PHILOSOPHICAL PSYCHOLOGY, VOL. 14, NO. 1, 2001

The bounds of cognition

FRED ADAMS & KEN AIZAWA

ABSTRACT *Recent work in cognitive science has suggested that there are actual cases in which cognitive processes extend in the physical world beyond the bounds of the brain and the body. We argue that, while transcranial cognition may be both a logical and a nomological possibility, no case has been made for its current existence. In other words, we defend a form of contingent intracranialism about the cognitive.*

1. Introduction

One way to find the product of 347 and 957 would be to apply the partial products algorithm in one's head. First compute the product in the rightmost column, carrying if necessary. Next compute the product of the second column, then add any carries. Repeat as necessary across the columns. The shortcoming of purely mental computation is that it taxes one's relatively fixed cognitive capacities. In particular, it makes serious demands on one's memory capacities. For each carrying operation, one has to remember the number to be carried while computing a product. Further, one needs to keep the various places straight. A four is in one of the tens places and a five is in the other; a three is in one of the hundreds places and a nine is in the other. Another method for finding the product of 347 and 957 would be to write the problem down on a piece of paper before applying the partial products algorithm. Since the numbers can be written one above the other, one can rely on vision to keep the ones, tens, and hundreds places coordinated. In addition, since one can write down the number to be carried above the column to which it will be carried, this would remove the burden of remembering the number to be carried. Further, by recording one's work at each step, one is spared the task of remembering where one is in the calculation. In making use of pencil and paper, one deploys a different set of cognitive capacities than that deployed in performing the computation in one's head. It is because the use of pencil and paper generally provides a faster and more reliable method of computing products that one so frequently turns to it.

The common sense way of understanding what is going on in this sort of situation is to say that tools, such as paper and pencil, allow us to work around, or

Fred Adams, Department of Philosophy, University of Delaware, Newark, DE 19716, email: fa@udel.edu; Ken Aizawa, Department of Philosophy, Centenary College of Louisiana, Shreveport, LA 71134, email: kaizawa@centenary.edu

ISSN 0951-5089/print/ISSN 1465-394X/online/01/010043–22
DOI: 10.1080/09515080120033571

move beyond, the limitations of the relatively fixed cognitive capacities residing in our brains. Address books, Rolodexes, and speed dialing on phones help reduce demands on long-term memory. Warning lights next to the car's gas gauge and kitchen timers help spare us the costs of inattention. Calculators, slide rules, and computers provide tools that enable us to perform logical and mathematical operations more quickly and reliably than we might when relying exclusively on the limited resources in our brains. Microscopes, telescopes, mass spectrometers, IR spectrometers, stethoscopes, and high-speed photography convert environmental energy into a form usable by our sensory apparatus. In all these cases, common sense has it that our cognitive faculties, restricted to the confines of our brains, can be aided in any manner of ways, by cleverly designed non-cognitive tools.

As so often happens, however, philosophers there are who would challenge common sense. Daniel Dennett, for example, has suggested, if not out and out asserted, that we should view the brain *and its paraphernalia* as a single cognitive system. He claims that the primary source of human intellectual superiority over animals:

> is our habit of *offloading* as much as possible of our cognitive tasks into the environment itself—extruding our mind (that is our mental projects and activities) into the surrounding world, where a host of peripheral devices we construct can store, process, and re-represent our meanings, streamlining, enhancing, and protecting the processes of transformation that *are* our thinking. This widespread practice of offloading releases us from the limitations of our animal brains (Dennett, 1996, pp. 134–135).

Dennett is, of course, right to emphasize the role and significance of tools in our cognitive lives. There is nothing contrary to common sense in this. Where he does make a radical departure from common sense is in thinking that, when tools are used, cognition is a "transcranial" or "extracranial" process. It's certainly a wild idea to suppose that to use a calculator is to have one's mind bleed out of one's brain into plastic buttons and semiconductors.

Andy Clark and David Chalmers defend a similar account of tool use. They provide a thought experiment in which a woman, Inga, believes that the Museum of Modern Art is on 53rd Street and, upon hearing that there is an interesting exhibition there, proceeds to 53rd Street. Otto, by contrast, suffers from Alzheimer's disease, hence carries around a notebook containing all sorts of useful information, including the location of the Museum of Modern Art. When Otto hears of the interesting exhibit, he consults his notebook, then sets out on his way. Clark and Chalmers contend that, while Inga and Otto differ in all manner of superficial respects, in relevant and important respects, Inga's memory and Otto's notebook are the same. They maintain that the essential causal dynamics of the two cases mirror each other precisely, that Inga's memory and Otto's notebook fall under one natural kind, namely, beliefs, and that a cognitive science that treats Otto and Inga together is the explanatorily simpler, hence superior, cognitive science.

Yet a third supporter of the radical interpretation of tool use is Merlin Donald. In an elaborate theory of the evolution of the human mind, Donald argues that the last dramatic transition in the cognitive evolution from australopithecines to *Homo sapiens sapiens* occurred with our use of "external symbol storage." This began some 40,000 years ago with the first uses of visuographic representations in the form of body decoration, grave decoration, and object arrangement, and continues today with the vast array of representational possibilities opened up with multimedia productions. Donald claims that these "exograms" (to be contrasted with Lashley's "engrams") constitute a vast memory store that has radically changed the architecture of human cognition. More specifically, exograms enable humans to engage in analytic thought not found in prior cultures.

> The major products of analytic thought … are generally absent from purely mythic cultures. A partial list of features that are absent include: formal arguments, systematic taxonomies, induction, deduction, verification, differentiation, quantification, idealization, and formal methods of measurement. Argument, discovery, proof, and theoretical synthesis are part of the legacy of this kind of thought. The highest product of analytic thought, and its governing construct, is the formal *theory*, an integrative device that is much more than a symbolic invention; it is a system of thought and argument that predicts and explains. (Donald, 1991, pp. 273–274).

So, in many respects, Donald's views are the same as Dennett's and Clark and Chalmers'.

In a very interesting discussion of navigation aboard a US navy destroyer, Edwin Hutchins (1995) develops a view of cognitive science that supports, or might be taken to support, the radical view of tool use. Part of Hutchins's concern is methodological, namely, to study "cognition in the wild," as it happens in its everyday natural environment, as opposed to its occurrence under laboratory conditions. The contrast here is roughly that between anthropological observation and experimental determination. The aim is to address such questions as, "What do people use their cognitive abilities for?" and "What kinds of tasks do they confront in the everyday world?" Such a shift, while a potentially radical methodological move, is perfectly consistent with the common sense metaphysical view of the bounds of cognition during tool use. There is nothing contrary to common sense in investigating the ways in which brain-bound cognition works in its everyday natural environment, rather than under experimental conditions. It is perfectly reasonable to ask what people use their brain-bound cognition for and what kinds of task they confront with their brain-bound cognition. Where Hutchins threatens to depart from common sense, toward Dennett's radical transcranial cognition, is in his analysis of the nature of cognition. Hutchins tells us:

> Having taken ship navigation as it is performed by a team on the bridge of a ship as the unit of cognitive analysis, I will attempt to apply the principal metaphor of cognitive science—cognition as computation—to the operation of this system. In so doing I do not make any special commitment to the nature of the computations that are going on inside individuals except to say that whatever happens there is part of a larger computational system.

> But I do believe that the computation observed in the activity of the larger system can be described in the way cognition has been traditionally described—that is, as computation realized through the creation, transformation, and propagation of representations states. (Hutchins, 1995, p. 49)

If cognition is simply computation over representational states, and if one's tools, such as paper and pencil, form or contain representations, then one has a case for the radical view that, in at least some cases of tool use, cognition extends beyond the boundary of the brain.

In this paper, we propose to defend common sense. Our view is that, as a matter of contingent empirical fact, in all actual cases of human tool use brain-bound cognitive processes interact with non-cognitive processes in the extracranial world. Think about the use of lopping shears. When a person chops thick branches from a tree with lopping shears, we might treat the human and her tool as a single system. If we do this, however, we still haven't the least reason to suppose that muscular processes inside the arms extend from her arms into the lopping shears. We have a relatively good scientific understanding of the biomechanics of muscular elements, such as actin, myosin, adenosine triphosphate, adenosine diphosphate, and creatin phosphate, and a very good scientific understanding of the physics of the levers, fulcrums, and cutting wedges in lopping shears. We know enough about these matters to know that the processes in muscles are quite distinct from the processes in the shears. In much the same way, we think that enough is known about psychological processes and other physical processes to rule out the possibility that, in real world cases, we have cognitive processes spanning the bounds of the brain.

In defending what we take to be common sense, we don't propose to challenge a principle articulated by Clark and Chalmers: "If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* … part of the cognitive process" (Clark & Chalmers, 1998, p. 2). To us, this means that the skull does not constitute a theoretically significant boundary for cognitive science. More specifically, it means that being inside the brain cannot be the mark of the cognitive. This seems to us to be true and obvious. The bounds of cognition must be determined by finding the mark of the cognitive, then seeing what sorts of processes in the world have the mark. Following this method, we see that, as a matter of contingent fact, the cognitive processes we find in the real world all happen to be brain bound. It appears to be just a contingent empirical fact that cognitive processes are not transcorporeal processes.

Just to emphasize the contingent nature of our claim, we might note that, in contrast to cognitive processes, it appears to us that, in some organisms, digestive processes *are* transcorporeal. That is, while cognitive processes, as a matter of contingent fact, are corporeally bound, it appears that digestive processes, as a matter of contingent fact, are not. Some spiders inject their prey with enzymes that digest their prey from the inside out, enabling the spiders to suck the partially digested material out of the victim's body. Further, a wide range of animals, including chickens and rabbits, practice coprophagy. As food passes through the digestive system, some components are broken down by the body's digestive acids and proteins and are assimilated into the body. This is responsible for some digestion. As the remaining food mass passes farther through the digestive system, it becomes laden with bacteria that inhabit the intestines of the animal. Once the remaining mass is excreted, the bacteria that have infiltrated the feces use digestive enzymes to break down components of the food that were not broken down by the digestive chemicals contributed by the animal. Once the bacteria have acted upon the feces for a sufficiently long period of time, the feces can then be consumed by the animal and the previously indigestible materials digested and absorbed. In these cases, the biological processes that occur outside the body are sufficiently similar to the biological processes that occur within the body that one can maintain that digestion is, in at least some organisms, a transcorporeal process. In our estimation, the common sense view is that, while cognition could have turned out to be transcorporeal, like various kinds of digestion, it in fact turns out to be intracorporeal, like muscular contraction. Cognitive processes are so different from the physical processes in the tools we use that a science that ignores this difference essentially ignores cognition.

Making good on this analysis will involve, in the first place, drawing attention to differences between the cognitive and the non-cognitive. It will involve saying something about the mark of the cognitive. Although we are in no position to offer a final understanding of the mark of the cognitive, we think that enough is at present known about it to render cases of transcranial/transcorporeal cognition most unlikely. The idea here is that, even though transcranial cognition is logically and perhaps nomologically possible, given what we know now about the human brain and cognition, the chances that some tool in our hands would have the cognitive properties our brains do is rather unlikely. An apt analogy: although it is logically and nomologically possible that Aizawa hits a hole in one the first time he tees off, given just some rather vague knowledge of physics and Aizawa's lack of athletic ability, objectively speaking, the probability of a hole in one is pretty low. When it comes to the analysis of examples, it appears that trancranialists largely ignore what is known about the brain and cognitive processing. In fact, we find that radical theorists tend to rely on a behavioral—not to say behaviorist—conception of the cognitive: if behavior B is normally produced by a set of cognitive processes G, then any other set of processes G**FF32* that produces B is also a set of cognitive processes. Cognitive scientists have generally rejected such behavioral conceptions of cognition, since they allow that gigantic look-up tables might count as cognitive models.

As our first step in making good on our analysis, Section 2 of our paper will sketch, not a theory of the mark of the cognitive, but two conditions we take to be necessary elements in the mark of the cognitive. With this orientation, it will be much easier, in Section 3 of our paper, to see what goes wrong with Clark and Chalmers' and Hutchins' reasons for taking a radical interpretation of tool use. Finally, in Section 4 of our paper, we shall try to forestall further arguments that might be brought forth against the theory of brain-bound cognition. Where Section

3 examines the case for transcranialism found in the cognitive science literature, Section 4 tries to anticipate some rejoinders on behalf of transcranialism.

2. Some necessary elements of the mark of the cognitive

The traditional question in the philosophy of mind is the mark of the mental. Here we assume without argument that the mental is not the same as the cognitive, hence that the mark of the mental is not the same as the mark of the cognitive. We take learning, remembering, sensing, perceiving, and thinking to be paradigm cases of cognitive processing. Further, we assume that qualia and phenomenal consciousness, while real and associated with cognitive process, are not in themselves essential elements of cognitive processes. Thus, we assume without argument that there can be cognitive processes that lack associated quale and are not phenomenally conscious to any agent. What we wish to examine here is a theory of the mark of the cognitive, more specifically, some necessary conditions on a state or process being cognitive.

2.1. Cognitive processes involve non-derived content

A first essential condition on the cognitive is that cognitive states must involve intrinsic, non-derived content. Strings of symbols on the printed page mean what they do in virtue of conventional associations between them and words of language. Numerals of various sorts represent the numbers they do in virtue of social agreements and practices. The representational capacity of orthography is in this way derived from the representational capacities of cognitive agents. By contrast, the cognitive states in normal cognitive agents do not derive their meanings from conventions or social practices. Despite possible interpretationist perversions to the contrary, it is not by anyone's convention that a state in a human brain is part of a person's thought that the cat is on the mat.

We may recall that a significant philosophical effort in cognitive science over the last 20 years has been to provide a theory of non-derived content for cognitive states, by and large within the context of some form of language of thought hypothesis. These approaches typically analyze thought contents in terms of the contents of mental representations, then analyze the contents of mental representations in terms of conditions on causal and/or historical relations. Dretske (1981, 1988), for example, develops a teleoinformational theory according to which an object or event R means some property P in virtue of the fact that R has the function within the system of indicating the presence of the property P. Fodor (1987, 1990), has a theory according to which, roughly, some symbol "X" means X, if it is the case that there is a law connecting Xs to "X"s, and all other laws connecting non-Xs to "X"s exist only in virtue of the existence of the X to "X" law. Searle (1980), has argued that content is an intrinsic causal property of things like brains. Millikan (1984), has a selectionist theory of content. There is, in short, a fairly broad consensus that cognition involves non-derived content. The principal point of contention in this area is what non-derived content is.

The thesis that cognition involves non-derived content should be contrasted with two stronger theses. The first is representationalism, the thesis that cognition involves representations having non-derived content. The second is a language of thought hypothesis, that cognition involves a combinatorial system of representations that possess non-derived content [1]. The theses here may be related in order of increasing logical strictness. The least restrictive theory requires only that the cognitive involve non-derived content, but allows that non-derived content may not come from representations understood as markers in a cognitive economy. Representationalism not only requires non-derived content, but also that such content be carried by markers. So, suppose there is a system with the intrinsic content that roses are red. If this content were distributed over the total cognitive state of the system in such a way that no subparts could be assigned intrinsic content, then representationalism would not be true of that system. Finally, a language of thought hypothesis requires that markers in the cognitive economy have combinatorial syntactic and semantic structure. We believe in both of these stronger hypotheses, but we need not avail ourselves of them in the present context. Instead, all we need presuppose is that cognition involves intrinsic content.

The hypothesis of intrinsic content, of course, has it critics. Most importantly, Dennett, at least at times, suggests that Mike's having a propositional attitude is a matter of some other person, say, Ike, taking a particular kind of stance toward Mike. The content in Mike's propositional attitude would therefore seem to be derived from Ike. The stance view, at least prima facie, conflicts with our thesis that there must be non-derived content for cognition. In our view, stance-taking, interpretativist approaches to content conflict with the idea of non-derived content, to whatever extent they do, by flirting with content instrumentalism. But, we take instrumentalism not to be a live option for a would-be science, such as cognitive science. This instrumentalism gets to be conceptually problematic when we ask how it is that Ike gets the propositional attitude contents he does. The content of Mike's attitude seems to depend on Ike's attitude, but whence comes the content of Ike's attitude? It would seem that there has to be some story of non-derived content for Ike, or for the person who takes an attitude toward Ike, or for the person who takes an attitude toward a person who takes an attitude toward Ike ... If so, then why not simply give a story about non-derived content for Mike? Or, perhaps Ike takes his stance toward Mike in virtue of some real property of Mike. In this case, it would seem that we have a case that Mike really does have non-derived contentful states, initial appearances to the contrary notwithstanding. As far as we can tell, fans of stance-taking and interpretivism cannot plausibly deny the need for non-derived content in cognition.

Another line of criticism of intrinsic content comes from certain areas of psychology. Some psycholinguists, for example, are sometimes moved to the view that parsing tasks are best viewed as purely syntactic or formal processes, hence non-representational. The following considerations, however, militate against such a view. We assume that parsing involves at least phonological, syntactic, and semantic levels. However these levels of processing work, it appears that they must involve some form of inductive inference. When a sound is heard in the environment, the phonological apparatus must conjecture some hypothesis as to what phonemes, if any, were present in the sound stream. The hypothesis will state, among other things, that the sound stream contains a particular set of phonemes. But, if the phonological apparatus is conjecturing some hypothesis, then there must be something meaningful in the area, namely, the hypothesis conjectured. Hypotheses are not purely syntactic items. The parser must also put forth some syntactic hypotheses about the nature of the sound stream in the agent's environment. Again, if there are hypotheses in syntactic processing, there must be meaningful items in syntactic processing. Further, if there is a logical form that the language faculty presents to central cognitive systems, then the language faculty will offer some defeasible conjecture as to the meaning of the sound stream produced in the environment. Again, these conjectures must be meaningful.

Various levels of linguistic processing involve the production of phonological, syntactic, or semantic hypotheses relevant to the incoming sound stream. More than this, the various components of the language processor appear sometimes to traffic in tentative preliminary hypotheses prior to final conjectures. Famous cases revealing this are so-called "garden path sentences," such as

The horse raced past the barn fell.

The idea is that people appear to construct hypotheses about the syntactic structure of a sentence as the words of the sentence come in. By the time one hears "The horse raced past the barn" one has constructed a parse tree corresponding to an entire sentence. When, however, the word "fell" comes in, this hypothesis has to be revised. In this particular case, it generally proves especially difficult for English speakers to assign the correct syntactic interpretation to the sentence.

So, we follow what we understand to be the theory of parsing properly understood, namely, that it is thoroughly committed to trafficking in hypotheses. We suspect that such psycholinguistic resistance to this idea as there is stems, in part, from a division of labor in cognitive science. Philosophers have, for their part, taken up the project of providing theories of non-derived content to figure into cognitive processes, where psycholinguists, and other psychologists, have, for their part, provided theories of cognitive processes that might have contentful states in them. While psycholinguists, and other psychologists, might bracket the problem of non-derived content, or leave it to philosophers, they cannot have it be a non-problem [2]. So, it seems to us, intrinsic content is a legitimate necessary condition on a state or process being cognitive.

Having argued that, in general, there must be non-derived content in cognitive processes, it must be admitted that it is unclear to what extent each cognitive state of each cognitive process must involve non-derived content. That is, it is epistemically possible that cognitive processes involve representations that include a closed set of non-representational functional elements, such as punctuation marks and parentheses. Such items might be included in the language of thought, based on the manner in which they interact with items having non-derived content. If this happens, then cognitive states will to some extent be less than maximally dependent on non-derived content. One might worry that this concession leaves some wriggle room for inserting extracranial states and processes into cognitive processes, but addressing such worries will depend on features of specific cases. There is no evident reason why extracranial cognition is an impossibility; it is, instead, merely a true empirical generalization that there is none.

2.2. Cognitive processes are causally individuated

The second necessary condition is a condition on the nature of processing. This point bears much more elaboration than did the preceding. The old saw is that science tries to carve nature at its joints. Part of what this means is that, to a first approximation, science tries to get beneath observable phenomena to find the real causal processes underlying them; science tries to partition the phenomenal world into causally homogeneous states and processes. Thus, as sciences develop a greater understanding of reality, they develop better partitions of the phenomenological. A range of examples will point out what we are driving at.

In the *Novum Organum*, Francis Bacon proposed a set of methods for determining the causes of things. According to one of these methods, to find the cause of X, one should list all the positive instances of things that are X, then find what is common to them all. As an example, Bacon applies this method to the "form of heat." On his list of hot things, Bacon includes the rays of the sun, fiery meteors, burning thunderbolts, eruptions of flame from the cavities of mountains, all bodies rubbed violently, piles of damp hay, quicklime sprinkled with water, horse-dung, the internal portions of animals, strong vinegar which when placed on the skin produces a sensation of burning, and keen and intense cold that produces a sensation of burning. Bacon conjectured that what was common to these was a high degree of molecular vibration and that the intensity of heat of a thing is the intensity of molecular vibration. Bacon clearly intended to carve nature at its joints, but it simply turns out as a matter of contingent empirical fact that the things that appear hot, or produce the sensation of being hot, do not constitute a natural kind. The rays of the sun, meteors, friction due to heat, body heat, and so forth, simply do not have a common cause. There is no single scientific theory that encompasses them all; the phenomena are explained by distinct theories. Friction falls to physics. Decomposition falls to biology. Exothermic reactions to chemistry.

As a second example, there are the late 19th century developments in the theory of evolution. By this time, Darwin's biogeographical, morphological, taxonomic, and embryological arguments had carried the day for evolution and many biologists had come to accept the theory of evolution by common descent. Despite this, the majority of biologists were reluctant to accept Darwin's hypothesis that evolution is caused primarily by natural selection. In this intellectual environment, biologists returned for a second look at Lamarckian theories of the inheritance of acquired characteristics. In support of their theory, neo-Lamarckians pointed to cases which, in retrospect, proved to be instances in which a mother would contract some disease, then pass this disease on to her offspring *in utero*. Phenomenologically, this looks like the inheritance of acquired characteristics, but, in truth, inheritance and infection involve distinct causal processes. Inheritance involves genetic material in sex cells of a parent being passed on to offspring; infection is the transmission of an alien organism, perhaps via the circulatory system in isolation from the sex cells. To a first approximation, inheritance is a process in the germ line of an organism, where infection is a process in the soma line of an organism. It is only after the true causal differences between inheritance and infection are made out that one can conclude that we have one less instance of the inheritance of acquired characteristics than we might at one time have thought. Throughout the episode, Lamarckians were aiming to carve nature at its joints, but in the absence of a true understanding of the nature of the processes underlying inheritance and infection, these distinct processes had to appear to be the same, both as instances of the inheritance of acquired characteristics.

The cognitive may, therefore, be assumed to be like other natural domains, namely, the cognitive must be discriminated on the basis of underlying causal processes. The point we have been driving at here might be approached in another way, namely, we believe there is more to cognition than merely passing the Turing test. Some of the mechanisms that might be used to pass the Turing test will count as cognitive mechanisms for doing this, while other mechanisms that might suffice will not count as cognitive mechanisms. A computer program might pass the Turing test by having a listing of all possible sensible conversations stored in memory. Such a program, however, would not constitute a cognitive mechanism for passing the test. This is presumably because we have sufficient ground for saying that the look-up table process is not of a kind with the complex of processes that go into enabling a normal human to carry out the same sort of conversation. The look-up table may, for example, answer questions in a constant amount of time for each sentence. Computer chess provides another famous sort of case where behavior can be carried out by both a cognitive and a non-cognitive process. In chess, there is a combinatorial explosion in the number of possible moves, responses, counter responses, and so forth. As a result, it quickly proves to be impractical to examine all the logically possible moves and countermoves. The most powerful chess playing programs, therefore, use special techniques to minimize the number of possible moves and countermoves they have to consider. Nevertheless, there is pretty strong reason to believe that the chess-playing methods currently employed by digital computers are not the chess-playing methods that are employed by human brains. Based on observations of the eye movements of grandmasters during play, it appears that grandmasters actually mentally work through an extraordinarily limited set of possible moves and countermoves, far fewer than the millions or billions considered by the most powerful chess-playing computer programs. The point is not simply that the computer processes and the human processes are different; it is that, when examined in detail, the differences are so great that they can be seen not to form a cognitive kind. The processes that take place in current digital chess playing computers are not of a kind with human chess playing.

2.3. The mark of the cognitive

We maintain a rather orthodox theory of the nature of the cognitive: cognition involves particular kinds of processes involving non-derived representations. So, for example, the representations involved in linguistic processes are handled differently than are the representation involved in visual processes. Within linguistic processing, phonological, morphological, and syntactic representations may be processed differently. Providing a further specification of these sorts of processes is one of the central endeavors of cognitive science. In this regard, cognitive science is much like planetary theory before Newton's discovery of gravity or the theory of electricity and magnetism before Faraday. Cognitive science has relatively little insight into the nature of cognitive mechanisms, hence relatively little can be said about the difference between the cognitive and the non-cognitive. Such assurance as we have that there will, however, be some significant processing condition on the cognitive comes from what we find in mature sciences. Mature sciences seek, and have historically found, taxonomies that carve nature at its causal joints.

As a final note here, it is worth emphasizing that nothing in our mark of the cognitive says anything about the locus of cognition. Nothing in the definition of a non-derived representation essentially requires that they occur only within a brain. Further, nothing about the kinds of processing in the brain conceptually, definitionally, analytically, or necessarily requires that they appear only with a brain. Thus, it is logically possible that there be transcranial or extracranial cognitive processes. Our view is simply that as a matter of contingent fact, when one looks at the processes which embed non-derived representations, such processes happen to occur almost exclusively within the brain. Insofar as we are intracranialists, we are what might be called "contingent intracranialists," rather than "necessary intracranialists."

3. The case for transcranial cognition

3.1. Clark and Chalmers' Tetris and Inga/Otto cases

The central fixtures in Clark and Chalmers' discussion are two putative examples of extracranial cognition. We have mentioned the case of Inga and Otto, but the first example Clark and Chalmers develop is based on three modes of playing the video game Tetris. We shall begin their example of three modes of Tetris play, then turn to their Inga/Otto case.

Clark and Chalmers propose a simple thought experiment involving three hypothetical modes in which one might play the computer game Tetris. In this game, blocks of various shapes descend from the top of the screen to a kind of wall of blocks. The aim is to rotate the oddly shaped, falling blocks in such a way as to form a complete horizontal row of blocks, which can then be eliminated from the wall. Clark and Chalmers describe these three possible modes of play:

(1) A person sits in front of a computer screen which displays images of various two-dimensional geometric shapes and is asked to answer questions concerning the potential fit of such shapes into depicted "sockets". To assess fit, the person must mentally rotate the shapes to align them with the sockets.
(2) A person sits in front of a similar computer screen, but this time can choose

either to physically rotate the image on the screen, by pressing a rotate button, or to mentally rotate the image as before. We can also suppose, not unrealistically, that some speed advantage accrues to the physical rotation operation.

(3) Sometime in the cyberpunk future, a person sits in front of a similar computer screen. This agent, however, has the benefit of a neural implant which can perform the rotation operation as fast as the computer in the previous example. The agent must still choose which internal resource to use (the implant or the good old fashioned mental rotation), as each resource makes different demands on attention and other concurrent brain activity. How much *cognition* is present in these cases? We suggest that all three cases are similar. Case (3) with the neural implant seems clearly to be on a par with case (1). And case (2) with the rotation button displays the same sort of computational structure as case (3), distributed across agent and computer instead of internalized within the agent. If the rotation in case (3) is cognitive, by what right do we count case (2) as fundamentally different? We cannot simply point to the skin/skull boundary as justification, since the legitimacy of that boundary is what is at issue. But nothing else seems different. (Clark & Chalmers, 1998, p. 1)

We take the point here to be that there is no principled difference among these cases, hence that in example (2), we have a case of transcranial cognition. The way to meet such "no principled difference arguments" is to provide a principled difference. Here, our necessary conditions on the mark of the cognitive—some fairly orthodox ideas in the philosophy of cognitive science—come to the aid of common sense. Cognitive processing is, of course, involved in all three cases, but in different ways. (1) and (2) differ in their use of non-derived representations and in the sorts of processes that go on in them, hence (2) does not constitute a "real world" case of transcranial cognition. Consider, first, the matter of non-derived representations. In case (1), the agent presumably uses mental representations of the blocks and their on-screen rotations in cognitive processing. By contrast, in case (2), the blocks on the screen that are physically rotated by pushing the button are not representations at all, either derived or non-derived. They do not *represent* blocks to be fit together; they *are* the blocks to be fit together. Consider, next, the differences in processing in the two cases. It seems to us safe to assume that the process that physically rotates the image on the screen at the push of the button as described in case (2) is not the same as the cognitive process that occurs in the brain. Pushing the button closes some sort of electrical circuit that, at some extremely short time delay, changes the way electrons are fired at the phosphorescent screen of a cathode ray tube. This sort of causal process is surely not the same as any cognitive process, or any fragment of a causal process, in the brain. In case (2), but not case (1), there is muscular activity, and the attendant cognitive processing associated with it, that is involved in pushing the button. The fact that, in case (2), the agent must decide between the two available methods for checking for fit—the method of mental rotation and the method of button pushing—entails numerous other cognitive differences in cases (1) and (2). In case (2) one must actually use the cognitive decision mechanisms, there must be attentional mechanisms the bring the decision mechanisms into play, and there must be memory mechanisms that store for the agent the information about the existence of the button and its use. So, even within the brain where the cognitive action is, there are cognitive differences in processing. Recognizing and controlling for these diverse cognitive factors under experimental conditions is absolutely foundational for psychology. Is it too much to say that the science of cognitive psychology as we know it would cease to exist without attention to such differences in experimentation?

We think that, by being underdescribed, Clark and Chalmers' case (3) simply muddies the waters about the differences between (1) and (2). In particular, by specifying the place of non-derived representations and the nature of processing involved in case (3), Clark and Chalmers can align (3) with (1), or they can align (3) with (2), but they cannot align (3) with both (1) and (2). We leave it to the reader to consider just how the differences between (1) and (2) mentioned above cannot be eliminated in (3).

So, Clark and Chalmers' "no principled difference" argument regarding the three modes of Tetris play runs afoul of our necessary conditions on the cognitive. The case of Inga and Otto does so as well, although more clearly. Recall that Clark and Chalmers contend that, in relevant respects, Inga's memory and Otto's notebook are the same. The essential causal dynamics of the two cases mirror each other precisely. They maintain that Inga's memory and Otto's notebook fall under one natural kind; both are beliefs.

> To provide substantial resistance, an opponent has to show that Otto's and Inga's cases differ in some important and relevant respect. But in what deep respect are the cases different? To make the case *solely* on the grounds that information is in the head in one case but not in the other would be to beg the question. If this difference is relevant to a difference in belief, it is surely not *primitively* relevant. To justify the different treatment, we must find some more basic difference between the two. (Clark & Chalmers, 1998, p. 6)

We, of course, do not want to point to the boundary of the brain as being of any essential theoretical interest in itself. Contrary to Clark and Chalmers' suggestion, however, it is obvious to us that there are important and relevant differences in the two cases. One obvious difference between the two cases involves non-derived content. Where the symbols written in Otto's notebook have merely derived content, the recollection in Inga's brain has non-derived content. Otto's notes do not, therefore, constitute beliefs or memories. There is, in addition, the fact that Inga and Otto carry out distinct processes in coming to arrive at the Museum of Modern Art. Otto's "memory recall" involves picking up the notebook and turning to the appropriate page in the notebook. This involves processes that have no analogue in Inga's memory recall. It seems to us not unreasonable to say that Otto's "memory recall" involves cognitive-motor processing not found in Inga's memory recall. In addition, Otto's "memory recall" involves visual processing for turning to the appropriate page of the notebook and reading the address. Inga's memory recall does not. Further, Inga's memory recall uses some capacity of the brain that Otto has lost due to his Alzheimer's disease. (By hypothesis, that is exactly why he is using the notebook.) It is fairly clear that Inga and Otto use distinct sets of capacities in order to produce similar behavior. They differ in the types of capacities they bring to bear on the problem of finding the address of the Museum of Modern Art and it is not unreasonable to suppose that some of these capacities are cognitive, where others are not. Further, it is hard to see how these differences could be scientifically unimportant or irrelevant, save by adopting a behavioral conception of cognition.

3.2. The coupling argument

Clark and Chalmers defend their analysis of the modes of Tetris play by recourse to a "coupling argument":

> In these cases, the human organism is linked with an external entity in a two-way interaction, creating a *coupled system* that can be seen as a cognitive system in its own right. All the components in the system play an active causal role, and they jointly govern behavior in the same sort of way that cognition usually does. If we remove the external component the system's behavioral competence will drop, just as it would if we removed part of its brain. Our thesis is that this sort of coupled process counts equally well as a cognitive process, whether or not it is wholly in the head. (Clark & Chalmers, 1998, p. 2)

To begin, we may observe that the mere causal coupling of some process with a broader environment does not, in general, thereby, extend that process into the broader environment. Consider the expansion of a bimetallic strip in a thermostat. This process is causally linked to a heater or air conditioner that regulates the temperature of the room the thermostat is in. Expansion does not, thereby, become a process that extends to the whole of the system. It is still restricted to the bimetallic strip in the thermostat. Take another example. The kidney filters impurities from the blood. In addition, this filtration is causally influenced by the heart's pumping of the blood, the size of the blood vessels in the circulatory system, the one-way valves in the circulatory system, and so forth. The fact that these various parts of the circulatory system causally interact with the process of filtration in the kidneys does not make even a prima facie case for the view that filtration occurs throughout the circulatory system, rather than in the kidney alone. So, a process P may actively interact with its environment, but this does not mean that P extends into its environment.

Now, Clark and Chalmers' are under no illusions about the foregoing point [3]. Their argument is somewhat more sophisticated. They contend that the active causal processes that extend into the environment *are just like the ones found in intracranial cognition*. This contention does not founder on the considerations just brought forth. At best, it begs the question against the traditional view that cognition is intracranial: tradition has it that the processes inside the brain are unlike the processes found outside of the brain, but Clark and Chalmers give us no reason to doubt this. But, in truth, things are not the best for Clark and Chalmers. We have good reason to think they are mistaken in their contention that there are actual cases in which active causal processes extending into the environment are just like the ones found in intracranial cognition. Intracranial cognitive processes involve non-derived mental representations, whereas extracranial processes such as rotating blocks on a computer screen, making marks on paper, tying strings around fingers, do not. Further, the causal processes in the brain appear to be nothing like the firing of electrons at a phosphorescing screen, generating marks on a paper through friction with a graphite shaft, generating marks on a paper with a ball point pen, rotating cards in a Rolodex, moving beads up and down on rods in an abacus, or pressing buttons on an electronic calculator. What could be more obvious than that the processes are causally distinct?

The coupling argument brings out the extent to which the mark of cognitive bears on the bounds of cognition. If Clark and Chalmers opt for the simplistic view that anything that is causally connected to a cognitive process is part of the cognitive process, then there is the threat of cognition bleeding into everything. This is sometimes called something like "the problem of cognitive bloat" or "cognitive ooze." These names do justice to the ugliness of the view, but not to its radical nature. The threat is of pancognitivism, where everything is cognitive. This is surely false. If, on the other hand, Clark and Chalmers opt for some more discriminating mark of the cognitive (and they have suggested none), then it is far from clear that this will allow the cognitive to cross the boundaries of the brain without extending to the whole of creation. The common sense view of the bounds of cognition at least has going for it the prospects of a reasonable, principled theory of the mark of the cognitive that has it that cognition is intracranial.

3.3. The explanatory argument

Clark and Chalmers also defend their analyses of Tetris and Inga/Otto by appealing to the idea that this provides for superior explanations of behavior.

> By embracing an active externalism, we allow a more natural explanation of all sorts of actions. One can explain my choice of words in Scrabble, for example, as the outcome of an extended cognitive process involving the rearrangement of tiles on my tray. Of course, one could always try to explain my action in terms of internal processes and a long series of "inputs" and "actions", but this explanation would be needlessly complex. If an isomorphic process were going on in the head, we would feel no urge to characterize it in this cumbersome way. In a very real sense, the re-arrangement of tiles on the tray is not part of action; it is part of *thought.* (Clark & Chalmers, 1998, p. 3, cf. pp. 5–6)

We find this sort of consideration hardly telling, given the pragmatics of explanation and the wide range of appeals we are willing to make to intentional ascriptions. Bear in mind, one may be inclined to say that one's car doesn't want to start or that one's plants are thirsty for water. That's certainly a lot easier than troubling with any real complexity having to do with the internal mechanisms of cars or plants. Perhaps one can get by with a "folk psychology" that uses such explanations, but one should hardly aspire to such an undiscriminating theory for a science. It is worth adding that isomorphism between one process and another does not imply that the two processes are of the same type. Philosophers these days seem not to appreciate that isomorphism is a relatively weak relation. There is, for example, an isomorphism between what water molecules do when they pass from the solid state into the liquid state and when they pass from the liquid state into the gaseous state. There is, none the less, a difference between melting and evaporating. They are distinct processes.

3.4. Donald's theory of exograms

Clark and Chalmers suggest that external symbols constitute a portion of human memory using the simple example of Inga and Otto. Donald, however, supports the view with a much more detailed account of the development of all manner of external representations, including body decorating, grave decorating, sculpture, Stonehenge, hieroglyphics, cuneiform, maps, graphs, and musical scores. Despite the richness of detail in Donald's account, the same sorts of considerations that undermine Clark and Chalmers' analysis of Inga and Otto undermine Donald's theory of exograms as part of the human cognitive architecture. External symbols, in all known cases, lack non-derived content. Further, the sorts of processes that govern external symbols are quite different than those governing internal symbols. In fact, Donald's theory of engrams and exograms makes part of the case for this analysis.

Clark and Chalmers suggest that the cognitive processes involved in internal and external memory traces are in essence the same; Donald, however, provides a relatively more detailed accounting of the sorts of ways in which the processing of exograms differs from the processing of engrams. Some of these points are summarized in Table 1, taken from Donald (1991). We believe that, to a first approximation, Donald is correct in his assessment of these differences and right to draw attention to them. So, for example, when Donald refers to limited size of single entries in memory, this implicitly refers to psychological laws concerning human memory storage which do not hold for forms of external memory storage. When Donald refers to constrained retrieval paths, again, he is implicitly relying on human psychological laws regarding memory retrieval that will not generally hold of external memory storage. So, Donald shares with us the view that internal cognitive processes differ from other processes in the external world. Where he differs with us is in his inattention to, or tacit rejection of, the implication that this diversity in processing provides grounds for thinking that there can be no cognitive science of transcorporeal processes.

3.5. Hutchins' conception of cognitive science

In the introduction to this paper, we mentioned how Hutchins' conception of cognitive science as the study of computations over representational states might be

TABLE 1. Some properties of engrams and exograms

Engrams	Exograms
Internal memory record	External memory record
Fixed physical medium	Virtually unlimited media
Constrained format	Unconstrained and reformattable
Impermanent	May be permanent
Large but limited capacity	Virtually unlimited
Limited size of single entries	Virtually unlimited
Not easily refined	Unlimited iterative refinement
Retrieval paths constrained	Retrieval paths unconstrained
Limited perceptual access in audition, virtually none in vision	Unlimited perceptual access, especially in vision; spatial structure useful as an organizational device

used to argue that human tool use might, in at least some instances, be analyzed as instances of transcorporeal cognition. The argument is that, if cognition is the study of computational operations performed over representational states, and if the tools a human uses involve representations, then we will have an instance of transcorporeal cognition.

As attractive as this argument may at first seem, its weakness should be evident upon closer examination. According to our analysis, Hutchins is studying what might best be thought of as naturally occurring computation, rather than cognition. In the first place, it appears that a principal source of difficulty for Hutchins' analysis is that it is likely that, as a matter of contingent fact, the kinds of computational processes we find operating over external representations, such as marks on a piece of paper, readings of meters and dials, indicator lights, warning lights, and so forth, will turn out to differ from the kinds of computational processes that we find operating over representations in brains. Compare the intracranial computation of the product of 347 and 957 from the computation of this product with pencil and paper. We may assume that there are computational processes at work in both cases, but that these computational processes are different. In particular, the internal processes are cognitive computational processes, where only some of the computational processes in the transcranial cases are cognitive. In particular, it will be only the internal portions of the transcranial computation that turn out to be cognitive. So, we should note that, in the purely intracranial case, there would be no processing corresponding to the visual processing of the marks on pencil and paper found in the transcranial computation [4]. Nor would the intracranial computation involve an analogue of the motor manipulations of the pencil used in the transcranial computation. Nor would the intracranial process involve anything like the process that rubs bits of graphite onto a sheet of paper. Although it is not known in any serious detail how the internal processing works in these diverse cases, we can reasonably maintain that they in fact are not all the same. That is enough for the present purposes of defending the common sense view that cognition is usually intracranial, even when tools are being used.

But, suppose that, by some very unlikely chance, a transcranial process for some task turns out to involve the same computational steps as those found in the brain. Would that be sufficient to establish the existence of a case of transcranial cognition? No. For, of course, there is the second necessary condition on cognition, that it must involve non-derived representations. Marks on a piece of paper, meters, dials, indicator lights, warning lights, and so forth have only derived meaning. They have such meaning as they do in virtue of the fact that we assign them meanings [5]. Insofar as extracranial states lack non-derived representations, they will not count as parts of cognitive processes. To make the point more vivid, imagine a possible future golden age of cognitive science when we have discovered the very computer program of the mind and we have it implemented in an electronic computer in the specific programming language of the brain. In this golden age, we may suppose we have a machine that goes through exactly the same computational processes in finding the product of 347 and 957 as some particular human does. Even in this case, we maintain, unless the computer satisfies the conditions under which a state becomes a non-derived representational state, the computer lacks cognition. Incidentally, this is why, even if one maintains that carrying out the partial products algorithm in one's head is computationally equivalent to carrying out the algorithm on paper, the externalized computation does not count as cognitive.

4. Transcranialist rejoinders

We anticipate two rejoinders to our de facto intracranialism, to our "contingent intracranialism." The first is brief, the second more involved. The multiple instantiability of functional categories is a familiar element in contemporary cognitive science. Applying this to our rather orthodox conception of the mark of the cognitive, one might think that non-derived representations and the sorts of functional processes that are found in the brain might also be instantiated in systems that cross the boundaries of the brain. One might think, therefore, that for all the proposed conditions on the mark of the cognitive show, transcranial cognition is still a live possibility. We agree with this line. Transcranial cognition is a live possibility. Our view is simply that, as a matter of boring contingent empirical fact, transcranial and extracranial cognition are not commonplace. We don't think that transcranial cognition is as prevalent as tool use. Unlike the radical interpreters of tool use, we don't think that transcranial cognition occurs every time a person picks up a hammer, answers the phone, writes in a notebook, logs onto the Internet, drives a car, or loads up a washing machine. To emphasize our point here, one might bear in mind that, while it is possible to build a universal Turing machine out of beer cans and pigeons or the population of China, this has probably never been done and never will be done.

To this point we have assumed that cognitive science will discover a domain in which processes in the world are individuated in terms of their underlying mechanisms. Yet, there are scientific ideas that do not seem to work out this way. Consider the study of animal communication. This covers a wide range of causal processes from pheromone communications in insects, to threat displays in mammals and fish, to territory marking in birds, and natural language in humans. Certainly this is not a discipline that has much divided processes in terms of underlying causal processes. Or consider the study of behavior. Behavior is sometimes divided into such categories as tropisms, fixed-action patterns, and habituation. These divisions are not laid out in terms of underlying causal processes. One might contend, however, that these are not mature sciences and that these categories are symptomatic of this fact. While there may be some truth to this claim, it appears that even mature sciences study such functional/mathematical kinds. Consider one way of defining "heritability" in population genetics [6]. It is possible to provide a mathematical equation relating parental mean for a quantitative trait (such as size) to an offspring mean for that trait. Such a measure of heritability abstracts from the mechanism that brings about that trait. In particular, such a measure of heritability abstracts away from any connection to genes. On such a mathematical measure of heritability, wealth can be highly heritable, even though it has nothing to do with genetic processes. An even more striking case of scientific concepts abstracting from mechanisms is the concept of an oscillator found in undergraduate physics texts. An oscillator of this sort is simply anything that displays periodic behavior. Anything from a pendulum to a hydrogen molecule to a radio signal to a binary star system can be an oscillator. These mathematical/functional categories cut across causal mechanisms. With such examples in mind, one might well ask what reason there is to think that cognitive science will not end up with non-causal categories in the way that these other sciences have. In fact, one might say that this sort of eventuality is just what is currently being entertained by Dennett, Clark, Chalmers, and Hutchins. In Hutchins' case, for example, the mathematics might be taken to be the theory of digital computation. Or, the mathematics might be that of dynamical systems theory. What reason is there to think such a cognitive science will not emerge?

The short answer is that, in contrast to intracranial processes, transcranial processes are not likely to give rise to interesting scientific regularities. There are no laws covering humans and their tool use over and above the laws of intracranial human cognition and the laws of the physical tools. Consider the diversity of possibilities in memory alone. Human memory displays a number of what appear to be law-like regularities, including primacy effects, recency effects, chunking effects, and others. Further, human memory capacities are task sensitive. Memories for images, faces, smells, and lists of words vary in properties of formation and recall. Consider, then, the range of tools humans use as mnemonic aids. There are photo albums, Rolodexes, computer databases, strings around the finger, address books, sets of business cards, bulletin boards, date books, personal information managing software, palmtop computers, hand drawn maps, and lists of "things to do." What are the chances of there being interesting regularities that cover humans interacting with all these sorts of tools? Slim to none, we speculate. There just isn't going to be a science covering the motley collection of "memory" processes found in human tool use.

What holds for memory and tool use likewise holds for perception and tool use. Human perception contains numerous regularities. Weber's law states that the intensity of a sensation is a logarithmic function of the intensity of the stimulus. There are also regularities concerning just noticeable differences among stimuli. But consider systems consisting of humans and their perceptual tools: humans with telescopes, humans with scanning electron microscopes, humans with sonar, humans with mass spectrometers, humans with high-speed cameras, humans with radar detectors, and humans with night vision goggles, to name but a few examples. Surely there will be no interesting regularities that cover the range of human/tool interactions in perception. Perhaps the theory of human perception would just collapse into physics.

And what about human thought and tool use? Human thought displays a number of cognitive regularities in such things as heuristics in probabilistic reasoning and deductive inference. But, compare humans coupled with computing devices. Such systems would seem to have practically boundless capabilities. Humans using digital computers would essentially have the cognitive capacities of computers. The theory of human/digital computer systems would threaten to collapse into the theory of digital computation. But, humans are also capable of using analog computing devices, such as slide rules, which might in principle extend the range of human computation beyond that of digital computation. Surely the capacities of humans with computing devices would be an unscientific motley of capacities.

The situation seems to be this. We do not wish the difference between our common sense view and the radical view of human tool use to degenerate into a debate over how one is to use the word "cognitive." We certainly do not want to define "cognition" as some sort of intracranial process. Being intracranial is merely an accidental, not essential, conceptual, definitional, or analytic, feature of cognition. Nor do we even wish to insist that the word "cognitive" apply only to certain sorts of processes involving non-derived representations, rather than to some possible mathematical/functional domain of the future. Setting aside fruitless debates over who gets to use what words, our claim can be seen to amount to the following. Let our old fashioned, common sense view of cognition delineated in terms of non-derived representations and kinds of processing be labeled *cognition*$_{cs}$. Let the transcranialists define some radical conception of the mark of the cognitive–about which they have been less than maximally forthcoming–and label that radical conception of cognition be labeled *cognition*$_{r}$. Our view is that cognition$_{cs}$ will produce a natural science, where cognition$_{r}$ will not. Insofar as we are doing cognitive psychology, there is little hope of finding a science of transcranial processes. The only sort of new science in the area will be a science of processes that happen to occur intracranially. This is not to say that intracranial processes, *qua* intracranial processes, form a natural kind. They don't. Rather, when we get to finding natural kinds that don't belong to physics or chemistry, they will as a matter of contingent fact be intracranial, rather than transcranial.

5. Conclusion

The question of the bounds of cognition in tool use seems to us to raise many of the same issues as does the question of the mark of the cognitive. Find out what separates the cognitive from the non-cognitive and you will go a long way toward determining the bounds of cognition in tool use. While it is unclear just exactly what constitutes the mark of the cognitive, we can see enough of this mark to see that cognitive processes are likely to be a brain-bound processes. Whatever is responsible for non-derived representations seems to find a place only in brains. Further, the sorts of processes that occur within brains seems to share certain sorts of regularities that they do not share with systems consisting of brains coupled with tools. Finally, systems consisting of brains coupled with tools would seem to form such a motley collection that they will not form the basis for any significant scientific theorizing. It seems to us that transcranial theories of cognition appeal to those who accept a behavioral principle regarding the cognitive: anything behaviorally equivalent to a cognitive system must be a cognitive system. When such a conception is clearly presented, cognitive scientists generally and correctly reject it. A rejection of such phenomenological conceptions is standard in science. There is no science of heat producing processes, rather physics, chemistry, and biology offer distinct theories of the production of heat. There is no science of the "inheritance of acquired characteristics." Rather, there is inheritance and infection. Insofar as we aspire to create a cognitive science, it seems reasonable to suppose that the science of cognition will resemble the science of physical, chemical, and biological processes. Such, at least, is our bet.

Acknowledgements

We thank Colin Allen, John Barker, Gerard O'Brien, Joel Pust, and two anonymous referees for helpful comments on earlier drafts of this paper.

Notes

[1] Here we ignore the possibility of drawing a distinction between imagistic and linguistic combinatorial systems of representations and a distinction between "Classical" and "Non-Classical" combinatorial systems of representation.

[2] As we shall explain, we think that Hutchins is among the cognitive scientists who have overlooked the need for non-derived representations.

[3] Although Van Gelder and Port (1995, p. 13) might be.

[4] It is, of course, possible that purely mental computation might involve visual imagery and that visual imagery might share processing mechanisms and resources with visual processing. Nevertheless, visual imagery processing and normal visual processing of information from the physical environment will not be identical.

[5] They, of course, have what has been called "natural meaning," as in smoke means fire, but this is not the "non-natural meaning" that is the stuff of cognition. Compare, for example, Grice (1957).

[6] See Brandon (1996) for an accessible discussion of this concept.

References

BRANDON, R. (1996). Phenotypic plasticity, cultural transmission, and human sociobiology. Reprinted in BRANDON, R. *Concepts and methods in evolutionary biology*. New York: Cambridge University Press.

CLARK, A. (1997). *Being there*. Cambridge, MA: MIT Press.

CLARK, A. & CHALMERS, D. (1998). *The extended mind*.

DENNETT, D. (1996). *Kinds of minds*. New York, NY: Basic Books.

DONALD, M. (1991). *Origins of the modern mind*. Cambridge, MA: Harvard University Press.

DRETSKE, F. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT Press.

DRETSKE, F. (1988). *Explaining behavior*. Cambridge, MA: MIT Press.

FODOR, J. (1987) *Psychosemantics*. Cambridge, MA: MIT Press.

FODOR, J. (1990). *A theory of content and other essays*. Cambridge, MA: MIT Press.

GRICE, H. (1957) Meaning. *Philosophical Review*, *66*, 377–388.

HUTCHINS, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.

MILLIKAN, R. (1984). *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.

SEARLE, J. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, *3*, 417–458.

VAN GELDER, T. & PORT, R. (1995). It's about time: an overview of the dynamical approach to cognition. In R. PORT & T. VAN GELDER (Eds). *Mind as motion: explorations in the dynamics of cognition*. Cambridge, MA: MIT Press.