

Twenty-five Theses against Cognitivism

Jeff Coulter

IN WHAT ensues, my intention is to focus the case against cognitive science and its derivative, cognitive neuroscience, in the form of (maybe) debatable propositions. I will use the method of paragraphing *à la* Ludwig Wittgenstein, but some of my arguments will extend beyond LW's aphoristic style, although hardly beyond his remarkable contributions to our understanding of issues which, since his death, still haunt us and to which his insights still apply.

(1) 'Mind' is either a theorist's reification (hence a fallacy) or a vernacular concept. Hence, we can have 'What is on your mind?' (meaning 'What are you [*qua person*] thinking' etc.); 'It slipped my mind' (meaning, *inter alia*, 'I forgot it'); 'I could not get it out of my mind' (meaning, 'I was obsessed by it'); 'Apply your mind' (meaning, *inter alia*, 'Concentrate on the problem'), etc., etc. 'Mind' has many vernacular uses, none of which commit the user to any form of reification (= the fallacy of misplaced concreteness, of hypostatization) of what the word 'mind' means. Capitalizing it does not make it into an entity.

(2) Since 'mind' is not the name of any *entity*, it cannot be isomorphic with, identical to, or reducible to, or functionally dependent upon, any other entity, such as a brain. 'Brain' is a name for a bodily organ whose operations, processes, states and functions indeed enable us *as persons* to behave in ways which license the avowal and ascription of 'mental' predicates. However, even the so-called 'mental' predicates are truly predicates whose locus of attribution is the person, and not any real (physical) or projected ('mental') component thereof.

■ *Theory, Culture & Society* 2008 (SAGE, Los Angeles, London, New Delhi, and Singapore),
Vol. 25(2): 19–32
DOI: 10.1177/0263276407086789

(3) The processes and states characteristic of neural functioning lack intentionality in Brentano's sense, hence cannot be equated to many of our 'mental' concepts. There can only be 'thoughts' of or about (etc.), and there is only 'consciousness' of something (etc.), while there are no neuronal firings that, of, about (etc.) anything. You can have a cup and fill it with coffee. Then you will have a cup of coffee. You can empty the coffee and be left with the cup. However, you cannot have a thought of *X*, delete the *X*, and still have 'the thought'! This is a conceptual, not an empirical, truth. It is a grammatical point.

(4) Computationally inspired conceptualizations of brain functions do not escape this constraint. Indeed, metaphorical constructions derived from computation are radically misleading and frequently incoherent in this domain. Brains do not receive, transform and process 'information', they form no 'representations' and they do not 'store' alphanumeric strings in digitized forms. Brains are energy-transforming organs whose functions enable *people* to receive, transform, process and store information of an alphanumerically expressible kind (wherever and whenever this is a criterially founded fact about their conduct). Even here, the expressions used ('receive information', 'store information', etc.) are not discrete and generic characterizations of conduct but are rather polymorphous characterizations subsuming hugely heterogeneous cases of situated conduct.

(5) Cognitivist conceptions are notoriously non-rigorous about the character of their *explananda*. They frequently aspire to operationalize or to substitute for, or to regiment, our mundane concepts of thinking, believing, hoping, expecting, remembering, forgetting, reasoning, understanding, perceiving and a host of others. They are typically insensitive to the conceptual articulations and rules of use which apply to their actual topics of inquiry, favoring modes of stipulative theorizing to the analysis of their grammars. One of the favored maneuvers of cognitivism is the 'globalization' of its *explananda*: thus, we are supposed to be considering Memory, Thought, Learning, Understanding, Perception, Recognition, etc. as though these were names for discretely identifiable and homogeneous empirical phenomena for scientific scrutiny and study. It is as if we were studying phenomena comparable to chemical elements, molecules, known species of plants or animals, etc. Nothing could be farther from the truth. We shall return to this theme later on . . .

(6) No one in his right mind thinks that *brains* can ride bicycles or drive cars. These are things that (able) people can do. However, the projection onto brains of person-level capacities and activities (*viz.*, their personification) is a central feature of much cognitivist theorizing, subserving the computational conception of brains as physical-symbol manipulators and information-processors. Brain functions properly described in the logically appropriate terms of biochemistry, electrophysiology and anatomy do indeed

facilitate persons' doings and accomplishments, but they do not engage in parallel 'activities' of, for example, parsing utterances, contextualizing behaviors, following rules, calculating distances and the myriad phenomena theoretically ascribed to them in a host of cognitivist models. Such purely conjectural attributions to brains and central nervous systems are made in complete independence of their actually satisfying the logical criteria for such ascriptions.

(7) One central cognitivist hypothesis has struck many as entirely reasonable, even though it is genuinely incoherent. It is the claim that memory is a matter of storing and retrieving information. Our mundane criteria for individuating memories (e.g. recollections, things recalled or remembered), that is, our criteria for distinguishing between one recollection and another, or the beginning and end of any particular recollection, are purpose- and context-dependent. Since many memories are available in their tellings, in their being recounted to an audience, a significant constraint upon their individuation is the redundancy principle (you seek to refrain from telling people what you can assume they already know). Details involved in a recollection are thus contingent upon the purpose(s) informing the telling and often the audience to whom or for whom the recollection is being produced. Brains, not being linguistic and social beings, are therefore not possessed of any such individuation criteria. Thus, the notion that they store memories (in the form of information, however neurally 'encoded') begs the individuation problem. It also assumes that what is recollected was originally stored in alphanumeric form, and this idea does not gibe with the fact that many things we can and do remember were not originally experienced in or as a symbolically 'packaged' form. Moreover, one can sometimes display the fact that one has remembered something *non*-discursively, as when, for example, one suddenly grasps one's car keys in one's left jacket pocket showing (context permitting) that one has remembered where one's car keys are. The same recollection could have been expressed in words, for example, 'My car keys are in my left jacket pocket' in circumstances where one had trouble locating them and then did. Storage-and-retrieval theorizing cannot accommodate the *diversity of manifestations* of memory. Further, remembering and recollecting are 'achievement verbs' (Ryle's concept), not 'process verbs': since 'retrieving' is a *process* (its speed is irrelevant here), it cannot be constitutive of remembering.

(8) The project of work in the field of Artificial Intelligence (AI) is directed toward the creation of (virtual) androids, simulacra of human beings insofar as their *conduct* is concerned. Cognitivism draws upon many ideas developed under the purview of this program, but does so in a way which strangely reverses the methodology of AI. Whereas AI seeks to decompose and componentialize various human capacities and practices in order to create programs which effectively simulate them computationally, cognitivism treats such regimentations (especially those which have an algorithmic

character) as if they were (formal) descriptions of the actual ‘underlying’ (*sic*) procedures by which the original human instances are accomplished by brains. The two approaches are utterly distinct, notwithstanding the frequent efforts made (mostly by philosophical proponents of both) to conflate them. The former is a valuable engineering (software and hardware) effort which has paid off in terms of, *inter alia*, the production of useful expert systems and increasingly manageable human–computer interfacing. The latter has been a source of consistent confusion for the behavioral sciences.

(9) The concept of a ‘rule’ of conduct (linguistic and otherwise) has played an important role in cognitivist theorizing, but it has been appealed to in misleading ways. That human practices are rule-governed is by now a truism, but cognitivism insists upon (i) assigning a location to ‘rules’ (the mind/brain) and (ii) treating them as essentially inaccessible to rule-following agents. The former involves rule-reification, while the latter is a function of mistakenly assuming that rules are essentially formal codifications of effective procedures, or can be adequately reformulated in such terms, of which lay folks are typically ignorant. Cognitivists either miss or conflate several important distinctions in this argument. There are distinctions to be drawn between rules that people follow, rules which govern (= constitute) what their conduct amounts to (= means) and rules with which their conduct is *in accord*. Only the latter is really significantly involved in efforts to simulate facets of human conduct (in AI). Competent cyclists can corner on their bicycles without the least effort, but their accomplishment in so doing is consistent with (i.e. *in accord* with) what Polanyi described as an ergonomic principle of which most of them are completely *unaware*, *viz.*: a cyclist must keep his balance by adjusting the curvature of his forward path in proportion to the ratio of his unbalance over the square of his speed. For a computational system to be designed to simulate cornering on a bicycle, that is all that is needed (at least as far as its foundational understanding is concerned). Competent cyclists do not (normally?) *follow* any such rule, nor is ‘cornering on a bicycle’ as a practice *governed* (= constituted as such, as intelligibly the case) by any such rule in the sense that one can ascribe such a bit of behavior to a cyclist on the basis of mundane criteria (e.g. he was tilting his bicycle to the right or left, did not fall off, turned to/the corner and kept on cycling). I cannot *follow* a rule I know nothing about, although what I do is constituted by rules I might be unable to formulate or reflect upon. *Constitutive* rules are not to be construed in Searle’s unfortunate terms as ‘necessary and sufficient conditions for the application of a concept’, but rather in most cases as guidelines, as criteria recognizable to agents, etc. There is no need to construe constitutive (or any other) rules as if they were necessarily algorithmic or poor cousins of algorithms. There is, nonetheless, a big difference between a rule (or algorithm) with which some bit of human conduct is in (logical) accord (a prerequisite for a successful AI-type simulation to work), a rule which the agent of the

conduct *follows* (which presupposes that he knows what it is), and a rule which enables someone to distinguish correctly between his *X-ing* and *Y-ing* or *Z-ing* in some circumstance or other.

(10) Of course, then there is the problem of ‘context’. In AI circles, this is best known as the frame problem (after Minsky). This is the problem (a technical one for AI, but a logical one for cognitivism) of being able to distinguish between a ‘context’ which abstractly co-constitutes the nature of the phenomenon in question (e.g. is that a question or an answer?) and the ‘relevant context’ which could, in principle, settle such an issue. (To get a grip on this, consider the following scenario: A says to B: ‘When did they get home?’ which utterance superficially satisfies the criteria for being a question. Now add the following contextual particular: B says, before A says anything: ‘What did she ask you?’ to which now A’s reply, as given above, comprises an answer [by virtue of its being also a quotation]). The frame problem has affinities with Gurwitsch’s ‘gestalt-contexture’ idea. Gurwitsch’s insight was that, while it is true that the sense of something (its intelligibility) depends upon its context of production/presentation, it is also true that whatever ‘context’ can mean in relation to the intelligibility of a phenomenon (the sense of an utterance, the nature of an activity, the characterization of an object, etc.) is not pre-given in any codifiable way. In fact, argued Gurwitsch, the relationship between intelligibility and context is akin to that between figure and ground in a gestalt structure. In gestalt configurations, no figure, no ground but also, no ground, no figure. They are *mutually* constitutive and thus logically inseparable. No one can formalize or codify the circumstantial particulars relevant to the intelligibility which something has in advance of its presentational ‘formats’. But its *relevantly* circumstantial materials depend entirely upon the sort of thing that the phenomenon might conceivably be. For example, if I flex my right index finger while pointing to someone with my right hand palm upward and my other fingers drawn back, this bit of behavior can be constituted variously as: ‘beckoning’, ‘exemplifying a bit of human behavior’, ‘flexing my finger(s) before playing the piano’, etc.: add a weapon, and it might be ‘killing’, ‘seriously/fatally wounding’, ‘assassinating’, ‘executing’, ‘combat/target practice’, ‘murdering’, etc. Point my right hand to my head while executing the same movement and it could be ‘trying to commit suicide’, ‘committing suicide’ and so forth. The action, the rules relevant to it and the relevant context (which grounds whatever rules are in fact operative) co-constitute what it amounts to. The purely physical components of whatever act of commission I might thereby be engaged in are entirely explicable in neurobiological terms, but the action that I am engaged in is not so explicable. Further, there is no sense to the idea that one could separately list or codify, on the one hand actions and their rules, and on the other hand, the contexts within which the rules apply so as to constitute what the actions are. They are imbricated, inextricably connected, as are figures and grounds in gestalt configurations. Thus, ‘contexts’ *qua* ‘relevant contexts’, resist formalization,

codification, enumeration and regimentation a priori. A *context* is not a set of pre-specifiable ‘conditions’, even though one can, for purely pragmatic purposes, list some bunch of circumstances and add a ‘similarity rider’ . . . But this is a far cry from the rigorous demands of strict formalization, and especially from the sort of regimentation advocated in the idea that ‘contexts’ might be formulable as ‘sets of necessary and sufficient conditions’ for an action to have been performed, with the intelligibility that it had. For example, waving to someone might be waving hello (greeting someone), waving goodbye, hailing a cab, signaling to someone, and a host of other non-enumerable albeit potentially envisageable scenarios: what makes the action one of those or *something else entirely* is not disclosable in any pre-specifiable fashion. The well-known defeasibility-of-rules argument in the philosophy of jurisprudence (Hart et al.) is akin to this one. After all, a ‘context’ could be simply ‘the surroundings’, but it could also be ‘identities of agents’ and it could also be ‘what just came before what was said/done’, etc. There is no logical exemption from this major constraint upon any effort at formalizing the nature of ‘contexts’ for human conduct, linguistic and otherwise. Having duly noted these issues, however, it is of course the case that, considered in a purely physical sense, ‘flexing the finger’ has a completely unassailable neurobiological explanation (see e.g. Carlson on the physiology of such a movement).

(11) Post-behaviorist computational mechanism (Shanker’s lovely expression) does not *exempt* itself from so many of the problems which beset its purely physicalistic precursors, *viz* behaviorism and mind–brain identity theorizing, whether we are speaking of Watson, Hull or Skinner, or of Carnap, Broad, Place, Smart, Armstrong or Fiegl. Before 1950, there was no computer science, no Turing ‘test’, no computationalist versions of mind, intelligence and conduct the burgeoning of which eventually nurtured the cognitivist revolution in the human sciences. That revolution was heralded by Newell and Simon and Chomsky at MIT in 1956. Hilary Putnam embraced this position and tried to advance it within the philosophy of mind for many decades before abandoning it. Also, we had the parallel discovery by y Cajal of the nature of neurons in the human brain, and their apparently ‘digital’ character (i.e. their on–off firing patterns). The stage was set for the amalgamation of all of these influences, which established cognitivism as a post-behaviorist computational mechanism. I want in this section to discuss the nature of its ‘mechanism’ or mechanistic-materialist (with due deference to Feuerbach) conception of the human agent. First of all, there is the constant ubiquity (from Chomsky to Fodor, from Kosslyn and Koenig to Pinker) of the concept of a ‘stimulus’ as the required notion of an ‘input’ to the CNS (central nervous system) such that subsequently theorized ‘cognitive operations’ can be construed as operating on such ‘input’ to produce whatever ‘output’ can be postulated. However, since Watson’s original formulation of the behavioristic project in psychology, a ‘stimulus’ has been either (i) rigorously defined as any quantum of energy impinging upon a

receptor organ (e.g. photons impinging upon photoreceptors in retinas, acoustic wave-forms impinging upon receptors in cochleas, etc.), or, more commonly, (ii) any object, situation or circumstance oriented to by an agent. These two alternative characterizations were always incompatible, and never interchangeable (cf. G.H. Mead and his insistence upon the distinction between the physical environment and the ‘symbolic’ [conceptualized] environment), although in practice there has been a considerable degree of ‘slippage’ between them all. The issue is that cognitivist theorizing is rarely if ever alert to this distinction, between a ‘stimulus’ as a physically definable input to the CNS and a ‘stimulus’ as a conceptually grasped ‘input’ in the experience of a person. Why is this so important? Because blurring these critical differences leads to incoherence in the characterization of what an ‘input’ is, and also to whatever ‘cognitive operations’ (such as interpretation, disambiguation, parsing, contextualization, representation, computation, etc.) are deemed necessary to account for the putative ‘gap’ between ‘input’ and (behavioral) ‘output’. Neural operations on energy quanta are describable, wherever possible, in the language of neurobiology and sensory neurophysiology alone. ‘Mental’ operations putatively ascribable to situations-as-experienced, actions-as-witnessed, utterances-as-heard, are entirely distinct from whatever physical and physiological processes subserve their being seen, heard, felt, etc. In fact, there is no evidence whatsoever to link the former to the latter, since the former are definable phenomena for empirical inquiry, and the latter are mere theoretical postulates ordained by a purely speculative and mentalistic metaphysics of the mind/brain predicated upon the appeal to the computationalist conception of mind and conduct.

(12) *Images* and *sounds* are not ‘inputs’ to the brain and CNS: energy quanta are, that is, photons and acoustic wave-forms. Thus, there is no mystery surrounding the fact that retinal images are two-dimensional and inverted, nor is there a mystery surrounding the fact that sounds are variously oriented to by hearers/listeners (one major source of this misconception was Helmholtz). Brains do *not* process *images* nor *sounds*, only energy quanta, thus there is no need to postulate any hypothetical intermediary processes of ‘adding a third dimension and re-inverting the image’ (Hyman, Hacker), nor is there any (logically ordained) requirement for conjecturing how brains ‘add meanings’ to sound waves. Just as there is no (logical space for, need for) a *physics* of chairs (whose major criterion is a functional one, *not* a physical one, *viz* sit-on-ability!), so is there no logical space for a (neuro-) physiology of *mind* nor for what – for the Cartesians – were the ‘mental’ attributes and properties of persons. In other words, brains and ‘minds’ are neither of them logically connectable ‘phenomena’ since only *one* side of the equation names a genuine *phenomenon*. Moreover, the putatively ‘mental’ predicates (which after Sprague I would prefer to call simply a subset of ‘personal predicates’) are not logical candidates for mappings, isomorphic testing, reducibility, etc. in relation to the human central

nervous system. This is going over old ground, but with new excavating tools.

(13) Neuroscientific imperialism was never a major issue for the non-biological human and behavioral studies until the advent of *cognitive* neuroscience. Of course, there were frequent polemical skirmishes between the disciplines, especially over the notion that anything non-biological in conceptualization was either unscientific or merely a place-holder for the eventual arrival at the deeper, truly *scientific* truths about human conduct. Such conceits linger on today, but usually as idle polemical positions (on *both* sides of the issue). However, the challenge of cognitive neuroscience to be the *all-embracing human science*, encompassing (at least in principle) *all* of human conduct, experience, mentality, emotion and even social interaction (concretely construed) is one not to be taken lightly, but also a challenge not to be missed! A caveat: just as Chomsky missed completely the revolution in understanding communicative interaction introduced by Harvey Sacks and Emanuel Schegloff (after the introduction by Harold Garfinkel of the concept of ‘ethnomethodology’ – the endogenous analysis of practical action and practical reasoning), preferring his insular algebraically inspired games with only orthographic concepts and sentential instances, so it is conceivable that our cognitive scientists just are not aware of what has been happening in the social sciences over the past 30 years. To read Steven Pinker’s lament about the so-called ‘Standard Social Science’ position on human agency (*viz* the moribund ‘*tabula rasa*’ philosophy of the past century) is enough to inspire doubt that some of our colleagues in adjacent disciplines who purport to characterize their opponents know what they are talking about, and underscores the possibility that they simply do not know the extent to which, and the degree to which, their arguments fail to impress us. But then, as members of presumably subordinate disciplines in the so-called scientific hierarchy, maybe we should be more circumspect than to go on the attack. Herein lies the province of, and the utility of, much modern analytic philosophy to the task at hand. Here, we need to revisit Wittgenstein.

(14) The ignorance of Wittgenstein’s argument against the logical possibility of a ‘private language’ might seem, at first blush, to be a matter of purely scholarly interest. However, this is far from the case, especially when considering the polemics attending a range of cognitivist claims as discussed here. In particular, we have to confront Fodor’s famous claim that there is indeed a ‘private language’, and that it is the unconscious and inaccessible ‘language of thought’. Wittgenstein’s original foil was Descartes’ notion that *only* an agent him or herself can truly understand what, for example, ‘red’ might be. Others can conceivably agree on what to ‘call’ such ‘experiences’, but then (so the classical metaphysical argument goes) who can tell if what you call blue is red to me, even though we share a public vocabulary for speaking of such matters? Wittgenstein made a manifold

counter-attack against such conceptions. First, he noted that while many public phenomena can be grasped by agents in and through their ostensive training in what they are commonly called, even ostensive training by itself is insufficient to inculcate a mastery of the concepts which are used referentially. For true referential mastery to occur, neophytes must acquire the rules for the use of such expressions as ‘red’ etc., in the practice of engaging in language-games. That is, they must learn how to use such expressions in ways common to (rule-governed in) the community of speakers of the language. But not all nouns have referents, and not all expressions are descriptive, formal similarities notwithstanding. Even a public ostensive definition can amount to only a sterile ritual of ‘associating a symbol with an object’ unless and until any such ostensive training can be elaborated into the mastery of a practice of usage. However, *private* ostensive definitions, whatever they might be, necessarily lack the criteria for distinguishing between what is thought to be the case and what actually is the case, even when what is thought to be the case is sincerely avowable. There is no such thing as (i.e. no logical space for) an essentially (in principle) *incommunicable* rule. But that is exactly what the model of an interior, radically, essentially, private, wholly ‘introspective’ ostensive definition presupposes. Wittgenstein argued for its complete incoherence. The ‘inner’, he reminded us, stands in need of *outer* criteria. Thus, samples involved in learning what is, for example, ‘red’ must be public, and, once the rules of use are mastered, there is no longer any room for an appeal to what something might *look like* to me, which might not *look* the same to you, which might subvert the intersubjectivity of the samples and rules subserving the use of the word (red). Since (merely) looking like *X*, appearing to be *X*, presupposes the mastery of what it is for something, criterially, to *be* (an) *X*, appearances do not have the last word in language learning: indeed, ‘appearances’ can only come into play *after* the grammar is mastered, which establishes what can count as a genuine case of (*not a mere appearance of*) an *X*. Once we have mastered what it is for something correctly to be characterized as, for example, ‘red’, *then* we can argue *in special cases* about whether or not *that* is red, etc. The bearing that this remarkable argument still has on cognitivism is simply this: cognitivism presupposes the isolated agent/speaker as a cognitive being with his ‘language of thought’, and tries to work outward. Wittgenstein taught us to begin with the social order of intersubjective communicative praxis and work back toward characterizing the properties of any individual agent/speaker as a cognitive being. ‘No man is an island’ is not just a poetic or a sociological shibboleth: it is a logical truth. Fodor’s putative ‘private language’ is no language, nor is it a ‘language’ of ‘thought’.

(15) Consider a founding assumption of cognitivism: machines can be made which think, understand, are intelligent, in the same sense that human beings are. SHRDLU and SOAR are AI accomplishments of considerable complexity and simulate effectively some facets of human comprehension and intelligence. However, the idea that one can extrapolate from the

mechanics of computational simulations of facets of human conduct to the explanation of such conduct in its original human form is the Achilles' heel of cognitivist doctrine. To borrow an example from my colleague, Wes Sharrock, if I sort through the mail mechanically and the Post Office sorts through the mail mechanically, have I done what the Post Office does? Can I use the one phenomenon to elucidate the nature of the other? It is fine to stipulate that, in the domain of computation, 'following rules' = 'implementing algorithms', as long as one realizes that such a stipulation carries no weight in terms of the elucidation of what it is for a *human agent* to 'follow a rule'. One can quite simply claim that a computer *can* understand something, so long as one realizes that one has changed the rules of use of the concept of understanding such that, given the change, the computer case is utterly distinct from the human case. The problem with cognitivist theorizing, yet again, is its propensity to conflate distinctive grammars of concept formation, which results in analogies which cannot be carried through in any coherent manner.

(16) Understanding is *not* the same thing as interpreting (an utterance, an action, etc.), although on occasion an interpretation can aid in understanding something. Interpreting can take time and can result in a textual product of some sort (i.e. the interpretation), while understanding is not a process but an achievement (Ryle) manifest in relevantly correct performance(s) in appropriate contexts. Some of these performances may be discursive, and some may not be. Nonetheless, it is the scenic performance (or its possibility) which comprises the criterion for having understood something, not any phenomena postulated as 'internal' to the person. Trying to understand may involve (polymorphous) processes, but actually understanding is neither an act nor a process of any kind. We distinguish between thinking that one has understood and actually having understood: thus, understanding carries *no subjective sovereignty* for its claimant. Since understanding in these cases (we exempt the notion of 'empathy' here) does not designate *any* process, the notion that understanding speech is a matter of 'processing' it (the concept favored by cognitivists infatuated with computational jargon) cannot pass muster.

(17) Recognizing is another achievement verb often misassimilated in cognitive models to a process, one characteristically involving some form of search-and-match procedure (cf. Marr). Retrieving data from buffers in computers is the background 'paradigm' for most more substantive cognitive theories of human recognitional capacities, but, again, as in the case of recollecting (and understanding), we are dealing here with achievements, with abilities, and not processes. Moreover, in the human case the search-and-match conception is circular: to be able to use a match between, for example, what is seen and heard now and what was seen and heard before to facilitate the recognition of something in front of one presupposes that the agent can recognize the match as such. Further, most cases in which we

recognize things do not require *assistance from* anything. Note as well that sometimes one can recognize something *as a case of a* something without ever having perceived that object before. Again, many cognitive models/theories in this domain fail adequately to unpack the many facets of their putative *explanandum*. Efforts to answer a generally posed question like: ‘How do we recognize phenomena?’ are bound to result in confusions, since the question makes no sense in the form in which it is (typically) posed. Instead, it is important to note that only particularized queries make sense, for example, ‘How did he recognize her in the crowd?’ or ‘How did you recognize me after all these years?’ and the like, and these can, naturally enough, have perfectly respectable vernacular answers.

(18) Technicalities aside, I want to address a few sources of opposition to the kind of arguments which I have been presenting here (and which I have adumbrated in considerable detail in many publications over the past 20 years or so). The first and most pervasive attack against those of us who oppose cognitivism is that our antagonism is based upon a notion of ‘anthropocentrism’. I plead guilty to this ‘charge’, and can only mutter an echo of the words of Ryle – we aren’t mechanisms, we are people. Those who advance the doctrines of materialist reductionism, cognitivism and sundry other related metaphysical positions about the nature of human agency frequently castigate the opposition as *anti-scientific*. This misplaced charge seems to me to embrace a conception about human beings that is not ordained by factual analysis nor by any significant adherence to the strictures of conceptual propriety. Theirs is not a matter of superior knowledge of neuroscience, nor of any other science, including computer science. If that were the case, and were sustainable by any evidence, then we would have made whatever concessions we might have been required to make a long time ago. To think otherwise is to accuse anti-cognitivists of either (i) being deliberately deceptive in our use of evidence or (ii) being merely blinkered to the true interpretation of whatever the neurosciences and the computational sciences have delivered. The challenge is this: anti-cognitivists have no dispute whatsoever with the established facts of neurobiology nor with computer science nor, I would add, with *genuine* advances in Artificial Intelligence research. Our dispute is with the dogmatic legacy we are bequeathing to our students and colleagues who buy wholesale a bill of goods because it is widely subscribed to by members of prestigious institutions such as MIT and so forth. Arguments from authority, and appeals to a majority consensus in *any* discipline, must be set aside when there are disputes about matters which have been debated for quite a while and which have so far *not* been resolved.

(19) The idea of the ‘human mind’ has been a profound distraction in the human sciences for more than a century. Materialist and behaviorist efforts to eradicate it, while well-intentioned, did indeed fail, but not because their principal objective was unsound. Their efforts were predicated upon the

idea that ‘mind’ is a phenomenon to be (empirically) investigated, but that whatever this peculiar phenomenon is, it certainly is not a Cartesian ‘ghost’, ethereal substance or other metaphysical entity. Before Wittgenstein, and the analytical tradition in the ‘philosophy of mind’ which he bequeathed to us, we simply lacked the logical and methodological tools with which to dispose of this pseudo-problem. It has now been more than 50 years since Wittgenstein’s *Philosophical Investigations* was published, and yet still we encounter professional philosophers and behavioral scientists struggling with problems derived from their preoccupation with conceptions of mind and mentality which are unnecessary and deeply misleading. The impact upon the neurosciences has also been problematic, as we can see from the trouble which Sir John Eccles, the last of the great Cartesian dualists in that discipline, encountered in making his claims intelligible. Now that we (at last) have works such as Max Bennett’s and Peter Hacker’s *Philosophical Foundations of Neuroscience*, that legacy can finally be put to rest.

(20) That ontological problems about ‘mind’ and ‘the mental’ have ‘grammatical’ resolutions is not something that many people with vested interests in cognitive science and cognitive neuroscience want to hear, especially in this age of Churchland’s ‘neurophilosophy’ and Searle’s pronouncements about ‘the rediscovery of the mind’, to say nothing of the legacy of Noam Chomsky et al. In spite of dozens of papers, chapters and books tackling these matters from a (broadly) Wittgensteinian perspective, it is still apparently troublesome to many to learn that a cloud of confusion can be dissolved in drops of grammar. Nonetheless, the fact remains that none of Wittgenstein’s opponents have shown that the entire range of ‘mental’ concepts and predicates cannot be analyzed *without residue* into properties of *personal* conduct, dispositions, capacities, achievements, circumstances and other ‘outer’ matters. Expecting, hoping, believing, imagining, understanding, remembering, recognizing, forgetting, perceiving, thinking, wondering, dreaming, intending – indeed, the whole array of what Descartes once construed as the ‘invisible’ things we humans do, requiring the postulation of the (non-physical) *res cogitans* (the human mind) – have *all* been shown, in many demonstrations, *not* to require any such reference whatsoever. Although in these brief summary comments, I cannot develop such a claim more extensively than I have hinted at (with examples of some of the so-called ‘mental’ concepts), there is a vast literature available to anyone willing to consult it which could at least arguably substantiate such an assertion. I will conclude with four examples to persuade the reader to look further into such (apparently still very contentious) claims.

(21) Consider ‘intending’. When I have an intention to do something, the ‘have’ is not the ‘have’ of possession (any more than it is in expressions such as: ‘I have an obligation . . .’ or ‘I have a bus to catch’). There is no entity corresponding to the intention I have, nor is there any inner process corresponding to (nor constituting) my intending to do something. Indeed, if I

have the intention of leaving town tomorrow, I need be in *no* particular, discernible ‘psychological’ nor physiological *state*, and my intention consists in what I said it did (if I am truthful). To intend to *X* is a kind of commitment, weaker than promising or swearing to do something, perhaps a bit stronger than merely being ready to do something. No one seems to think that ‘commitments’ might have neural correlates, and it is clear that the sincerity of one’s (avowed) intention is exhibited in what one then proceeds to do in relevant circumstances.

(22) Stich sub-titled one of his books: ‘The Case Against Belief’, claiming that since ‘beliefs’ are hard to align with cognitivist postulates, their ontological status is dubious at best. However, to believe that something is the case is often (although not always) to be tentative about one’s claim, to be, as it were, hedging one’s bets about what one declares to be the case. (Here we locate the distinction between ‘knowing’, which is factive, and ‘believing’, which is not.) Believing *in* something, by contrast, is usually a matter of *conviction* (religious, political, etc.). What one truly believes is a matter of what one says and does. Contrary to Dennett, we do not carry our ‘beliefs’ around in the neural equivalent of a compartment in our heads, but rather what we believe is shown in, displayed by, what we are disposed to say/do and what we actually, in relevant circumstances, say/do. Again, note that ‘having a belief’ is to use ‘have’ *non*-possessionally.

(23) As with (discursive) ‘thoughts’, ‘dreams’ are, roughly, what they are of or about. We can dream of *X*, dream about *X*, dream *X* up, but (unlike with thoughts) we do not ‘dream’ things over or dream things through. Nonetheless, the concept of ‘a dream’ is relational, not free-standing (cf. the discussion of the cup with and without coffee in paragraph #3 earlier). We learn to predicate ‘dreaming’ of ourselves in virtue of a particular sort of ‘teaching link’ (Pears). A child spontaneously, after waking, avows a perceptual claim which its guardians cannot ratify, even in principle, but they refrain from simply dismissing it: instead, they take the opportunity to teach the child a new language-game involving the insertion of a prefatory expression (e.g. ‘I dreamt that . . .’) before the perceptual account. Nocturnal dreams (as compared to daydreams) involve an array of physiological correlates, some of them perhaps approximating even to markers, such as specific patterns of alpha rhythm activity or REM (rapid eye movement) sleep, but ‘dreams’ themselves are accounts of a special logical status, and the truth of such accounts is a function of the truthfulness of the agent who avows the accounts, nothing further. There is nothing neurophysiological about the *nature* of dreams, even though neurophysiology may teach us much about whatever neural correlates there may be when people dream.

(24) It has been argued in various places that the Kornhuber–Libet experiments revealed the nature and scope of a neurophysiological sequence of (electrical) events referred to as a ‘readiness potential’, which has been

interpreted as a neurological correlate of an ‘act of will’ prior to an instance of voluntary conduct. However, such an interpretation is flawed. First, the experimental circumstances in which the evidence of the ‘readiness potential’ was derived depend upon giving subjects instructions to, for example, ‘flex your right index finger at will’. The periodicity of the response to the instruction is clearly arbitrary, but then only within the limits of the temporal parameters of the experiment itself. The subjects’ finger-flexings were *instructed actions*, and, elsewhere, instructed actions can be *contrasted* to instances of voluntary ones. Second, the idea that conduct engaged in voluntarily comprises all conduct *not* engaged in *involuntarily* is mistaken. There are more than two applicable concepts in play here, and, as Austin reminded us, not all actions are ‘adverb-hungry’. One may do something reluctantly, although not out of coercion. One may do something out of an obligation or a requirement, but not because one wants to do it or feels like doing it or would *otherwise* (= such conditions *not* obtaining) do it. In addition, note that ‘choosing’ is not itself a feature of even most instances of what we would call ‘doing something freely’: ‘choosing’ is *itself* a kind of activity presupposing the availability of real alternatives in some situation(s) and not an ‘underlying’ property of all voluntary conduct.

(25) In the foregoing, I have been concerned to survey a range of issues and arguments relevant to the adjudication of claims made by cognitive theorists. Very little of what I have said here is original, although it often appears that very little of what has been discussed is given its proper consideration or weight in many circles where cognition is the central focus of intellectual interest. Elaborations and defenses of many of these points may be found in the writings of, *inter alia*, and by no means exhaustively: Gilbert Ryle, Ludwig Wittgenstein, Norman Malcolm, J.L. Austin, P.M.S. Hacker, G.P. Baker, Max Bennett, Bruce Goldberg, Stuart Shanker, J.F.M. Hunter, David Cooper, Gilbert Harman, John Hyman, Meredith Williams, Hubert L. Dreyfus, Roy Harris, Keith Gunderson, Elmer Sprague, W.W. Sharrock and myself. If the reader’s interest has been sufficiently piqued to stimulate him or her to explore any of these issues further and in greater argumentative and exemplary detail, then this effort will have been worthwhile.

Jeff Coulter is Professor of Sociology and Associate Faculty Member in Philosophy, College of Arts and Sciences, Boston University.