

Is it a bird?
Jerry Fodor

THE BIG BOOK OF CONCEPTS. Gregory L. Murphy. 555pp. MIT Press. Pounds 29.95. - 0 2621 3409 8.

Problems with old and new approaches to the theory of concepts

When the winter nights draw in, cognitive scientists gather around their campfire. They sing cheerful songs to keep their spirits up; and they recount to each other the legends of their tribe, such as, in Gregory L. Murphy's words, that "important progress has been made in the psychology of concepts. Many interesting principles and generalizations have been discovered . . . and even if the field does not agree on the overarching theory that encompasses all of them, that does not deny that those discoveries have been a real advance".

That is good news if it's true. It is among the great metaphysical mysteries how cognition is possible; how a mind can think about, or perceive, or learn about, the world. Arguably, concepts are the pivots such achievements turn on.

On one hand, concepts are the constituents of thoughts; you can't think about eating a tomato unless you have and exercise the concepts tomato and eating. So "a person who has no concepts starves while surrounded by tomatoes", because it never occurs to him that he might eat one. And, on the other hand, thoughts connect to the world only if their constituent concepts do. The thought that tomatoes are edible is about tomatoes being edible because the concept tomato applies to tomatoes and the concept edible applies to comestibles. What with one thing and another, it is very plausible that a good theory of cognition must consist, in large part, of a good theory of concepts.

Well, the consensus is that the cognitive psychology of concepts has achieved a breakthrough in the past several decades.

"The classical view (of concepts) has taken a big fall. Into this vacuum other theories developed." In fact, according to the received account, only a handful of theories of concepts remain as serious options, and much has already been done to choose among them.

The Big Book of Concepts is a lengthy exposition of this consensus. It summarizes, and comments on, an impressive variety of the key experimental findings, and it is a reliable guide to the standard interpretations of these data. Murphy has provided a really invaluable resource for students and researchers, and the merely curious will benefit from skimming. At a minimum, the nearly forty pages of references are a godsend for those of us who are bibliographically challenged.

That said in praise of the book, however, I find that I believe practically none of its conclusions. I think the differences between the "classical" view of concepts and the current revisionist alternatives have been oversold; they are variations on much the same theme, and they are hopeless for much the same reasons. I am prepared to bet that, when the paradigm really does shift, hardly anything recognizable will be left standing. I should emphasize, however, that Murphy's optimistic view is the consensus, and that what I have to say against it is widely considered eccentric.

According to the classical story, most concepts are definitions. The caveat is because, on pain of circularity, if all the concepts that are "complex" are defined, the rest end up undefined and "primitive". Which concepts are primitive was a traditional bone of philosophical contention. Typical empiricists held that they are exclusively "sensory" or "observational", hence that conceptual content derives entirely from experience. Typical rationalists denied that the limits of experience are the limits of conceptualization; rather, the primitive concepts are innate and might include, for example, cause or object or God. That most concepts are definable was, however, common ground in this debate, so definitions are important, according to either view. And definitions purport to express necessary truths. If it is definitional that bachelors are unmarried, then there couldn't be a married bachelor. If it is definitional that dogs are animals, then it is not possible for anything but an animal to be a dog. Murphy misses the connection between definition and necessity entirely, as do most of his colleagues in cognitive psychology. We will presently see something of what that costs them.

It turns out, in fact, that there are a host of embarrassments for the view that typical concepts are definitions. For example, a definition is supposed to provide necessary and sufficient conditions for something to fall under the concept it defines. If bachelors are, by definition, unmarried males, then being a male and unmarried is what you need, and it's all you need, in order to be a bachelor. But, in a vast number of cases even of quotidian concepts, it is surprisingly hard to articulate such conditions. My dictionary says that a chair is a seat with four legs, but it's wrong. For one thing, it's not a chair if all the legs are in the same corner. For another, there could be -in fact, there are -one-legged chairs, the one leg being in the middle. Niggling counter examples keep cropping up, "bachelor" to the contrary notwithstanding. Strict definitions are very thin on the ground. But then it is hard to believe that they could be the key to the theory of concepts.

Also, the classical theory is mute on a point that cognitive psychologists have recently come to care about a lot, namely that there is often a distinction between the good ("typical", "stereotypical") instances of a concept and those that are less good or marginal. No doubt it follows from the definition of "penguin" (assuming there is one) that penguins are birds. Still it is intuitively plausible that a penguin is a less good case of bird than, say, a sparrow or a robin; if you ask them, subjects in experiments will tell you so.

As Murphy exhaustively recounts, recent work has convincingly demonstrated the pervasiveness and reliability of such "typicality effects". They aren't, perhaps, strictly incompatible with the classical story about concepts, but they are ubiquitous and the classical story does not predict them. There is lots more wrong with the classical story, but this much should do to be getting on with.

Hence the recent call to revolution. It used to be said that a penguin is a bird because bird is part of its definition; likewise *mutatis mutandis* for all the other concepts supposed to be complex. By contrast, current revisionism says that complex concepts are held together by a relation of typicality or (subjective) probability. Accordingly, to satisfy a concept is to be relevantly similar to its good or stereotypical instances. Penguins are birds, but they are bad examples of their kind because they lack so many of the "features" that birds are generally supposed to have. They don't fly, for example, but "good" birds do; and "good" birds don't dress for dinner. As Murphy says, "almost every conceptual task has shown that there are unclear examples and variation in typicality of category members . . . the concept's (presumed definitional) core is simply not explaining most of the data. As a result, most researchers have argued that the concept core can simply be done away with, without any loss in the ability to explain the results".

So that is the new proposal: to grasp a concept is to know what features typical things that it applies to generally have; and, in the great majority of cases, the structure of complex concepts is not definitional but statistical.

Some of the old problems disappear at once. Since (unlike truth by definition) probability and typicality are graded, there can be better or worse examples of the things a concept applies to. In effect, what makes robins and penguins both birds is that the concept bird is associated to both; but what makes robins better birds than penguins is that the association is stronger in the one case than in the other. It is no accident that the traditional language of associationism sounds natural in this context. The revisionist story about concepts is, among other things, a recidivist empiricism. And, in their deepest heart of hearts, empiricists have generally thought that the strength of a belief is just the strength of associations among the concepts it contains. So, back to Hume. As one grows old, the cyclical theory of history really does come to seem quite plausible.

I don't propose even to try to adjudicate between classicists and revisionists.

Suffice to say that, pace Murphy, it seems a stand-off at best. Murphy is right that typicality-based theories account for a lot of experimental data which definition-based theories don't. On the other hand, the definition theory of concepts makes room for a distinction between necessary truths and mere common-or-garden truths, and this distinction has intuitive force: it is not just that triangles are very good examples of three-sided figures, or even that the probability that a triangle has three sides is very high; it is that nothing that lacks three sides is even possibly a triangle. Typicality theories have trouble explaining why that is so, much as definition theories have trouble explaining why penguins are bad birds. Also, since typicality theories say that a bird is something that is sufficiently similar to a robin, they need a viable account of similarity.

This proves hard to find. As all philosophers know, everything is similar to everything else in some way or other, so only relevant similarities count for category membership. Just now there are birds in the garden, and there is also a bird bath; so birds and bird baths are similar in respect of the current locations of some of each. Still, no bird is a bird bath; no bird is even an atypical bird bath. How, then, does one decide which similarities are relevant and which aren't? The awful possibility looms that the only kind of similarity to a robin that makes a thing a bird (good, bad, or indifferent) is similarity in respect of being a bird. That is no help at all, of course, if the project is to say what the concept bird is.

Like most cognitive psychologists, Murphy is out of sympathy with such perplexities, indeed with considerations of theory in general. It appears he is worried that, closely pursued, they might require his reading "Quine, or even Frege I feel that we should consider a different way of settling the question: seeing where the data lead us". Yes, if you like. But which data? Data about necessary truths? Typicality data? Data of yet some other kind? The sad truth is that, if you put the wrong questions to Nature, Nature will give you the wrong answers. Nothing substitutes for thinking. Not even experiments.

My point, so far, is that the choice between the classical theory and the current alternatives is less clear than proponents of the latter are wont to advertise.

But perhaps that doesn't matter since, if I'm right, it's Tweedledum and Tweedledee. I think both kinds of theory are deeply wrong, and that assumptions they share are a large part of the trouble.

For example, both are up to their ears in the postulation of "features".

According to the one, bachelors are unmarried men because the features "unmarried" and "man" belong to the definition of bachelor. According to the other, penguins are marginal birds because they lack the feature "flies", which is part of the bird- stereotype. Certain questions thus suggest themselves:

- What, exactly, are features? How are they individuated? How many of them are there? Do primitive concepts have features (does the concept red have the feature "red")?

Is it that the feature "is generally a creature that flies" is part of the concept bird? Or is it that the feature "is generally a bird" is part of the concept creature that flies? If what makes a concept complex is its being a bundle of features, mustn't we suppose, on pain of regress, that the features themselves are all simple? But if "red" is a feature (because tomatoes are typically red), mustn't the feature "red" contain the feature "colour" (because red things are typically coloured)? What, if anything, is the difference between the features of birds that constitute the concept, and the many, many other things that one believes about birds? Or, if everything one believes about birds is part of one's bird concept, how is it possible even to consider the possibility that some of one's bird-beliefs aren't true? And, if conceptual truths are just statistical, what is the difference between "'bird' is a feature of the concept penguin", which sounds pretty impressive, and "it is generally supposed that penguins are birds", which sounds pretty flat?

Murphy hardly even contemplates such issues, nor (to my knowledge) does anybody else in cognitive psychology. Instead, he thinks the uncritical postulation of features is justified because "it is difficult to understand exactly what it would mean to comprehend a word (concept) if it is not to bring to mind its components (features) Every concrete theory of semantic representation does use components of some kind, because there is no other obvious way to represent the meaning". But obvious is one thing, true is another. In fact, the concept of a feature is a blank cheque. I see no prospect of its being cashed.

Nor is what is wrong with the current psychology of concepts just its playing fast and loose with features. There's a bookful to say about this; all I have room for here is a gesture in what I take to be the right direction. Consider, then, the question of what it is to have a concept. The classical story says that it is to know the concept's defining features. That would be plausible except that most concepts don't have definitions (see above). Typicality theories have a corresponding problem: although there are plenty of graded concepts like bird, not every property a thing has (graded or otherwise) is a determinant of its kind. I think, for example, that it is wildly unlikely -out of the question, really -that there are penguins on Mars. Certainly, no typical penguin lives there. But my thinking this isn't necessary for my having the penguin concept. I could have had "penguin" even if I'd never heard of Mars.

Correspondingly, if someone does believe there are Martian penguins, it is likely to be his Mars concept rather than his penguin concept that wants revision. (It really is very cold up there, after all; and there aren't any Martian fish for Martian penguins to eat.) The long and short is that having "penguin" can't consist just in holding some high-probability beliefs about which penguins are typical; not even if the beliefs are true. What, then,

should a typicality theorist say about the nature of concept possession?

We approach a great historical divide. Sometime early in the twentieth century, the dominant empiricist tradition in theories of mind took a pragmatist turn (Dewey was a paradigm). The idea is that having a concept is a kind of "knowing how". At least in many cases, and at least in part, it is knowing how to sort things into the ones that the concept applies to and the ones that it doesn't.

Clearly, this account of concept possession comports comfortably with the theory that concepts are structures of typicality beliefs.

The two are connected by the idea that one's concepts are constituted by the beliefs that one uses in sorting the things that concept applies to. Though I am certain that there are no penguins on Mars, I don't generally use that fact in deciding whether a thing is a penguin. That is why my belief that it is typical of penguins not to live on Mars doesn't count as part of my penguin concept.

So, having a concept is knowing how to use it; in particular, it is knowing how to use it to sort things. This variety of pragmatism was, I think, the defining catastrophe of twentieth-century theories of mind, and practically everybody still believes it. Ask an experimental psychologist how to test for concept possession (or concept acquisition), and you will be told to run a "categorization experiment": see whether the mind in question can distinguish the things that the concept applies to from the things that it doesn't. "The majority of research on concepts has focused on the process of learning to categorize", just as Murphy says. Or, ask a philosopher what possessing a concept consists in, and more than likely you will be told that it is mastery of "criteria" for the concept's application. Concepts express clusters of family resemblance, and having them is knowing how to tell who is in which family.

No wonder typicality theorists often view themselves as latter-day Wittgensteinians.

This view of concept possession is now so pervasive in philosophy and cognitive science that it is hard to remember it is a view and not a revelation. Things weren't always so. If you had asked Hume (or, indeed, practically anyone else before around 1900), you would have been told that having a concept is being able to think about whatever it is the concept of. What you need, and all you need, to have the concept penguin is to be able to make penguins "present to the mind"; you do not, in particular, need to be able to sort penguins. I must say, that strikes me as extremely plausible. I myself claim to have the concept penguin, though I doubt that I could tell them from many of their close relations.

It is a long, sorry story how pragmatism gained the upper hand. Suffice to say that no behaviourist -not even a closet behaviourist -can allow "thinking about" to stand at the heart of a theory of cognition; for, unlike sorting, thinking happens on the inside and is hard to observe. One of the tribal legends, widely credited, says that the "cognitive revolution" in the 1960s swept out that closet.

In fact, though, practically all experimental psychologists and philosophers of mind continue to be behaviourists of one kind or other. They have just ceased to notice that they are.

So here is what I take to be the situation. Definitional theories of concepts perished for lack of examples. Typicality theories are committed to pragmatism for their account of concept possession, and pragmatism is certainly untrue.

Since, at present, these exhaust the options, we really are in the dark about concepts. Maybe somebody will find the light switch sooner or later. Or maybe not. I do, however, owe you an argument that pragmatism is indeed certainly untrue. I will conclude by sketching one that I take to be pretty conclusive.

Everybody (excepting, maybe, a handful of "connectionists" who are beyond the pale) agrees that some concepts are complex. For example, "brown cow" clearly is; and, on the face of it, its constituents are "brown" and "cow". (Whether these are its primitive constituents is moot; fortunately, we needn't decide that question for the purposes in hand.) It is because concepts have this sort of combinatorial structure that they are "productive"; that is, it is possible to make up new ones ad lib (brown penguin, purple cow). Moreover this combinatorial structure is "compositional": the content of complex concepts is determined entirely by the content and arrangement of their constituents.

Because all there is to brown cow is what it inherits from brown and cow, if you have brown and cow you get brown cow for free. In particular, you don't have to examine any cows in order to acquire it. However, exactly, the details go, some such account of the productivity of concepts is widely, and rightly, considered to be essential. A theory of concepts must get compositionality right or it is expelled from the club.

But typicality theories cannot get compositionality right. The point -what is often called the "pet fish" problem -is transparent. Typicality doesn't compose. Typical pet fish are neither typical pets nor typical fish. (There is, of course, no reason why they should be. The intersection of two sets needn't contain the typical members of either.) Correspondingly, somebody who is able to sort typical pets versus everything else, and typical fish versus everything else, is not thereby enabled to sort typical pet fish versus everything else. So concepts cannot be typicality structures, and pragmatism is certainly false. QED.

Murphy devotes a chapter to this sort of problem, but it is thoroughly unsatisfactory. He says, in effect, that the content of complex concepts is sometimes "emergent" from the content of their constituents. In these cases, grasping the concept requires consulting "real world" knowledge. But this misses the point, which is precisely that you don't have to examine brown cows to acquire brown cow. All you need to grasp is the constituents. If concepts were typicality structures, then one would have to enquire at pet stores to acquire the concept pet fish. But they aren't, and so one doesn't.

Having a concept is being able to bring to mind whatever it is the concept of; and all one needs to bring pet fish to mind is mastery of the concepts pet and fish. Since pet fish is the concept of a pet that is a fish (and vice versa), people who haven't a clue what fish are kept as pets can fully grasp that if something is a fish that's kept as a pet, then it must be a pet fish, whether or not it is typical. (See how the necessity truths keep coming back.) If, in order to have pet fish you had to know which pet fish are typical, you couldn't so much as enquire whether guppies (say) are pet fish: you already need to have pet fish (and guppy) even to raise the question.

It is part of our not knowing how the mind works that we don't know what concepts are or what it is to have one. Just about everything that current cognitive science says about either topic is wrong. But at least it is clear that concepts aren't typicality structures and that having them is

not being able to sort things. Except for leaving that out, Gregory Murphy's book tells you most of what there is to the psychology of concepts. Read it, therefore, by all means; but don't even consider believing it.