

Review of Jose Luis Bermudez, 'Thinking without Words'

Jerry Fodor

'Dr. Livingstone, I presume?' Thus, famously, Sir Henry Morton Stanley. Stanley was, of course spot on; it *was* Dr. Livingstone. Elsewise his presuming so wouldn't have become the stuff of legend. A question suggests itself: How did he manage to presume so cleverly? Of all the things that Stanley might have presumed, how did he hit upon the one that was both pertinent and true? Why didn't he presume Queen Victoria, for example? Or Tower Bridge?

First blush, that sounds like an easy sort of question. In fact it's an abyss. Though philosophers and psychologists have been working on such matters for a couple of millennia, the best they've got is less a theory than a program of research. That is the background for Bermudez' book, so let's start with it.

This seems safe: Stanley must have done some thinking. He must have inferred, on the basis of his beliefs, memories, hunches (etc.) about the situation in which he found himself, that it was most likely Livingstone that he ought to presume. 'The situation in which he found himself' thus included not only whatever was perceptually available at the scene, but also a lot of cognitive commitments that Stanley brought with him. If he inferred that it was Livingstone, it must have been from those sorts of premises that he did so; he had, after all, nothing else to go on. I suppose that's all pretty much truistic; still it prompts some useful reflections.

Notice, to begin with, the intimate relation between thinking and inferring. At the crucial point, Stanley's thinking must have consisted of drawing inferences from what he independently believed. It's plausible that (at least some kinds of) thinking just *are* processes of drawing inferences. Likewise for a lot other things the mind does like learning, perceiving, and planning. The picture that emerges is of the mind (or the brain if

you prefer) as some kind of inferring machine; perhaps some kind of computing machine, since computations are themselves plausibly construed as chains of inferences.

Second, if the mind is in the inference-drawing line of work, there must be symbols in which it formulates its premises and conclusions; there are no inferences absent a medium (or media) in which to couch them. That matters because you can't say just anything you like in whatever kind of symbols you choose. Pictures, for example, can't express negative or contingent propositions (*it's not raining* and *if it's raining, that will spoil the picnic*). But negative and conditional thoughts play a central role in the kinds of inferences that minds routinely carry out. (*It's certainly not Queen Victoria; If it's certainly not Queen Victoria, then perhaps it's Dr. Livingstone. So perhaps it's Dr. Livingstone*). Such considerations suggest, at a minimum, that the mind doesn't do all of its thinking in pictures. In fact, they suggest a strategy for empirical research: find out what kinds of inference minds can make, then figure out what kinds of symbols they would need in order to make them. You will arrive, if all goes well, at a theory of those kinds of mental representations that figure in thinking, perceiving, learning and the like insofar as these are inferential processes. It turns out that this kind of research is feasible, and not without significant results.

It seems likely, for example, that the kinds of representations required as the vehicles of thought are not very different from what 'natural languages' (English, German, whatever) provide as vehicles of communication: sentences, or something of the sort. Hence the talk in cognitive science of a language of thought in which our cognitive processes are carried out. This seems hardly surprising. English is used to communicate our thoughts, so it must be that English is rich enough to express their content. So English, or something like it, is *prima facie* plausible as a model of the system of symbols that we think in. That's very convenient because we already have in hand quite a powerful account of (some of) the kinds of inferences that natural languages can be used to formulate; we have logic. The inferential account of mental processes thus offers a nexus between the kinds of inquiries that cognitive psychologists pursue, and the kind that logicians do. This is good news. We can all use all the help that we can get.

Again, if the inference/language of thought story about thinking is even close to being right, a lot of what goes on in the mind must be unconscious, which is to say: unavailable to introspection. This is true not just of the etiology of neuroses and the like, but also of the higher cognitive processes that are supposed to mediate our higher cognitive capacities. The thesis that most of what goes on in the cognitive mind is unconscious follows, pretty directly, from the thesis that a lot of what goes on there is the drawing of inferences.

Thus some of the intuitions that underlie a venerable tradition of theorizing about the mind. It certainly isn't the only one, nor is there a current consensus. Over the years, psychologists have worried about explanations that postulate unobservables, of which the theory that much thinking is unconscious inference is a paradigm. So psychology suffers from from recurring bouts of behaviorism and associationism, to both of which the thinking-as-inferring story is anathema. And philosophers have worried that, if perceiving consists of inferring from some mental representations to others, perception won't ever make contact with a mind-independent world. These (and other) familiar objections to inferentialism may have merit, but we can ignore them for present purposes since Berumdez' book looks kindly on the language of thought version of the inferentialist tradition. At least the first half does.

One more preliminary: I've been speaking rather freely about what 'the mind' does and how 'the mind' does it, and you may wonder just whose mind it is I mean. Stanley's? Mine? Yours? The Average Man's? On anybody's story there are lots of individual differences among cognizers, not just in how well they cognize, but also in what kinds of cognizing they are able to do (and also, perhaps, in how they are able to do it.) Even if we're prepared to make the quite radical assumption that all human thinking is more or less the same, what about other species? What about dogs and cats and chimpanzees? Does their kind of thinking and perceiving also consist of drawing inferences in a language of thought? If so, is it the same language of thought that we think in? And are they the same kinds of inferences that we draw? Which bullets are you prepared to bite?

Now we're ready for Bermudez, whose main thesis is that inferentialism may be ok for the cognitive psychology of adult humans; but it doesn't work for other species. (It doesn't even work for prelinguistic human infants.) At best it has a handle on the cognitive psychology of creatures that have mastered a natural (public) language; which, de facto, presumably means just us. According to Bermudez, all of that is knowable more or less a priori, and he offers a more or less a priori argument to prove it. Much of his book is devoted to preparing and presenting that argument; much of the rest is devoted to asking just what kind of minds animals and babies might have if they don't have the inferentialist/language of thought kind.

Bermudez has done his homework; he has read a lot of psychology (and neurology; and anthropology) all of which he is prepared to mine for philosophical payoff. That's admirable, and you'll like the bibliography even if you don't like text. And, in fact, I'm pretty sure that not much of what the text says is right. The situation is awash in irony: Though Bermudez is certainly no behaviorist, he endorses the very kinds of proposals that were floated by behaviorists in philosophy and psychology back in the '50s and '60s. And, unsurprisingly, the arguments that worked against them also work against him. Some examples presently; but even at the start the whole idea that there are two (or more?) fundamentally different kinds of minds might strike one as unparsimonious. If, as Bermudez concedes, some sort of inferentialism is likely to work for our minds, isn't the least hypothesis that it is also likely to work for the minds of other kinds of creatures? Surely it's reasonable, absent contrary evidence, to suppose the differences between our minds and theirs are largely quantitative. The latter, after all, are widely supposed to have evolved from the former; and, indisputably, our babies turn into us. The gap can't be impassable in either case.

Bermudez, however, sets his face against this line of thought. His reason for doing so is deeply characteristic: It "doesn't help us with the epistemological issues that must be confronted by an account of thought in the absence of language. It doesn't give us any clues as to how we might go about attributing thoughts to nonlinguistic creatures." His

own method, by contrast, is relentlessly epistemological. Passages like the following are ubiquitous: “I propose to make a start by asking [for example] ... what evidence might there be that a creature is representing the consequences of two or more different courses of action, rather than the actions themselves. Once we have a clear set of operational criteria in view, it will be easier [to answer such questions.]” I think Bermudez’ insistence on this methodology is very surprising; and I think it is disastrous.

In fact, scientific theories don’t usually provide “operational criteria” for attributing the states, events, processes, (whatever), that they purport to describe; that’s why experimental ingenuity is required to test them. The Big Bang Theory didn’t tell us what sorts of evidence we could use to confirm it; we found that out more or less by accident. Likewise, string theory doesn’t tell us how to tell whether string theory is true; quite possibly there isn’t any way to tell whether string theory is true. That is disappointing but it shouldn’t be surprising. The *truth* of string theory depends on whether there are strings. But the *testing* of string theory depends on relations between strings and us. It may be that we’re too big to test string theory; or that we don’t live long enough; or that we live in the wrong part of the universe; or that we lack the required equipment, or that we’re too stupid. Other possibilities suggest themselves. But none of this epistemological stuff is relevant to whether there are strings; ontology is one thing, epistemology is quite another.

I would have thought this sort of point is widely recognized these days. Scientists have, by and large, stopped looking for ‘operational definitions’ of their theoretical terms, and philosophers have, by and large, stopped telling them to do so. Bermudez’ himself remarks that “as a hypothesis about the machinery of cognition, the language of thought hypothesis is under no obligation to provide its own epistemology.” Exactly. Nor (pace Bermudez) would its failure to do so mean “that the language of thought hypothesis cannot provide a straight forward solution to the problem of explaining how ... nonlinguistic creatures might be thinkers.” To the contrary, it provides precisely such a solution *if it’s true*. What it doesn’t do is tell us how to find out whether it’s true; that, to

repeat, requires experimental ingenuity. Possibly it requires more experimental ingenuity than we have, though it is permissible to hope not.

Suppose, however, that one does demand “operational criteria” for applications of the explanatory vocabulary of psychological theories. Since there is no reason to doubt that how an animal behaves depends on what it believes, wants, thinks, intends, and remembers, we will need operational criteria for each of those; that’s what Bermudez’ program shares with behaviorism. But the behaviorist program didn’t work for them, and it’s won’t work for him. There simply aren’t such criteria, nor ought we expect there to be. Scientific theories are about what there is in the world; they are not about how to tell what there is in the world.

Here’s an example of the program not working; one of very many. According to Bermudez “the content of a desire is given by the state of affairs that serves as its satisfaction-condition, and a state of affairs serves as the satisfaction-condition of a desire if its being the case would bring about (in the right way) the cessation of the behavior to which that desire gave rise.” (So, hunger is a desire for food because getting fed is what terminates one’s hunger behavior.) Note, first, how very counter-intuitive this suggestion actually is. Surely the satisfaction-condition for a desire isn’t the state of affairs that would terminate *the behavior that the desire gives rise to*; at best, it’s the state of affairs that would terminate *the desire*. Getting food terminates your hunger. Whether it also stops your hunger behavior depends on the circumstances; notably on what you have in mind. It won’t stop your scrounging for food if you have in mind not just to do some eating, but also to do some hoarding.

In fact, the psychologist Clark Hull (c. 1950) suggested precisely Bermudez’s sort of trick as a way to make the notion of a drive behavioristically respectable: A drive is something the satisfaction of which terminates the behavior that it causes. That came to no good. People pointed out, for example, the ‘salted peanuts’ effect: It’s a truism that the desire for salted peanuts is directed towards eating salted peanuts. Still, eating salted peanuts doesn’t stop you from wanting to eat more salted peanuts; a fortiori, it doesn’t stop you

from scrounging for more. What Shakespeare said of Cleopatra is true of salted peanuts too: they make hungry where they most satisfy.

Behaviorism failed because it couldn't bear the cost of epistemologizing animal psychology. Bermudez has the same problem, and I doubt that he can afford it either. It would be instructive to go through that in detail, but I'll settle for three quick remarks.

First, the argument that Bermudez offers to show that the language of thought kind of inferentialism won't work for animals depends on a lemma: namely, that our kind of thinking requires a public language (which, of course, animals don't have.) The argument for the lemma is that without a public language you can't have 'second-order' thoughts; you can't think about thinking. But the argument for the lemma is unconvincing. It's plausible that you can't think about your thinking unless you can represent your thoughts to yourself; indeed, it's plausible that you can't think about anything that you can't represent to yourself. But what, exactly, shows that you have to represent them to yourself in a natural language? Since maybe they are represented in the language of thought, we seem to be back where we started. Bermudez' rejoinder, after a very long slog, is disappointing (not to say question-begging). It's that "all the propositional thoughts that we consciously introspect... take the form of sentences in a public language" So what? Introspection, has turned out to be a bad tool for investigating minds; including one's own mind. It's among the important results in cognitive psychology that much of what you introspectively believe about your mind just isn't true. Indeed, if what you mean by thinking is whatever actually explains one's higher cognitive capacities, it's an open question whether one can introspect one's thinking at all. To be sure, it does seem that something is going on in consciousness when one thinks; but what's the argument that what seems to be going on actually has much of a role in explaining cognition? What's the argument that it has *any*? The fact seems to be that we are mostly strangers to ourselves.

Second, as remarked above, there are kinds of creatures that start out not having a public language and end up in one: namely, our children. How do they do it if, as Bermudez

believes, there are principled (not just quantitative) differences between our minds and theirs? Bermudez doesn't face this, and I don't blame him. But it's germane to his main thesis. Plausibly, for example, learning English requires learning that the form of words 'it's raining' is properly used to communicate the thought that it's raining. How do you learn that sort of thing if you have the kind of mind that can't, even in principle, think about thoughts?

And finally, it's arguable that (again like the behaviorists before him) Bermudez looks in the wrong place when he tries to estimate the cognitive capacities of animal minds. Practically invariably, he looks at what animals can learn; and it is quite possible that learning isn't where animal minds do their best stuff. Try, for example, perception. There's every reason to believe that even relatively crude perceptual capacities require inferences of very formidable complexity; certainly more complex than what could be done with the resources that Bermudez is prepared to grant. But animals do see. (I know there are philosophers who are prepared to deny that they do, but really...) It appears that the complex inferential processes that subservise perception are typically 'domain specific' and 'encapsulated'. That is, perception may depend on computationally very elaborate processes that are unavailable for other projects like, for example, learning. It could be that the evolution from animal minds to ours consisted in loosening such domain specificity constraints, with the consequence that we can think about a lot more kinds of things than they can. Bermudez doesn't face the computational complexity of the perceptual capacities of animals. Instead, he borrows J.J. Gibson's notion that perception consists (not of drawing inferences but) of "picking up" information about what the distal environment "affords". I don't know what that means, and I'd guess that Bermudez doesn't either. Or Gibson, come to think of it.

The moral of all this is that behaviorism really is dead. Even fancy, sophisticated, philosophical behaviorism really is dead. And the kind of behaviorism that seeks to impose epistemic constraints on the ontology of psychological theories is especially dead. Bermudez gives it an exhilarating run for its money, but it just won't do. Perhaps that's surprising. Who'd have guessed that animals and infants may not wear their minds on

their sleeves? Who'd have guessed that their their behaviors might manifest their thoughts only very indirectly? Who'd have guessed that they are, in these respects, so similar to us?

