

The Bounds of Cognition. Frederick Adams and Kenneth Aizawa. Malden, Massachusetts: Wiley-Blackwell, 2008, 216 pages, \$74.95 hardcover.
Reviewed by Justin C. Fisher, Southern Methodist University.

Fred Adams and Kenneth Aizawa have long been the loyal opposition in the debate about extended cognition. Contemporary humans regularly use external devices to process information. Many of us store telephone numbers in our cell phones rather than our brains. Alzheimer's patients use trusted notebooks to store all kinds of information (Clark and Chalmers, 1998). Expert Scrabble players continually reorganize their letters to more quickly see possible words they might play (Kirsh, 1995). Fans of extended cognition have held that the information processing performed partly within such external devices is enough like traditional cases of cognitive processing that it also deserves to be called "cognitive processing."¹ Adams and Aizawa have been two key figures to stand against this tide, arguing that we should instead view these as mere cases of external tool use, and that, at least for the time being, we should reserve the term "cognitive processing" for processes that occur inside creatures' heads.

The Bounds of Cognition compiles and updates Adams' and Aizawa's attempts to defend against this tide, and it gives the authors the opportunity to go on the offensive themselves, and give careful arguments for why we should stick to their more conservative construal of "cognitive processing." This book does very well to give the reader a thorough overview of the state of play in the debate over extended cognition. As a consequence, the present paper is as much a critical commentary on the whole debate as it is a review of the book itself.

Requests for reprints should be sent to Professor Justin C. Fisher, Hyer Hall 207, Southern Methodist University, P.O.Box 750142, Dallas, TX 75275, fisher@smu.edu.

¹ The many fans of extended cognition include van Gelder and Port (1995), Clark and Chalmers (1998), Rowlands (1999), Gibbs (2001), Noë (2004), Wilson (2004), and Rockwell (2005).

Is this Debate Merely Terminological?

Adams and Aizawa rightly note that the question of whether cognitive processing extends into external devices hinges upon what exactly you mean by “cognitive processing.” If by “cognitive processing” you mean the process of gathering, storing, and manipulating information, then the answer is clearly yes: humans use many external devices in this sort of processing. If you instead construe “cognitive processing” to include only processes that have conscious phenomenology, or that display characteristic patterns of human neural activation, then it’s equally clear that the answer should be no—cell phones, notebooks, and scrabble racks almost certainly don’t have a conscious phenomenology of their own, nor do they display the characteristic patterns of human neural activation.

To someone not invested in these debates, there’s a strong inclination to read “cognitive processing” as a term of art, and to doubt whether there is any fact of the matter as to which, if either, of these proposals captures what “cognitive processing” really means. All parties agree that many humans use external devices to manipulate information, and all parties agree that this external information manipulation differs in various ways from neural information processing. When we agree on all this, then it seems we agree on everything important, and it’s hard to see why it should matter whether we choose to label these external processes as “cognitive processing” or not. Hence, there is a real concern that this might turn out to be a merely terminological debate, no more substantive than debates about whether cat claws deserve to be counted as a sort of “fingernail.”²

² A similar concern is pressed by Susi, Lindblom, and Ziemke (2003).

Adams and Aizawa offer two responses to this concern, but neither is satisfying. First, they note that “no party to the debate over extended cognition wants the matter to degenerate into a mere terminological dispute” and they muster “at least some evidence that the other parties to the dispute wish to avoid a mere battle over linguistic usage” (p. 83). This sounds like an argument from wishful thinking. It might be good for the careers of the various parties if this debate were substantive and not just terminological, but this doesn’t give the reader any positive reason to believe it actually is substantive and not just terminological. A psychologist friend of mine put her response to this argument quite succinctly: “You might *like* it to be substantive but tough noogies!”

Second, Adams and Aizawa suggest that there is a particular theory of cognitive processing “implicitly at work in a lot of cognitive psychological research” (p. 84). The research that they have in mind is research on particular aspects of cognition happening in human brains, e.g., Miller’s (1956) finding that human short-term memory typically can hold just seven items. It is indisputable that many cognitive psychologists have been especially interested in the sorts of cognitive processing that occur in human brains. However, it’s not at all clear that this research implicitly presupposes that the peculiar features of human brains must be displayed by *all* cognitive processes.

Here’s a useful analogy (one we will revisit several times). Many ornithologists have carefully studied the features that are distinctive of flight in birds, including birds’ use of feathers to modulate airflow. However, no ornithologist implicitly presupposes that all flight must involve feathers. Just as ornithologists happily accept that flight can occur in many ways that don’t involve the characteristic features of flight in birds, most cognitive psychologists would (or at least should) happily accept that cognitive processing can

occur in many ways that don't involve the characteristic features of cognitive processing in human brains. So, Adams and Aizawa's observation that cognitive psychologists have focused primarily upon cognitive processing in human brains is, by itself, no reason to think that cognitive psychologists are implicitly committed to any particularly restrictive construal of "cognitive processing."

It would have served Adams and Aizawa well to take concerns about the substantiveness of this debate much more seriously. It would have been especially nice if they could have offered (and defended) some sort of positive story regarding which features of our use of the term "cognitive processing" determine the range of things we are committed to counting as instances of cognitive processing. Had they given us such a story, this would have set up a clear game plan for resolving the debate: first, they could look at our usage to discover our implicit commitments regarding the meaning of the term "cognitive processing," and then they could check to see whether these commitments are consistent with counting various cases of tool use as instances of "cognitive processing." Sadly, they don't offer a clear story about this, and hence both the rules of the game and its broader significance are left unclear.

Victory Conditions

With the rules of the game left unclear in this way, it's hard to say what exactly is at stake. We can, however, say something in abstract about what each side might do to win.

Adams and Aizawa aim to establish the conclusion that, at least in contemporary humans, cognitive processing happens only in our heads. One clear way to establish this conclusion would be to identify some sort of necessary condition for cognitive

processing, and to argue that, at least in contemporary humans, this condition is met only by things in our heads. This is Adams and Aizawa's offensive game plan.

In contrast, extended cognition theorists aim to establish that, for some contemporary humans, our cognitive processing occurs partly outside our heads. One clear way to establish this conclusion would be for the extended cognition theorist to identify some sort of sufficient condition for cognitive processing, and to argue that this sufficient condition is met by some of our existing external devices. Adams and Aizawa need to thwart any such attempt by the extended cognition theorist. This is their defensive game plan.

Unfortunately, Adams and Aizawa don't keep distinct which side needs to identify necessary conditions and which side needs to identify sufficient conditions. Instead they point out that, ideally, any theorist would be able to identify both necessary and sufficient conditions. This sets the bar for each team much higher than it would need to be to resolve the present debate. It also makes the dialectic at various points more confusing than it would need to be. For example, Adams and Aizawa acknowledge shortcomings in one of their attempts to identify a necessary condition, but argue that their opponents have nothing better to offer (p. 72). This complaint is misplaced because extended cognition theorists have no immediate need to identify a necessary condition for cognition—a sufficient condition would suffice for their immediate purpose of showing that some cases of external tool use should count as instances of cognitive processing. Consider again the analogy of flight. It is difficult—perhaps impossible—to specify interesting *necessary* conditions for flight. Despite this, the Wright brothers did very well to demonstrate that machines can meet a *sufficient* condition for flight. In much the

same way, extended cognition theorists might hope to demonstrate clear cases of extended cognition without having any aspirations to articulate necessary conditions for cognition. In this discussion, and elsewhere, the dialectic would have been much more clear had the two teams' victory conditions been kept distinct.

Now that *we*, at least, are clear on each team's victory conditions, we may proceed to the game. We'll begin by considering the two relatively brief attempts Adams and Aizawa make at playing offense. Then, after a halftime report, we'll consider their much more protracted attempts at playing defense.

Playing Offense I: Non-Derived Representations

In their first attempt at playing offense (Chapter 3), Adams and Aizawa argue that a necessary condition for cognitive processing is that it “involves non-derived representations, representations that mean what they do independently of other representational or intentional capacities” (p. 31). The intuitive idea behind this offense is that the states of external devices have representational content only because they have been assigned representational content by people, whereas states in the brain have representational content that is not derived from the representational capacities of anything else. For this offense to work, Adams and Aizawa would need to establish both that non-derived representations are indeed a necessary condition for cognitive processing, and that the various external devices that contemporary humans use don't satisfy this necessary condition.

Adams and Aizawa list thoughts, experiences, and perceptions as “paradigm cases of items bearing non-derived content” (p. 32). However, it's not clear that even these “paradigm cases” satisfy their definition of non-derived representations as

“representations that mean what they do independently of other representational or intentional capacities” (p. 31). For example, one plausible view has it that the representational content of our thoughts and experiences is derived from the representational content of perceptual states that produce them. Another plausible view has it that many perceptual states are intrinsically relational, and represent various aspects of one’s environment in a way that displays their relations to other represented aspects. These views might be wrong, but they aren’t obviously wrong, and they shouldn’t be ruled out without argument. And on these views, even Adams and Aizawa’s “paradigm cases” would turn out not to meet their definition, for each would depend, in part, upon other representational capacities. This bodes poorly for hopes of establishing that non-derived representations are a necessary condition for cognitive processing.

Things get worse as we consider other clear instances of cognitive processing. It’s quite plausible that the representational content of long-term memory states derives from the representational content of the perceptual states and/or short-term memory states that cause them. It’s also quite plausible that the representational content of various imaginational states derives from the representational content of associated perceptual states. Again, these views might turn out to be wrong, but they aren’t obviously wrong. And if they are right, this means that other clear instances of cognition don’t meet Adams and Aizawa’s proposed necessary condition either.³

³ Adams and Aizawa might respond to these concerns by leaning on the weasel-word “involves.” Even if long-term memory *contains* only derived representations, it still *involves* non-derived representations, as these played a role in producing the long-term memories. It might well be a necessary condition that cognition “involves” non-derived representation in this broad sense of “involves.” However, the Alzheimer’s patient’s notebook also involves non-derived representations in this same broad sense. Hence, this maneuver won’t help Adams and Aizawa to rule out this notebook as being cognitive.

Thus far, we've seen that Adams and Aizawa have fallen short of showing that including non-derived representations is a necessary condition for something to count as cognitive processing. Even if they could somehow establish this, they would also need to establish that existing external devices don't involve non-derived representations.

To evaluate this claim, we would need an account of what makes something a non-derived representation. In lieu of offering an account of their own, Adams and Aizawa gesture approvingly towards various approaches in "naturalized semantics" including Fred Dretske (1981, 1988), Jerry Fodor (1987, 1990), and Robert Cummins (1996). Shockingly, they don't even consider the question of whether external devices like an Alzheimer patient's notebook might satisfy the various accounts they gesture towards. At first blush, it seems likely that many existing devices might satisfy each of these accounts. For example, Cummins attributes non-derived content to anything whose relational structure is isomorphic to that of something else, a condition that hand-drawn maps can very easily meet—indeed maps are a textbook example of representational isomorphism! Similarly, it is quite plausible that some existing devices will satisfy Fodor's and even Dretske's accounts. Hence, Adams and Aizawa fall far short of showing that our external devices lack non-derived content.

Adams and Aizawa's first attempt at offense falls short of their victory conditions. They fail to establish that non-derived content is a necessary condition for cognitive processing—instead it is quite plausible that much clearly-cognitive in-the-head processing involves only derived content. They also fail even to argue, much less to establish, that existing external devices lack non-derived content—instead it is quite

plausible that some of these devices will actually satisfy the accounts of non-derived content they gesture towards.

Playing Offense II: Cognitive Mechanisms

In their second attempt at playing offense (Chapter 4), Adams and Aizawa attempt to identify a different necessary condition for cognitive processing, namely “the kinds of mechanisms that are involved” (p. 57). The intuitive idea behind this offense is that we might discover a particular kind of mechanism that human brains contain—cognitive mechanisms—and then argue that external devices that don’t contain such mechanisms shouldn’t count as doing cognitive processing. The victory conditions for this second line of offense match those of the first: Adams and Aizawa need to establish both that certain kinds of mechanisms are necessary for cognitive processing, and that the various existing devices that contemporary humans use don’t involve these kinds of mechanisms.

Adams and Aizawa think that the best way to identify the mechanisms that are distinctive of cognition in general is to look at the mechanisms at play in human cognition. The bulk of this chapter (pp. 61-70) reads like a psychology textbook, detailing generalizations that cognitive psychologists have found regarding human brains—e.g., human short-term memory can hold only seven items, and unaided human memory is most likely to remember the first and last items displayed in a long list. These are the sorts of “mechanisms” that Adams and Aizawa take to be distinctive of cognition.

Adams and Aizawa acknowledge that “a definition of the cognitive exclusively in terms of normal human cognition is too parochial” (p. 70) and that a full understanding of cognition must somehow allow for abnormal human cognition, animal cognition, and cognition in aliens. Their positive response to this concern is to suggest that we might

define “cognitive processing in general” as being anything that bears a *family resemblance* to the cognitive mechanisms at play in normal human brains (pp. 72-73).

Of course, there will be great challenges in determining what sort of resemblance to normal human brains something must bear in order to count as “cognitive processing.” Adams and Aizawa will need to be quite liberal regarding the sorts of resemblance they will admit if they are to encompass cases of abnormal human cognition, animal cognition, and the many sorts of alien cognition that we would surely want to accommodate.

Notice that there is a great risk here, for one needn't be all that liberal to see a strong resemblance between ordinary cognition and proposed cases of extended cognition. In Clark and Chalmers' (1998) much-discussed example, the Alzheimer's patient Otto uses his notebook to perform tasks for which the ordinary human Inga would use biological memory. Otto and Inga both use perceptual mechanisms to gather information; both translate this information into a durable passive format so that they can carry it with them for long periods of time; both have procedures by which they can retrieve that information on demand; and feeding either person false perceptual input will lead them to misadventures, as would directly modifying their memories by hypnosis or surreptitious inking. If all these similarities count as “resemblance” then it will turn out that Otto's notebook does bear a family resemblance to ordinary human long-term memory, and hence will meet this necessary condition for cognitive processing after all.

Adams and Aizawa must strike a fine balance to achieve their victory conditions. They need to construe “resemblance” liberally enough that “resembling” human cognitive mechanisms is a necessary condition for cognitive processing. And they need

to construe “resemblance” conservatively enough that existing external devices won’t count as “resembling” human cognitive mechanisms. Adams and Aizawa need to identify a sort of resemblance that is broad enough to encompass *all* the possible ways that aliens’ memories might work, yet narrow enough to exclude Otto’s notebook. They offer no positive reason to think they can strike this balance successfully. Worse, there are strong reasons to think they can’t.

The sorts of resemblance that Adams and Aizawa dwell upon are ones involving the features in their “psych-textbook” overview mentioned above. Most of these features are *limitations*—e.g., our short-term memory can hold *only* seven items, and we do *less well* recalling earlier items than items presented more recently. They clearly think of these limitations as being the sorts of features that should be counted in whatever similarity comparisons we use to determine whether or not to label something as a “cognitive process.” For example, they take it to weigh against counting Otto’s notebook as cognitive that the notebook doesn’t suffer the same sorts of “recency effects” as does biological memory (p. 137).

Since Adams and Aizawa count similarity to human cognitive limitations as a relevant sort of resemblance, they are committed to saying that, *ceteris paribus*, creatures that aren’t bound by human cognitive limitations are less likely to count as doing cognitive processing. But this is highly counterintuitive! Most of us are impressed by the cognitive abilities of (real and fictional) creatures that surpass our own cognitive limitations. We are impressed by autistic savants, by people with photographic memories, by Mr. Spock, and by God. The fact that these creatures transcend our own cognitive limitations weighs *in favor* of thinking that they employ cognitive processing,

not against it. This suggests that Adams and Aizawa are wrong to count suffering from cognitive limitations like ours as a necessary condition for cognitive processing.

Here's another way to make this point. Imagine two Alzheimer's patients, PenOtto and PencilOtto. PenOtto uses a pen to record his memories in permanent ink, whereas PencilOtto writes in pencil which smudges and becomes illegible over time. PencilOtto's notebook suffers a recency effect—it is more likely to correctly recall recently presented items than older items. Hence, by Adams and Aizawa's standards, PencilOtto's notebook is closer to counting as a cognitive process than is PenOtto's notebook. And if these notebooks are fancy enough (Adams and Aizawa allow that some future devices might be fancy enough to count as cognitive—see pp. 25-29) then they would say that PencilOtto's notebook is cognitive while PenOtto's isn't, even though the only difference between them is that PenOtto's "non-cognitive" memory works *better* than PencilOtto's "cognitive" memory. This absurd result shows that there is something wrong with Adams and Aizawa's approach.

To diagnose where their approach went wrong, it is helpful, once again, to return to the analogy of flight. Imagine an approach that defines flight as anything that bears a family resemblance to bird flight. This approach would yield the same sorts of absurd results as Adams and Aizawa's approach to cognition: e.g., something might be deemed "non-flying" because it lacks feathers or can accelerate more quickly than birds. The problem is that flight (in general) is not something that we should analyze in terms of underlying mechanisms. Instead, flight apparently needs to be understood functionally, perhaps as a scheme for locomotion through atmosphere. As birds, insects, airplanes, and helicopters illustrate, flight can be accomplished by many different kinds of mechanisms,

and these mechanisms needn't bear any interesting sort of family resemblance to one another.

There is strong reason to think that cognition, too, should be understood functionally, in terms of accomplishing certain sorts of information-processing goals, and not in terms of the particular mechanisms that any one species happens to use to accomplish these goals. We admire creatures that surpass our own cognitive limitations as being *especially good cognizers*, just as we would admire creatures that surpass birds' aerodynamic limitations as being *especially good flyers*. We also freely allow that animals, aliens, and deities might have cognitive processes implemented by mechanisms quite different from those in human brains, just as we allow that many creatures and devices might fly, despite being implemented by mechanisms quite different from those in birds. These observations weigh against Adams and Aizawa's proposal that cognition should be linked to particular underlying mechanisms, and in favor of the proposal that cognition, like flight, should be understood in general functional terms.

This blunts Adams and Aizawa's second line of offense. This offense presumed, wrongly, that cognition should be understood in terms of similarity to the particular mechanisms at play in human brains. Instead, it seems, cognition should be given some sort of general functional characterization, and Adams and Aizawa offer no compelling reason to deny that existing external devices would satisfy such a characterization.

It is worth stressing that Adams and Aizawa's failure to achieve victory does not amount to an immediate victory for the other team. We have not determined exactly what functional specification should be given for "cognitive processing," nor have we determined whether any of our existing external devices would meet this specification.

We can imagine more liberal and more conservative functional specifications that one might propose for “cognitive processing,” and existing external devices will meet only the more liberal of these proposed specifications. Depending upon which proposed specification we end up adopting, either team might yet win.

The analogy to flight gives some grounds for pause here. It’s quite doubtful that there is a single correct functional characterization for flight. Flying squirrels fulfill many of the functions characteristic of flight (steering, balance, controlled descent) but it’s not clear whether they should count as flying or just gliding. Blimps also provide controlled locomotion through atmosphere, but it’s not clear whether they deserve to count as flying or just floating. We can easily imagine useful functional specifications for “flight” that include or exclude either of these cases. It seems likely that there’s no deep fact of the matter which of these captures “what flight really is,” and which serve only to capture scientifically interesting “flight-like” phenomena. Once we agree on all the details about how flying squirrels and blimps locomote through the air, the debate about whether or not to call this locomotion “flight” is merely terminological.

At this point, your alarm bells should be ringing, saying that the parallel debate about “cognitive processing” is probably merely terminological too—there are surely many possible functional characterizations of “cognition-like” phenomena, and no deep fact of the matter as to which of all the possible characterizations is correct. Let’s construe those alarm bells as the half-time buzzer. When play resumes, the extended mind theorists will take the offensive and try to establish that, whatever understanding of “cognitive processing” we end up accepting, it should be liberal enough to encompass at least some existing external devices.

Halftime Report

It's often said that the best defense is a strong offense. If Adams and Aizawa had managed to mount a strong offense, then their defensive task would be quite easy. A strong offense would make the reader confident that their conclusion was correct, and all that would remain would be to point out where their opponents' arguments had gone astray. If Adams and Aizawa had a winning position, you would expect them to spend most of their efforts developing their offense, and you would expect their defense to come as a brief afterthought. If they instead had a losing position, you would expect their offense to flounder, and you would expect them to get bogged down in prolonged defensive battles, quibbling with their opponents' attacks, waiting for the hammer to fall.

The statistics suggest that Adams and Aizawa have a losing position. They spend only 45 pages playing offense (pp. 31-75), and then 97 pages playing defense (pp. 76-173). Just to be clear, these 97 pages do not defend premises needed for their offense. Instead these are 97 pages of gritty trench warfare—they go piecemeal through various positive arguments their opponents give and defend against each of them. It's hard to win when your offense receives half as many resources as your defense.

Our detailed assessment of Adams and Aizawa's offense agreed with what the statistics would predict: their offense floundered far short of their victory conditions. Having failed to achieve victory themselves, they now must fend off each individual attempt to establish that something-or-other is an instance of extended cognition. The best they can hope for now is a tie—if they can blunt their opponents' offense, then it will turn out that neither team had an effective argument for its conclusion, and we will be left not knowing which conclusion is correct. But if Adams and Aizawa allow even one line of their opponent's attack to go through unblunted, the game will be lost.

Playing Defense

It's beyond the scope of the present paper to slog through 97 pages of philosophical trench-warfare. Instead, I'll make a few general comments about Adams and Aizawa's defense, and then focus upon what I take to be their opponents' strongest attack.

Adams and Aizawa are to be commended for attempting to find common threads among their opponents' many arguments, as this helps to map the space of possible argumentative strategies, and it helps to reveal the strengths and weaknesses of these strategies.

Unfortunately, one often gets the feeling that Adams and Aizawa's approach is too Procrustean—that they had to distort their opponents' views to make them fit their preconceived categories. For example, (in Chapter 6) they interpret many of their opponents as committing what they call “the Coupling-Constitution Fallacy”—the fallacy of presuming that anything a brain causally interacts with must therefore be involved in its cognitive processing. They offer a menagerie of nine quotes (pp. 89-91) that “strike us as presenting the shortest, clearest articulations of the move from a causal or coupling kind of dependency to a constitution or supervenience kind of dependency” (p. 89).

Among these nine quotes, I see only three (van Gelder, 1995, p. 373; Haugeland, 1998, p. 217; Noë, 2004, p. 221) as clearly endorsing this sort of move. The other quotes either don't endorse this move (Rockwell, 2005, p. 46), allow that additional conditions may be required to make this move (Clark and Chalmers, 1998, p. 2; Menary, 2006, p. 331), or don't make this move at all (Rowlands, 1999, p. 23; Wilson, 2004, p. 194; Noë, 2004, p. 220). Adams and Aizawa are right that this move is fallacious, regardless of how many or few authors make it. However, one suspects that most of the authors Adams and Aizawa consider are clever enough not to commit a simple fallacy like this, and hence

that Adams and Aizawa might be so busy fighting straw men that they will overlook the stronger arguments that these theorists can and do make.

Concerns about straw men seemed especially pressing in Adams and Aizawa's discussion of what I took to be extended cognition theorists' strongest line of attack. This attack is based on a principle proposed by Clark and Chalmers (1998), which Adams and Aizawa call the "Parity Principle":

Parity Principle: 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* (so we claim) part of the cognitive process. (Clark and Chalmers, 1998, p. 8)

Notice that the Parity Principle is not a full analysis of the term "cognitive process." Instead, it is just a sufficient condition for recognizing further cognitive processes on the basis of ones we can already recognize. Adams and Aizawa state no objection to the Parity Principle as a sufficient condition. Indeed, they gloss it as saying "[I]f some transcorporeal process is cognitively just like an intracranial cognitive process, then the transcorporeal process is a cognitive process too" (p. 28) and they cheerfully grant that this is "something like a logical or conceptual truth" (p. 133).

Hence, Adams and Aizawa have allowed Clark and Chalmers to achieve the first of their two victory conditions: establishing a sufficient condition for something to count as a cognitive process. So it all comes down to whether extended cognition theorists can also achieve their second victory condition: establishing that some existing external devices meet this sufficient condition. Adams and Aizawa's defensive strategy is to consider Clark and Chalmers' two flagship cases, and argue that neither satisfies this sufficient condition.

They begin with Otto the Alzheimer's patient (pp. 134-141). To apply the Parity Principle, they would need to consider some creature who employs an intracranial process that is otherwise just like Otto's use of his notebook, and they would need to ask whether that intracranial process would count as a cognitive process. They choose a surprisingly poor candidate for this role: the ordinary human Inga. They then argue at great length that Inga does not have an intracranial version of Otto's transcorporeal process. Instead Inga has an intracranial process that displays many features—like recency effects—that Otto's process does not display. They then hastily conclude that Otto's notebook therefore does not meet this sufficient condition for counting as a cognitive process.

What Adams and Aizawa *should have* concluded from this is that Inga wasn't a relevant case to consider in applying the Parity Principle to Otto. Instead, they should have considered a creature who, unlike Inga, had an intracranial process that *was* like Otto's transcorporeal process in all relevant ways. And the question they should have asked is whether *this creature's* intracranial process would count as a cognitive process.

Suppose we someday encounter intelligent space aliens, and find that they have Otto's level of competence at navigating the world. We can, for example, tell them the location of a museum, and they can later navigate to the location we told them. Suppose we dissect one of these aliens and discover that the way in which they "remember" such information is actually by using tiny pencils inside their heads to inscribe this information onto tiny slips of paper which are later read by tiny eyes when the information is to be retrieved. Unlike Inga, these aliens do seem to have an intracranial version of Otto's transcorporeal process. Now, the relevant question is whether this process inside the

aliens' heads should count as a cognitive process or not. If it would count as a cognitive process, then by the Parity Principle, so too should Otto's use of his notebook.

My own inclination is to say that the imagined aliens' memory *would* count as a cognitive process. To be sure, the aliens' memory would work quite differently from ordinary human memory, and would have strengths and limitations quite different from those of ordinary human memory. But these seem to be irrelevant differences of fine-grained implementation. What's being implemented strikes me as a clearly cognitive capacity—the capacity to remember things.

Here's another way to drive this point home. Suppose psychologists someday discover that the way that humans remember song lyrics is actually by having internal mechanisms inscribe those lyrics onto tiny sheets of neural tissue, and then using tiny optical sensors to read those inscriptions at the time of retrieval. I doubt that we actually will discover that this is how human memory for song lyrics works, but for all I know *we could* discover that this is how it works, at least for some people. What should we conclude if we do someday discover this? I suppose Adams and Aizawa might maintain that, if we discover this, then people don't actually *remember* song lyrics but instead use tiny internal tools to keep track of them. However, I think the much more natural conclusion in response to this discovery would be to say that we *do* remember song lyrics, and that we do so via an internal cognitive process that happens to involve inscription and optical reading.

I have proposed two sorts of creatures (aliens and conceivable humans) that store information by inscribing it on inner slips of tissue. Unlike Inga, these creatures have intracranial processes that are very similar to Otto's transcorporeal process. Hence,

unlike Inga, these creatures are good candidates for applying the Parity Principle to Otto. For both creatures, it is very plausible that we should count their intracranial processes as cognitive processes. Hence, by the Parity Principle, we should count Otto's use of his notebook as a cognitive process too. Victory for the extended mind theorists!

I can see three potential places for Adams and Aizawa to resist this argument. First, they might rethink their acceptance of the Parity Principle as "something like a logical or conceptual truth" (p. 133). Second, they might insist that there are relevant differences between Otto and the creatures I described. One potentially relevant difference is that Otto's optical sensors are hooked into rich systems for perception, visualization, and motor control, whereas the internal optical sensors in the creatures I described are most naturally construed as having much more limited uses. Another difference is that Otto can think about his information retrieval as a deliberate motor activity, whereas the other creatures are most naturally construed as not being able to do this. Such maneuvers might force us to shift, yet again, to aliens whose intracranial processes are even more like Otto's transcorporeal process. Third, they might hold that the creatures' intracranial processes (especially as elaborated in response to the second concern) shouldn't really count as fully cognitive processes. Insofar as these processes involve what looks, for all the world, like a rich internal system of motor manipulation and optical perception, maybe these processes don't really deserve to be counted as cognitive, after all.

I have no idea how this debate would play out, but it strikes my ears as a merely terminological dispute about which sorts of intracranial processes really deserve to be called "cognitive." I was disappointed, however, that Adams and Aizawa didn't apply the Parity Principle correctly, and hence didn't follow this interesting debate out.

Perhaps they can be excused for not choosing an appropriate case for applying the Parity Principle to Otto. After all, Inga is the most commonly discussed comparison to Otto, and it takes a bit of charity (though not all that much) to recognize that she isn't really the right case to use here.

Much less excusable is Adams and Aizawa's poor treatment of the other case they consider: a Tetris player who rotates pieces on-screen in order to more quickly recognize potential ways in which those pieces might fit onto an uneven surface below. Considering this case, Clark and Chalmers (1998, p. 1) describe three possible individuals:

- (1) there is the analog of Inga, who uses mental rotation to predict how a given piece will fit;
- (2) there is the analog of Otto, who can use mental rotation or use an external rotate button and visual display to predict how a given piece will fit; and
- (3) there is the analog of my space alien, a person in the cyberpunk future who can use mental rotation or use a neural implant that performs the same processing as in case (2), but entirely intracranially.

Clark and Chalmers argue that, even though case (3) involves further cognitive abilities than case (1), there is no reason to suppose these abilities are any less cognitive than case (1). They then apply their Parity Principle to cases (3) and (2), which differ only with respect to where this extra information processing occurs: "If the rotation in case (3) is cognitive, by what right do we count case (2) as fundamentally different?" (Clark and Chalmers, 1998, p. 1).

Adams and Aizawa quote this argument in full, but then interpret it in a way that completely ignores the role that case (3) is supposed to play:

We take the point here to be that there is no principled difference among these cases, hence that there is a cognitive equivalence between the three cases. [...]

The way to meet such “no principled difference” arguments is to provide a principled difference. (p. 142)

Adams and Aizawa go on to point out various cognitive differences between (1) and (2)—e.g., that case (2) involves muscular control processes while case (1) does not. Having found a “principled difference” between cases (1) and (2), they then dismiss Clark and Chalmers’ argument without even mentioning case (3) (pp. 142-143).

This is a serious misreading of Clark and Chalmers. Clark and Chalmers explicitly acknowledge that there is a cognitive difference between case (1) and the others—in case (1) the agent has available only one means (old fashioned mental rotation) for determining how a piece will fit, whereas in cases (2) and (3) the agent must choose between two available means (old fashioned mental rotation and a faster alternative). Clark and Chalmers make clear that they do not include case (1) for direct comparison to case (2), but instead to justify thinking that case (3) involves a cognitive process—a fancier, but no less cognitive, version of the cognitive process in case (1). They then use the Parity Principle to argue that case (2)—the transcorporeal version of case (3)—should be counted as cognitive as well. By ignoring Case (3), Adams and Aizawa completely miss the explicit structure of Clark and Chalmers’ argument, and the role that case (3) and the Parity Principle were supposed to play in it.

What if Adams and Aizawa had read this argument appropriately? I think the upshot would be much the same as in our discussion of the space aliens above. Adams and Aizawa’s best recourse would be to highlight potential differences between cases (3) and (2)—e.g., the different roles that visual processing systems might play in these two cases. This might force Clark and Chalmers to build more structure into case (3) in order to make clear that it really is an intracranial version of case (2). Adams and Aizawa could

then dig in their heels and insist that the intracranial processing in case (3), thus elaborated, shouldn't really be labeled as "cognitive" afterall.

At this point your alarm bells should be ringing once again, saying that we're probably entering into another merely terminological dispute about which intracranial processes we should count as "cognitive." Let's construe those alarm bells as a final buzzer, and move on to a post-game analysis.

Post-Game Analysis

I would like to highlight four points that I think deserve additional consideration by the participants in this debate. First, there is the shadow hanging over the entire game: the worry that this might be a merely terminological debate. Why take the game into overtime when we could all agree that Otto's notebook does cognitive-processing-liberally-construed but not cognitive-processing-conservatively-construed, and go home now? It would behoove all parties to make clear what they think could serve to make a particular construal of "cognitive processing" be the correct one, and/or why it should matter which construal of "cognitive processing" we end up adopting.

Second, there is Adams and Aizawa's most valuable player: the strong intuition that there is *some sort* of important difference between the representational capacities of brains and the representational capacities of artifacts. I argued above that their particular attempt to cash out this intuition in terms of "non-derived representations" won't work. A good challenge for all parties is to explore the strong intuitions underlying this argument, and to determine the extent to which they are correct and the bearing this should have on the debate.

Third, there was the extended cognition theorists' most valuable player: Clark and Chalmers' Parity Principle. I argued above that this principle has more muscle than has often been recognized. Once we find appropriate intracranial processes for comparison, there is a real prospect for using this principle to draw conclusions about extended cognition. However, I noted that this may end up transforming an intractable debate about which *mundane transcorporeal* processes we'd want to count as cognitive into an equally intractable debate about which *fanciful intracranial* processes we'd want to count as cognitive. It would behoove all parties to consider these issues carefully.

And fourth, there was my own most valuable player: the analogy between cognition and flight. This analogy suggests that, in our attempts to understand cognition in general, it may be a mistake to dwell too much upon the particular features of cognition in humans. This analogy also suggests that we might want to seek some sort of general functional characterization for cognition. However, this raises hard questions about which functional characterization(s) we should adopt, and whether liberal characterizations and conservative characterizations might each turn out to be useful in their own ways, as they are for "flight." It would be good for all parties to explore these issues, and the analogy to flight may be a useful guide as they do.

References

- Clark, A., and Chalmers, D. (1998). The extended mind. *Analysis* 58, 7-19.
- Cummins, R. (1996). *Representations, targets, and attitudes*. Cambridge, Massachusetts: MIT Press.
- Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, Massachusetts: MIT Press.
- Dretske, F. (1988). *Explaining behavior*. Cambridge, Massachusetts: MIT Press.
- Fodor, J. (1987). *Psychosemantics*. Cambridge, Massachusetts: MIT Press.

- Fodor, J. (1990). *A theory of content and other essays*. Cambridge, Massachusetts: MIT Press.
- Gibbs, R. (2001). Intentions as emergent products of social interactions. In B.F. Malle, L.J. Moses, and D.A. Baldwin (Eds.), *Intentions and intentionality* (pp. 105-122). Cambridge, Massachusetts: MIT Press.
- Haugeland, J. (1998). Mind embodied and embedded. In J. Haugeland (Ed.), *Having thought* (pp. 207-237). Cambridge, Massachusetts: Harvard University Press.
- Kirsh, D. (1995). The intelligent use of space. *Artificial Intelligence*, 73, 31-68.
- Menary, R. (2006). Attacking the bounds of cognition. *Philosophical Psychology*, 19, 329-344.
- Miller, G. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 81-97.
- Noë, A. (2004). *Action in perception*. Cambridge, Massachusetts: MIT Press.
- Rockwell, T. (2005). *Neither brain nor ghost*. Cambridge, Massachusetts: MIT Press.
- Rowlands, M. (1999). *The body in mind*. Cambridge: Cambridge University Press.
- Susi, T., Lindblom, J., and Ziemke, T. (2003). Beyond the bounds of cognition. In K. Forbus, D. Gentner, and T. Regier (Eds.), *Proceedings of the 25th Annual Conference of the Cognitive Science Society* (pp. 1305-10). Mahwah, New Jersey: Lawrence Erlbaum.
- van Gelder (1995). What might cognition be, if not computation? *Journal of Philosophy*, 91, 345-381.
- van Gelder, T., and Port, R. (1995). It's about time: An overview of the dynamical approach to cognition. In R. Port and T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp. 1-43). Cambridge, Massachusetts: MIT Press.
- Wilson, R. (2004). *Boundaries of the mind: The individual in the fragile sciences*. Cambridge: Cambridge University Press.