“And no matter how quasi-pictorial or imagistic the process are that eventuate in the Vorsetzer’s verbal answers to questions, they will not be generated in an inner place where the lost properties (the properties merely ‘talked about’ in the bit map) are somehow restored in order to be appreciated by a judge that composes the answers.” (p. 297)

This is tied, of course, to his insistence in Chapter Five that the brain only needs to make discriminations once, and doesn’t need to re-present them to a master appreciator in the Cartesian Theater.

The fact that a CADBLIND system needn’t employ a Cartesian Theater to do its work in no way proves that brains also do not. However, the exercise does prove that a Cartesian Theater is not an a priori requirement for any theory that explains the human talent for manipulating mental images.

2. Words, Pictures, and Thoughts

In earlier sections of the book, Dennett pointed out that verbal thinking is not exactly like talking to one’s self, and mental images aren’t really pictures in the head. We don’t really think in words or pictures—we think in thoughts, as John Maynard Keynes once quipped. What could this mean?

In this section, Dennett briefly outlines the sense in which thoughts require a sort of control structure laid down in the brain. In the case of language, we have an abstract structure which has a social or meme-like component, in that it must be taught, but also has genetic, biological base, in that to some degree it requires evolved structures in brains. The degree to which a structure falls into one category or the other is unimportant.

Language is just one example of a control structure. Any time you learn a new skill, you are in essence installing a control structure of sorts in your mind, one which helps you quickly recognize and organize relevant facts, and gives you a repertoire of responses for manipulating aspects of your world.

Dennett gives examples of other control structures: learning a game, telling time, arithmetic, money, the bus routes, using the telephone (p. 300), but language certainly looms far above any other because it “infects and inflects our thought at every level” (p. 301). He quotes with approval a passage written by philosopher Justin Leiber:

One could say that natural language was our first great original artifact, so natural language, with our brains to run it, was our primal invention of the universal computer. One could say this except for the sneaking suspicion that language isn't something we invented but something we became, not something we constructed but something in which we created, and recreated, ourselves.”—*Invitation to Cognitive Science*, p. 8.

Students of cognitive science and philosophy may be inclined to read this and immediately think of language of thought approaches of the sort promoted by Jerry Fodor. There is an important distinction to be made here. In a sense a language is like a high-level programming language such as Lisp, but that is not what any version of “mentalese” was intended to represent. What Fodor and other proponents of a language of thought are talking about is a much deeper and inaccessible level of language-like representation, something more like machine language than any programming language. The level at which Dennett intends to pitch natural language as a language of thought is more like “talking to one's self”. It's how we cajole, induce and focus our efforts and concentration for more complex tasks extended over longer periods of time.

3. Reporting and Expressing

If you recall, Dennett introduced the character “Otto” to represent the voice of folk intuition and, at times, traditional philosophy in its approach to the mind. Here is what Otto has to say regarding speech acts about mental states:

I can go back and forth between the experience and the candidate report, checking the words against the experience to make sure I have found *les mots justes*. . . I attend to my particular conscious experience and arrive at a judgment about which words would do the most justice to its character. When I am satisfied that I have framed an accurate report, I express it. From my introspective report, you can come to know about some feature of my conscious experience. (p. 303)

Up to a point, this is compatible with the Pandemonium model of speech production discussed in chapter eight, in which a complicated interplay between “content demons” and “word demons” eventually resulted in an output of a speech act. What is missing from that model but present in Otto's is an “Inner I” responsible for editing.

However, very few of our utterances resemble Otto's description above. Mostly, we simply speak without the careful deliberation Otto seems to hold up as typical. We may find ourselves correcting what we originally were about to say, but we don't do so all the time. And at any rate, we have no privileged insight into how we do it.

Nevertheless, Otto thinks something is being left out. He will have no truck with an inner Cartesian Theater, agreeing (with most philosophers, of course) that such a thing is absurd. But, he says:

It's not enough for words to get strung together by some internal mating dance and then uttered. If they are to be reports of someone's conscious mental states, they have to be based somehow on an act of inner apprehension. What the Pandemonium model leaves out is the speaker's state of awareness that guides the speech.- (p. 304)

So goes the common wisdom, and even year's ago when this discussion thread started in another forum, I had not only forgotten Dennett's rejection of this model, but I went as far as to attribute Otto's views to Dennett himself!

Probeman correctly surmised I had gotten things wrong, but I stubbornly persisted until he contacted Dennett himself and confirmed my mistake. This, even though I had read the book and various papers by Dennett several times. Such is the allure of intuition!

In this section, Dennett uses the analysis of philosopher David Rosenthal on our everyday concepts of expressing and reporting to examine this picture and, in effect, explode it from the inside. I should add that for purposes of this section alone, I'm going to follow Dennett by using these terms in a precise technical sense which was not followed in earlier chapter summaries and may not be followed in future posts.

Imagine you have a cat which stands before the refrigerator, looking up at you and perhaps issuing a plaintive meow. Perhaps speaking to a friend or just to yourself, you declare, "The cat wants his supper."

What you have just done is express your own mental state--a belief about the cat's mental state, being a desire for supper. In doing so, you have issued a report about the cat's mental state, or what you take it to be. You are not here reporting your own belief or expressing the cat's desire:

The cat is expressing his desire by standing anxiously by the refrigerator, and you, noticing this, use it as the basis—the evidence—for your report. There are many ways of expressing a mental state (such as a desire), but only one way of reporting on, and that is by uttering a speech act (oral or written or otherwise signaled). - (pp. 305-306)

You can express without reporting. In the example, expressing your belief about the cat's desire could have been done merely by fetching food from the refrigerator or rolling your eyes with exasperation at the cat's timing (p. 306). All you are doing by expressing is providing ample evidence for another observer to infer your mental state. Reporting is always an intentional act and most often involves language. However, in some contexts—for instance, an experiment where the subject must press a button in the presence of some specified stimulus—reporting may be done other ways. And a noteworthy way of expressing a mental state is to report another mental state, which is where the example begins.

Folk psychology models reporting one's own mental state along the lines of reporting events in the external world, and Dennett sees this as a possible source of a problem. But some more stage-setting by way of Rosenthal is necessary.

Let's call mental states about other mental states (beliefs about desires, beliefs about beliefs, hopes about fears, desires about beliefs, and so on) second-order mental states. All we need is some tolerance for recursive-ness, and we can introduce a third-order mental state—such as when "I (1) believe that you (2) think that I (3) want to have a cup of coffee" (p. 306).

So, if I report that I want to read a book instead of working on this chapter summary, I must be expressing a second order belief about my own desire. And if I announce something like, "I believe I want to read a book instead of working on this summary" that report is an expression of yet another, third-order belief. On the standard model, all of these are distinct mental states.

Folk psychology gives us more--we often distinguish thoughts from beliefs. Thoughts in this commonsense view are "occurrent or episodic states—transient events" (p.307). For example, you have no doubt believed that soccer is a game most of your life, but only now having read this sentence has a thought to that effect occurred in you. So there must be corresponding second, third, and other-order thoughts—thoughts about thoughts. Or if you will, thoughts about beliefs, desires about thoughts, etc. And now we come to a head spinning revelation:

When I express a belief—such as my belief that I want to be fed—I do not express the higher-order belief directly; what happens is that my underlying belief yields an episodic thought, the higher order thought that I want to be fed, and I express that thought (if I choose to do so). All this is involved, Rosenthal argues, in the common sense model of saying what you think. (p. 307)

According to Rosenthal, then, since the hallmark of a conscious state is that (barring injury, drugs, or other disruptions of context) it can be reported, this leads to the conclusion that conscious states must be accompanied by the right sort of higher order thoughts in order to be conscious. Yet paradoxically, "The second-order thought does

not itself have to be conscious in order for its first order object to be conscious. You can express of a thought without being conscious of it, so you can express a second-order thought without being conscious of it—all you need be conscious of is its object, the first order thought you report.” (p 307) That is (and there are exceptions one can invent), one typically attends to the state or object the thought itself is about rather than the thought itself.

A simple example helps to reduce the sense of paradox in the idea of an unconscious thought. Say you tip over a cup of coffee at your desk and immediately leap up from your chair. Doing this did not require a conscious reflection that “the desk top would not absorb the coffee, or that coffee, a liquid obeying the law of gravity, would spill over the edge, but such unconscious thoughts must have occurred—for had the cup contained table salt, or the desk been covered with a towel, you would not have leaped up.” (p. 308)

Thus Rosenthal’s is an attempt to define conscious states un-mysteriously within the confines of folk psychology, not by attributing to them “some inexplicable intrinsic property, but the straightforward property of having a higher-order accompanying thought that is about the state in question” (p. 309).

4. Zombies, Zimboes, and the User Illusion

There are many different sorts of philosopher’s zombies. The general concept is that a zombie is an unconscious being capable of producing intelligent speech acts and rational behavior.

Dennett here introduces another variation on the theme: “A zimbo is a zombie that, as a result of self-monitoring, has internal (but unconscious) higher-order informational states.” (p. 310) So, like us, the zimbo can issue reports expressing its own second-order states, but those states are not conscious. Like a very sophisticated version of Shakey, the zimbo can use self-reflection to answer questions about how it (for instance) processes visual information or rotates mental imagery.

(One could argue, or at least I would, that all traditional philosophers’ zombies were zimboes to begin with. Dennett’s distinction merely focuses our attention.)

But how does this picture look in light of Rosenthal’s notion that when a mental state is accompanied by a higher order thought, this guarantees that it is a conscious state? Either the plausibility of a zimbo or Rosenthal’s theory must be discarded.

At any rate, zimboes and zombies (falsely) believe they are conscious. They have a sort of “user illusion” to this effect. A user illusion, Dennett explains, is the sort of illusion one gets from any sort of modern computer interface. The interface provides a sort of virtual reality which the user manipulates to perform various tasks. The real workings of the machine remain hidden behind all these visual and audio layers. He writes:

This is one of the facts that makes a virtual machine such a tempting analogy for consciousness, for it has always seemed that our access to what is happening inside our own brains is limited; we don’t have to know how the backstage machinery of our brains does its magic; we are acquainted with its operations only as they come clothed for us in the interactive metaphors of phenomenology. - (p. 312).

But doesn’t the notion of a user illusion in this sense require a Cartesian Theater with an inner observer as the user? Not necessarily. The Vorsetzer from the Mark I CADBLIND system described earlier is an unconscious observer completely ignorant of the workings of the CAD system it is pointed at, and recall that it could eventually be replaced by a system using “more modest transactions” (p. 312). Rosenthal’s analysis shows that folk psychology has robust tolerance for the unconsciousness of higher order states.

Now back to the zimbo and the tension it introduces: which do we disqualify, it or Rosenthal’s rendition of folk psychology?

5. Problems with Folk Psychology

Well, to a degree it depends on how far we want to credit Rosenthal’s idea that we have “an ever-expandable hierarchy of higher-order thoughts, conceived to be salient, independent, contentful episodes occurring in real time

in the mind.” (p. 314)

Consider cases that appear to confirm his picture: “And then it occurred to him that he was looking right at the missing cuff link” or “He loves her—he just doesn’t realize it yet.” (p. 314). Such sentences allude to transitions from one state to another, in which one becomes conscious of a first order state. But as Dennett mentioned in an earlier section, folk psychology models this on the picture of first observing such states, resulting in belief states the onsets of which are marked by expressible thoughts. In other words, it all happens just as it does with beliefs about external events. This picture was also urged by Otto.

This could be how things work, but it isn’t necessarily the case, as the model of speech production in chapter 8 demonstrated. According to that model, folk psychology actually has things backwards:

It is not that first one goes into a higher order state of self-observation, creating a higher order thought, so that one can then report the lower level thought by expressing the higher order thought. It is rather that the second order state (the better informed state) comes to be created by the very process of framing the report. We don’t first apprehend our experience in the Cartesian Theater and then, on the basis of that acquired knowledge, have the ability to frame reports to express; our being able to say what it is like is the basis for our “higher-order beliefs. - (p. 315)

Not only does the pandemonium model pose an alternative, but Otto’s folk psychology has problems on its own, for the clarity of distinct, hierarchical states is an illusion.

Says Otto, “On the one hand, there is the truth about how it is with me, and then on the other hand, there is what I eventually say about how it is with me. . .although I tend to be a reliable reporter, there is always room for errors to creep in.” (p. 317)

Yes, that is the problem: for each transition we posit, the possibility always exists for error. There is how things are with us, what we believe, what we eventually say, and an episodic thought or two mixed in. Our beliefs in light of episodic thoughts could be mistaken, misremembered over time, or perhaps mis-expressed, with the mistaken expression coming to be (falsely) believed.

So then we have a temptation: why not stop this threatened explosion or regress right at the beginning by declaring that “my thought (or belief) about how it seems to me is just the same thing as what my experience actually is. . .they are one and the same thing.” (p. 318) This would mean that logically, believing that you are in pain is always the same thing as being in pain. We can give in to this temptation, but only if we are very careful about what it ultimately means.

For instance, we could continue to avail ourselves of talk about second order beliefs about first order beliefs, but only under the understanding that they are really the same thing but appear to be different owing to a verbal distinction (as in “Jones set out to marry his fiancée and ended up marrying his wife”-p. 318).

But this will still not do. The problem of error cannot be made to go away—even if we want to posit that it is at least highly implausible that you could be mistaken about how things are with you now, the passage of time virtually guarantees that you can be mistaken about things once were with you. And as the various interpretations of color phi revealed earlier in the discussion of the Multiple Drafts Theory, the possibility of error in remembering shows up even when the time range is measured in milliseconds. If we want to maintain folk tradition here, we are forced to posit differences in states that are impossible for anyone to detect—by a subject or experimenter, even in principle (ie, Orwellian versus Stalinesque interpretations discussed earlier).

Dennett’s suggestion is then to simply abandon the tradition of folk psychology that “individuate states (beliefs, states of consciousness, states of communicative intention, etc)—by their content” (p. 319). Instead of discrete states, in their place we simply have an ongoing “process that serves, over time, to ensure a good fit between an entity’s internal information-bearing events and the entity’s capacity to express (some of) the information in those events in speech.” (p. 319)

So what a good theory will have to do is not assume that folk psychology is correct in positing hierarchies of discrete thoughts about beliefs and beliefs about thoughts, but explain why we come to describe ourselves as being

structured this way. After all, the tendency to posit processes of internal observations with events to observe is a regress, one which must eventually be “stopped by a process that unites contents with their verbal expression without any intermediary content-appreciator.” (p. 320)

And in the next chapters, Dennett begins the most ambitious steps in *Consciousness Explained* to show how we can do this.

By Faustus (Brian Peterson)