Phenom Cogn Sci (2010) 9:605–618
DOI 10.1007/s11097-010-9185-8

The holy grail of cognitivism: a response to Adams and Aizawa

Richard Menary

Published online: 6 November 2010

Abstract Adams and Aizawa (2010b) define cognitivism as the processing of representations with underived content. In this paper, I respond to their use of this stipulative definition of cognition. I look at the plausibility of Adams and Aizawa's cognitivism, taking into account that they have no criteria for cognitive representation and no naturalistic theory of content determination. This is a glaring hole in their cognitivism—which requires both a theory of representation and underived content to be successful. I also explain why my own position, cognitive integration, is not susceptible to the supposed causal-coupling fallacy. Finally, I look at the more interesting question of whether the distinction between derived and underived content is important for cognition. Given Adams and Aizawa's concession that there is no difference in content between derived and underived representations (only a difference in how they get their content) I conclude that the distinction is not important and show that there is empirical research which does not respect the distinction.

Keywords Cognition · Representation · Content · Cognitive integration · Extended mind · Distributed cognition

Introduction

Adams and Aizawa, in this volume, are arguing against a position that I do not hold: "a cognitive process extends from an individual's brain into that individual's body and a set of environmental tools at time t_0." (2010b) In particular, I do not think the following: "So, even Menary's strong hypothesis that cognitive agents are never

The research for this paper was conducted whilst on research leave awarded by the Faculty of Arts. It is supported by the Australian Research Council Discovery grant: *Embodied Virtues and Expertise*.

R. Menary (✉)
Department of Philosophy, The University of Wollongong, Wollongong, New South Wales, Australia
e-mail: rmenary@uow.edu.au

without their cognitive processes extending into tools is not enough to avoid the coupling-constitution fallacy." (Adams and Aizawa 2010b) What I propose to do is to show where I think the debate has gone wrong.

In the first section, I look at the plausibility of Adams and Aizawa's cognitivism. In the second section I explain why cognitive integration is not susceptible to the supposed fallacy. In the third section I look at the more interesting question of whether the distinction between derived and underived content is important for cognition. When reviewing Adams and Aizawa's arguments they don't provide an argument for why all cognitive content must be underived, in fact they assume it. There are examples of empirical research on cognition that don't respect Adams and Aizawa's stipulation that all cognitive content is underived. If Adams and Aizawa persist in holding to their stipulation they will either have to deny the validity of the empirical research, an entirely unconvincing move, or provide some ad hoc explanation for why the empirical research fits their stipulation. That move, I shall show, is a zero sum game.

Searching for the cognitivist holy grail

Adams and Aizawa (2010b) make much play of the virtues of cognitivism as a model of cognition.[1] However, they present a strong and weak version of that model:

Weak cognitivism: cognition involves the processing of representations.

Strong cognitivism: all cognition involves the processing of representations with underived content.

Note that even the strong version is ambiguous, because it allows that there may be some underived representations that are processed, for example, in an extended process that also involves conventional representations (for example see Menary 2006 and Clark 2010). In which case there really is no case left to answer. However, Adams and Aizawa are inconsistent in the strength of their claims as they also assert that cognitive content must be underived (in this world at least). If they hold to that stipulation—that cognitive content must be underived[2]—then they cannot, as they do in this issue, allow the possibility of weak cognitivism (in this world).

The weak version is consistent with most of the examples of integrated cognition that I present in *Dimensions of Mind* (Menary 2010). It only conflicts with examples that do not involve the processing of representations. In fact, the claim that Clark and Menary argue against cognitivism is somewhat strange since both of them allow that representations play an important role in cognitive explanations, they just don't believe that all cognition involves the processing of representations. So their positions are quite consistent with a weak cognitivism, that cognition often involves the processing of representations, but not always. Furthermore, at least Menary (2007, 2009) gives conditions for when some phenomenon ought to be considered a

[1] We should not confuse cognition with cognitivism. Cognitivism is supposed to tell us what cognition is. Standardly, it is the idea that cognition involves the processing of representations. Historically, it has been associated with a methodological individualism which takes cognitive processes and representations to supervene on physical states of an individual. However, one can be a cognitivist and not be an individualist.

[2] I don't take seriously talk of full stops in the language of thought (see Adams and Aizawa 2001, 2010a).

representation, which is something that Adams and Aizawa have neglected to do. This brings me to the holy grail part of the cognitivist enterprise.

What is a mental or cognitive representation? There is no philosophically or empirically agreed upon account of what makes something a cognitive representation. This is quite a stunning fact. Imagine genetics without a model of genes, this is the position in which cognitivism finds itself.[3] Adams and Aizawa are not alone in having no criteria for determining when something is to be counted as a representation (oh the irony!) This brings us to the second leg of the grail quest, a theory of content. Adams and Aizawa make much play of a purported naturalistic theory of content for cognitive representations. However, they have no convincing theory available to them, and this explains why they do not attempt to explain how cognitive representations get their contents. This is also quite stunning. In my 2006 paper which Adams and Aizawa discuss, I tried to work out their arguments in terms of a naturalistic theory of content, a causal one, but they deny that they are committed to that theory and that any of my criticisms of that theory count against them.

Adams and Aizawa need to commit to a naturalistic theory of content to get their arguments to work. They have no criteria for when something is a representation, and no theory of naturalistic content determination, but their arguments rest upon these accounts being complete. They are not and we are only offered a promise that they will be. So the grail is empty and awaits a champion (or two) to discover its missing contents.

Causal coupling

Adams and Aizawa suppose that I, along with Clark and many others, are susceptible to the causal-constitution fallacy: "If a cognitive agent causally interacts with some object in the external world in some "important" way—if that agent is coupled to an object—then that agent's cognitive processing is constituted by processes extending into that object" (2010b) something has clearly gone wrong here. The outline of cognitive integration that I gave in *Dimensions of Mind* (Menary 2010) should, I hope, be enough to dispel the idea that I think that cognitive processes extend into tools (I confess to not understanding what is meant by this notion, but see below for further discussion).

Something has certainly gone horribly wrong in the literature concerning so called 'coupling' arguments for the extended mind—as Adams and Aizawa call them. Simply put, Adams and Aizawa's objection to the extended mind (and other neighbouring positions such as cognitive integration) is that they mistake causal relations for constitutive relations. They identify this as a mistake (or fallacy); causal relations should not be confused with constitutive relations. One thing to note, before I continue, is that Adams and Aizawa give no indication of how we are supposed to make the distinction. An account of the difference between causation and constitution would be helpful here, but there is none forthcoming (see Hurley 2010, Ross and Ladyman 2010 for a critique along these lines).

[3] See Ramsey 2008 for a book length treatment of this problem.

I take the following to be the commitments of those working in the extended mind framework: cognitive processes are causal processes (cognitive processes cause certain effects, presumably cognitive states and behaviour), the extended mind holds that certain kinds of coupled processes (processes that span brain, body and local environment) are cognitive processes. So it is certainly true that coupled processes are taken to be cognitive processes but this is a statement of identity. I think it may be accurate to say that a potentially misleading way of stating this has entered the literature—the environment drives cognitive processes—and maybe there are ways of stating it that fall prey to the causal-constitution fallacy, but Adams and Aizawa are wrong to think that *this is* the central argument for the extended mind. It is an entirely different matter to provide evidence for, as I have elsewhere put it, hybrid cognitive processes (see Menary 2006, p. 334) and I consider this to be the point of extended, distributed and embodied cognition as an empirical movement in the cognitive sciences.

A brief bit of historical recounting might help here. The coupled process as cognitive process idea arose out of the empirical work done by those using dynamical modelling techniques in the cognitive sciences. This included work by philosophers such as Tim Van Gelder, linguists such as Robert Port (Port and Van Gelder 1995), psychologists such as Esther Thelen (Thelen and Smith 1996) and Michael Spivey (Spivey 2007), and roboticists such as Randall Beer (Beer 1995). Anyone familiar with this kind of work should be unsurprised by the thrust of what I am saying. Dynamicists focus on the coordination dynamics of interacting components of a system. The global behaviours of a system are a product of the coordinations between system components (which may themselves be complex systems). The system is constituted by its components and their interactions and its successful functioning requires all its proper parts to be in good working order.

The interesting thing about dynamical work on cognition is that the interacting components of cognitive systems are sometimes located spatially outside the central nervous system of an organism. This is where Adams and Aizawa and sympathisers will prick up their ears, but the point at issue here is not a causal-constitution fallacy (of the form I am causally related to object X, therefore X is part of me), but is an issue about whether those components should be counted as proper components of cognitive systems (I take it that this is also Rupert's worry).

The dynamicist must here show why we should think of these components as system components (usually they are just interested in explaining, or modelling the causal dynamics between components), they might give an evolutionary[4] (phylogenetic) account of the system. They might give an ontogenetic[5] account of the system as part of the developmental history of an organism. Whatever explanation they give, it must allow that there is a synchronic account of how the system components are so coordinated that they produce system behaviour.[6] This is the job that cognitive integration does.

In the light of these remarks let's look a little closer at Adams and Aizawa's account of causal-coupling to see where they have gone wrong. "If a cognitive agent causally interacts with some object in the external world in some "important" way—

[4] See Sterelny's paper (Sterelny 2010).
[5] See Stotz's paper (Stotz 2010).
[6] For a more detailed exploration of these issues see Menary 2007 chapters 2, 5 and 7.

if that agent is coupled to an object—then that agent's cognitive processing is constituted by processes extending into that object." (2010b) This is precisely what I was objecting to in 2006. It assumes an already formed cognitive agent with, presumably, internal representations manipulated by computational processes, who just happens to interact with the environment. Adams and Aizawa are here leveraging their argument on a premise that I think is false—that we can consider cognitive agents independently of their environments (apart from the inputs from and outputs to the environment). The point of the ecological turn in cognitive science is to show that there are deep consequences for our conception of cognitive agents once we consider their embodiment and embedding in an environment. The whole purpose of Adams and Aizawa's arguments is to deflate those consequences, but you can't do so just by assuming that agents are cognitive independently of their environmental embedding.

One only needs to take a look at the papers by Sterelny (2010), Stotz (2010) and Sutton et al. (2010) in this issue, let alone the many papers and books by Clark, Wheeler, Rowlands, Wilson and others, to recognise that even those who don't find the extended mind a very compelling position, such as Sterelny, don't argue from a position like this.

Here is another example of how Adams and Aizawa have failed to understand the point. Their example of Otto before he acquires a notebook and Otto after he acquires a notebook is otiose (2010b). It is not the notebook that is important, it is Otto's capacity to write down and retrieve information as written sentences that matters. This might be in a notebook, or it might be in some other medium.

Indeed, Sterelny argues that our cognitive capacities are deeply dependent upon the cognitive niche such that the cognitive capacities of the agent are transformed by interactions with the niche. It would be closer to the mark to talk of the niche 'extending' into the agent,[7] in the sense that the agent acquires new representations and new capacities by being 'coupled' to or scaffolded by the niche. For Sterelny, there is no cognitive agent without embedding in the cognitive niche.

Then there is the puzzling metaphysics implied by "extending into the object". As I mentioned above, I have no idea what this means but it sounds positively mediaeval. Adams and Aizawa think that they clarify things by saying: "Take a burning match to a piece of paper and the combustion process will extend into the paper, but take the same match to a steel wrecking ball and it will not extend into the ball. What will extend and what will not depends on what is coupled." (2010b) I confess to not understanding this kind of metaphysics of things extending into other things. I really wonder whether there is anyone who believes that cognition is a process in the head that then gets extended out into the body and then penetrates into objects.

There are two things one might say at this point: the first is that Clark and others, when in a more metaphorical mode say things like "Cognition leaks out into body and world." (2008, p. xxviii) maybe this kind of metaphorical flourish gives the misleading impression that Clark thinks that the mind is in the head and then (miraculously) extends out into the body and into objects. To be fair to Adams and Aizawa, one can find examples of such metaphorical talk, but that is to ignore the real arguments and examples, most of them empirical, which do not support the

[7] But I don't take the metaphor seriously, see the section on transformation in Dimensions of Mind for a serious account of the relationship between a cognitive agent and the developmental niche.

straw man characterisation that Adams and Aizawa present to us.[8] I take it that the term 'extended' here is a synonym for 'distributed', the mind is extended or distributed across the brain, body and world,[9] in a way reminiscent of the term 'distributed representation'—representations may be distributed across large areas of the brain and not localised to any particular point. It does not, on my understanding, mean that mind or cognition is first in the head and then extends or moves somehow out into the world and into things (as I was at pains to point out in both Menary 2006, 2007).

The other thing one might say, and the one I imagine that Adams and Aizawa will be keen to say, is that even if I claim not to be guilty of a coupling fallacy, I am nevertheless guilty of such a move. This seems to be the point of their discussion of my rejection of the causal-coupling fallacy in my 2006 paper (2010b). They quote me as saying that I am committed to their causal-coupling fallacy, so here is my response. Firstly, it is clear from the page before the quoted section (where I briefly explain the commitments of cognitive integration) and the page and a half after it (where I continue that explanation) that I am not arguing for the position that: "a cognitive process extends from an individual's brain into that individual's body and a set of environmental tools at time t_0." (2010) What I do say is that a clear way of motivating cognitive integration was first presented by Mark Rowlands which he calls the manipulation thesis: "Cognitive processes are not located exclusively in the skin of cognising organisms because such processes are, in part, made up of physical or bodily *manipulation* of structures in the environments of such organisms." (Rowlands 1999 p. 23)

It is very clear, to me at least, that the manipulation thesis does not depend upon any kind of causal coupling (in Adams and Aizawa's sense) and does not make any claims about cognitive processes extending from brains into bodies and tools. What I do say in the 2006 paper is that internalists like Adams and Aizawa deny that the manipulations of external vehicles (which might well be representations) are cognitive processes. Certainly, integrationists do think that they are and the coordination of internal and external vehicles and processes is at least part of that story, but it is not the only one. Adams and Aizawa, strangely, fail to notice the discussion of four different theses that are supposed to motivate integration: the manipulation thesis, the hybrid mind thesis, the transformation thesis and the cognitive norms thesis.[10] Therefore, I think that Adams and Aizawa have been looking too hard for a fit to their 'fallacy' schema and not hard enough at the actual commitments of integrationists. I haven't 'abandoned' an argument that I never made!

"Suppose instead, that what Menary has in mind is that what we are supposed to do is explain why the manipulation of the notebook (X) and the brain processes (Y) are so coordinated that they together constitute the process of remembering (Z) which causes further behavior."[11] (Adams and Aizawa 2010b) They then go on to say: "But wait! The internalist is not going to accept the obligation to explain this,

[8] Personally, I have never used these metaphors.
[9] I have already explicitly made this point in Menary 2009, fn p. 41 and discuss it in Menary 2007, p. 55–60.
[10] See Menary 2007 and this issue for discussion.
[11] Actually, this is just what I do say in 2006 p. 334.

because the internalist rejects the idea that the manipulation of the notebook and the processes are so coordinated as to constitute the process of remembering." (Adams and Aizawa 2010b) However, I didn't claim that the internalist was supposed to explain that. It should be quite obvious that they cannot. Instead I was asking Adams and Aizawa to first of all get the position that they are attacking right before attacking it. This is something that they are still failing to do.

Adams and Aizawa have failed to hit the target again. It is the inscriptions, the publically accessible and meaningful representations in the social niche that integrationists are interested in. They are not interested in claims like "cognitive processes in the head extend into the pencil and make it cognitive" This is a straw creation of Adams and Aizawa's. I have repeatedly stated my position that cognitive integration starts from the following positions:

- That we are actively embodied in a socially constructed cognitive niche and
- That phylogenetically and ontogenetically there is good evidence to suppose that we acquire cognitive capacities to create, maintain and manipulate the shared cognitive niche and
- That this has led to the development of hybrid cognitive systems where the bodily manipulation of vehicles (some of them representational) in the niche involves the coordination of neural, bodily and environmental vehicles.
- Cognitive processing sometimes involves these online bodily manipulations of the cognitive niche, but also collaborative thinking and offline private thinking.

Consequently, I propose that I am not committed to the view that cognition is first in the head and then gets extended into tools. Nor does it follow that I am committed to the idea that pencils can think for themselves.

Here is an open empirical question with which Adams and Aizawa are free to disagree (on empirical grounds): are there socially integrated cognitive systems of the sort I have here described? I take it that they would say no (although I would be happy to find out that this is a misinterpretation of their position). Given their cognitive internalism (that only the brain does cognition) it seems unlikely that they could accept the cognitive integrationist framework. However, they appear to agree with me that our cognitive capacities, at least those that we are born with, are limited (Adams and Aizawa 2010b). Yet they are unwilling to take the next logical step in the argument, that humans have adapted to their niches by actively structuring them and that structuring has led to changes and improvements in our wonderfully plastic brains. This has led to the evolution of the hybrid human mind; thinking can be offline and private, or online by actively manipulating the cognitive niche.

Content

Adams and Aizawa have stipulated that the mark of the cognitive is representation with underived content. The actual reasons they give in support of this are scarce. In their 2001 introduction of the condition, it is simply assumed that because for the last 30 (or more) years philosophers have been trying to provide a naturalistic theory of content, cognitive science must be working with such a notion (Adams and Aizawa 2001, p. 48). Firstly, it seems odd to claim that a science uses a principle of

demarcation because of the work of a few philosophers (maybe all of us secretly hope that this is the case). Secondly, it is odd to use the demarcation as a principle of defining cognition, since the demarcation was introduced by John Searle (1980) precisely to attack classical cognitivism and to show that the symbolic representations beloved of classical cognitivism, including cognitive psychology and AI, were examples of derived representations. In the next subsection, I will show why Adams and Aizawa have made an important concession and how this concession contradicts their stipulation. For, they have conceded that derived and underived representations have the same contents. If the content of a sentence of natural language in my head and a sentence that is spoken or written down have the same content, then it is unclear why only the former should be allowed to be counted as cognitive.[12]

In the second subsection, I show how their position depends upon a completed theory of naturalistic semantics to which they do not have access. Indeed most of the content that is relevant to cognition will be formulated in terms of conventional representations such as sentences of a natural language and depends upon meaning that is publically created and shared. In the final subsection I argue that cognitive scientists don't care about Adams and Aizawa's stipulation. There is empirical research in cognitive neuroscience that does not respect the stipulation, therefore Adams and Aizawa ought to give it up.

An important concession

In their joint paper for this issue, Adams and Aizawa respond to some of my concerns about their usage of the intrinsic/extrinsic underived/derived distinctions as applied to representations and content (Menary 2006). Their concern is how something gets its content, the pertinent question is: how does the way that a representation get its content matter to whether or not a representation can count as a cognitive representation? It turns out not to matter much, because Adams and Aizawa really think that there is no difference in kind between underived and derived content. Both underived and derived representations have the same content, they just get those contents determined differently (2010b).

This is what I have been asking Adams and Aizawa to clarify (Menary 2006, p. 337, Menary 2009, p. 40, Menary 2010, p. 16–17) My concern with Adams and Aizawa's use of the distinction has always been that they appear to be making underived and derived content too distinct and that there would consequently have to be differences between the contents of representations in the brain and representations in the environment. Their stipulation that all cognitive content must be underived depends upon there being a difference. Since there is not a difference in content the difference must lie somewhere else (in the method of content determination).

I will now take it that Adams and Aizawa ought to concede that there is no difference in content between my thought that 'the harbour Bridge looks beautiful in the sunlight this morning' and my utterance to you that 'the Harbour Bridge looks

[12] Assuming that we avoid the question begging assertion that only when the sentence is in the head should it be counted as cognitive.

beautiful in the sunlight this morning.' Indeed I might think that very sentence to myself in my head before uttering it to you. If Adams and Aizawa are happy to accept this conclusion, then we really have no disagreement, because we have an example of thinking in natural language (English in this instance). Similarly it seems obvious to me that a Venn diagram that I am imagining now has the same meaning as the Venn diagram that I am drawing on the page now.

However, this ecumenical turn contradicts some of their earlier claims, where, for example, they "presuppose that thoughts have non-derived content, but that natural language has merely derived content." (Adams and Aizawa 2008 p. 35)[13] Given that commitment in conjunction with their stipulation, they cannot really help themselves to my example above. That is why I have always found Adams and Aizawa's position so strange, it leads to clear problems and even contradictions. Given that they very clearly stipulate that no representations with merely derived content can count as cognitive and that natural language has only derived content, we are left with the puzzling claim that my conscious linguistic thoughts (what I am saying to myself) are not really mental or cognitive at all (see Menary 2009 p. 40).

Is conscious thinking in natural language not really thinking at all? Of course they can't really be endorsing that (can they?) Therefore, it is hard to see what the importance of the distinction is here. They might claim that thinking in language in the head involves underived content, but spoken sentences are derived from them, so only when natural language is in the head can it be counted as cognitive. Perhaps they might just assume this, but then they are assuming the truth of the conclusion that they should be arguing for.

They could abandon the distinction and allow that thinking is often conducted in natural language or with concepts that are learned in natural language, and that we think mathematical thoughts with natural language words for numbers and so on.[14] However, in doing so they would be giving the green light to the integrationist project. For integrationists are interested in how we come to think with public symbols by mastering, or internalising, those symbol systems. It makes no difference whether we are thinking privately to ourselves or publically by working through a problem on paper. Given that there is no important difference in content this cannot be used as grounds for denying that public thinking ever takes place. As I put it in 2006, if this is true then the distinction is vacuous (at least when considering what is and isn't cognitive).

Therefore, it ought to be clear that Adams and Aizawa should give up on the claim that derived content cannot be cognitive content, at least in the case of natural language. We'll need to keep this in mind for the ensuing discussion: (1) Adams and Aizawa concede that there is no distinction in content between my thought that "the Harbour bridge looks beautiful in the sunlight this morning" and my saying it out loud. (2) That consequently thinking in natural language is possible. (3) That once this is allowed there is no good reason for holding that there cannot be public thinking in natural language.

There is a real dilemma here, either having derived content precludes a representation from being cognitive or it does not. If Adams and Aizawa think

[13] They do point out that this position is not unproblematic, but endorse it anyway.
[14] This would be to agree with Dehaene, see my Dimensions of Mind this issue.

that derived and underived representations can have the same contents then they do not think that there is any difference in content, but there must be some difference in content for the distinction to have any force.

I will make my position clear: meaning is public and determined by the patterns of interaction between language users; it's the same meaning whether or not language is being used privately or publically. It follows that even if the image of a Venn diagram is derived from a naturalistic perceptual process—I see the Venn diagram and the image is stored for later offline thinking in imagination—the conventional meaning of the Venn diagram will need to apply to the image just as much as the diagram on the page. All the discussion hinges on this point. In 2006, I was trying to point out to Adams and Aizawa that the image of the Venn diagram when internally processed, is processed in terms of its publically determined meaning. So it doesn't matter at all that the image is causally produced by perceptual processes, or that it is neurally realised.[15] What matters is that Venn diagrams have a certain publically determined meaning. This is crucial for Adams and Aizawa's case, since if I am right then when we think about Venn diagrams we are thinking with meanings that are publically determined, but Adams and Aizawa deny that this is possible.

At times in their paper Adams and Aizawa (2010b) appear to think that the importance of the distinction is at the level of representational vehicle: representations in the brain are not conventionally determined and representations outside the brain are. So then this is just a version of the differences argument that Adams and Aizawa have laboured painfully over. Representations in the brain have different properties than representations outside the head, therefore representations outside the head are not cognitive. This is a bad argument as I have just shown, because the contents are the same and it is the content that is doing the work.

I suspect that it just turns out that Adams and Aizawa are making a merely trivial claim (which no one I take it denies): a pencil drawing of a Venn diagram on a piece of paper cannot literally be found in the brain—you wouldn't find pieces of paper in my head if you opened it up and had a look. It doesn't follow from this claim that the meanings of Venn diagrams aren't the same whether they are represented as images neurally or as circles on a page. It cannot turn out that way for Adams and Aizawa by their own admission, so it turns out that all the hullabaloo about the underived/derived distinction amounts to nothing. Adams and Aizawa are free at this point to retreat to their weaker position that cognition involves the processing of representations. That's fine, but it's no objection to Clark's central example of Otto and his notebook. It may be an objection to epistemic actions, where there are no representations being manipulated and Adams and Aizawa highlight the Tetris example as one of these. This is a vanishingly small objection and even here we should be wary.

Naturalistic semantics

I will not engage with Adams and Aizawa in a goldilocks argument about whether or not underived cognitive content can be just right or not (2010b). The point is simple,

[15] I assume that there will be a complex explanation of how images are produced in the brain, but this is irrelevant for the present discussion.

there is no naturalistic theory of content determination which gives an account of how richly conceptual cognitive representations can have the richly conceptual contents that they do.[16] Therefore, if you hold to a theory of cognitive content that is simply, or only, naturalistic you will be at a loss to account for this richness of representations. Strangely enough, so called conventional[17] theories of content determination for conceptually rich representations are not so constrained. Adams and Aizawa need to stump up with a naturalistic theory that works, otherwise they will continue to be left with a gaping hole in their underived version of cognitivism.

Adams and Aizawa need to show us how a naturalistic semantics can give us a maximally rich conceptual semantics for cognitive states. In their paper they do not commit to giving such an account, they even concede that it might turn out that naturalistic semantics will give us a very minimal cognitive semantics and that the richness will be made up for by the semantics of public representational systems. I think that this is the more likely scenario, indeed this is the position that I think we are in now. If cognitive semantics is to be explained, in part, in terms of the semantics of public representations then Adams and Aizawa's stipulation fails.

Naturalistic theories of content are good at showing us how intentional directedness can be a natural relation. Take Millikan's bee dance example, the consumer bees are intentionally directed at the waggle dance for some end—"they react to them appropriately" as Millikan puts it (Millikan 1984, p. 71). The intentional icons do not imply content because they can be entirely extensional. However, if you want intensionality, you need propositions so that you can have something represented under an aspect. Then you get intersubstitutivity of co-referring terms salva veritate effects. Superman is the strongest man in the world and Clark Kent (by identity) is the strongest man in the world. 'Superman' and 'Clark Kent' are intersubstitutive only if one knows that Superman is identical to Clark Kent.

There are then two compelling explanatory options[18] the first is that mental representations with propositional content get their contents (e.g. Clark Kent, Superman, man, strongest, etc.) by the combination of atomic representations that get their content by some naturalistic relation such as asymmetric dependency, teleology, or indicator semantics. The problem is that none of those theories can do that job. None of these naturalistic relations can adequately explain how the atoms in the following get their contents: π is the ratio of a circle's circumference to its diameter (and which is the same as the ratio of a circle's area to the square of its radius).

Alternative two is to say that the propositional contents of thoughts like these are inherited from the propositional contents of the natural language sentences which they express. That is to say that we often think in language—it doesn't follow from this that all thinking is in language. Alternative two allows that there may be some content that is naturalistically determined, but that there is also content that is publically, or socially determined. It follows that some of our thought contents will have publically determined meanings—notably for conceptually rich representations such as π.

[16] Dan Hutto has long championed such a position, see his recent 2008 book for discussion.
[17] I prefer the term public meaning to conventional meaning.
[18] There may be others I shan't try to be comprehensive about variations on a theme.

Cognitive scientists don't care about the stipulation

"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." (Adams and Aizawa 2001, p. 43).

This is how Adams and Aizawa start their first foray against transcranial theories of cognition. They go on to say that: "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." (p. 43) The difference between the mental arithmetic and the pen and paper arithmetic is that the former "makes serious demands on one's memory capacities" (p. 43). This difference is enough for Adams and Aizawa to discount our application of the partial products algorithm to numbers on the page as being a matter of cognitive processing. There are several things that are instructive about this example, first of all the algorithm is a conventional procedure for doing basic arithmetic—it is something that we learn and learn how to apply. Secondly it applies to symbolic numerals of a number system—in this case Arabic—which are also part of the same public system of mathematics. Finally there is the really interesting part, we can apply the same procedure (process) to numerals (representations) online by creating and bodily manipulating numerals in physical space or we can apply the same procedure offline to the same numerals as some brain based analogue of the bodily manipulations of numerals. Indeed, even Fodor thinks something like this is true: "I wouldn't be in the least surprised, for example, if it turned out that some arithmetic thinking is carried out by executing previously memorised algorithms that are defined *over public language symbols for numbers* ("now carry the '2' and so forth").[19]" (Fodor 1998, p. 72) It turns out that there is even substantial research which shows that this kind of thing is going on whenever we do offline arithmetic. Stanislas Dehaene (a cognitive neuroscientist) and colleagues have produced stunning results and a compelling model that shows that we think in an internalised conventional representational system.[20] "The model that emerges suggests that we all possess an intuition about numbers and a sense of quantities and of their additive nature. Upon this central kernel of understanding are grafted the arbitrary cultural symbols of words and numbers…. The arithmetic intuition that we inherit through evolution is continuous and approximate. The learning of words and numbers makes it digital and precise. Symbols give us access to sequential algorithms for exact calculations." (Dehaene 2007 p. 41).

I take it that Dehaene and colleagues are doing empirical research on cognition and yet if we are to take Adams and Aizawa's underived content condition seriously they cannot be. This seems to me to be decisive. Scientists often ignore the strictures and demarcations laid down a priori by philosophers. Philosophers, such as Adams and Aizawa, may think that they are providing a theory of underived content for cognitive scientists (2001, p. 50), yet cognitive scientists may not be at all interested in these theories because their work does not need them, or they may not be working only with representations with underived content.

[19] This is quoted by Clark in his 2008 paper p.54–5 but he doesn't martial it against Adams and Aizawa.
[20] See *Dimensions of the Mind* (this issue) for more details.

It is interesting at this point to note what Adams and Aizawa and I agree upon:

1. We agree that there is no difference in the content of sentences I say to myself privately and those that I say publically (out loud)
2. That human cognitive capacities are limited
3. That we have developed public systems of representation and methods for manipulating those representations which help us to overcome our limitations.

But then there is disagreement about what follows from this. I think that it shows that we sometimes think publically by manipulating external representations, but that we also internalise them and that this transforms our offline cognitive capacities. This is what I take the importance of research like Dehaene's to be. Adams and Aizawa conclude that public systems of representation and their manipulation cannot be cognitive; only the manipulation of representations in the brain with naturalistically determined underived content can be cognitive. I think that is a bad move, it is a zero sum game: it forces us to deny Dehaene's claim that the "arbitrary cultural symbols of words and numbers" can count as being cognitive, it forces us to say that the analogical representations of quantity in the parietal lobe are the only cognitive representations of interest.

Adams and Aizawa once wrote that "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."[21] I find myself in complete agreement with them and it is precisely what I claim to be committed to. Tools aren't cognitive, our manipulations of public representations[22] can be.[23]

References

Adams, A., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology, 14*, 43–64.

Adams, A., & Aizawa, K. (2008). *Defending the bounds of cognition*. Oxford: Basil Blackwell.

Adams A. and Aizawa K. (2010a) Defending the bounds of cognition, In: Menary, R. (ed.) *The Extended Mind* Bradford Books/MIT Press.

Adams, F., & Aizawa, K. (2010b). The value of cognitivism in thinking about extended cognition. *Phenomenology and the Cognitive Sciences*. doi:10.1007/s11097-010-9184-9

Beer, R. (1995). Computational and dynamical languages for autonomous agents. In R. Port & T. Van Gelder (Eds.), *Mind as motion: Dynamics, behavior, and cognition*. Cambridge Mass: MIT Press.

Clark, A. (2008). *Supersizing the mind: Embodiment, action and cognitive extension*. Oxford: Oxford University Press.

Clark A. (2010) Memento's revenge, In: Menary, R. (ed.) T*he Extended Mind* Bradford Books/MIT Press.

Dehaene, S. (2007). A few steps towards a science of mental life. *Mind, Brain and Education, 1*(1), 28–47.

Fodor, J. (1998). Do we think in mentalese: Remarks on some arguments of peter carruthers. In J. Fodor (Ed.), *Critical condition: Polemical essays on cognitive science and the philosophy of mind*. Cambridge, MA: MIT Press.

Hurley S. (2010) Varieties of externalism, In: Menary, R. (ed.) T*he Extended Mind* Bradford Books/MIT Press.

Hutto, D. (2008). *Folk psychological narratives: The sociocultural basis of understanding reasons*. Cambridge Mass: MIT Press.

[21] Adams and Aizawa 2001 p.43.

[22] Sometimes by means of tools.

[23] Whether or not they are manipulated in the head or in a shared public environment.

Menary, R. (2006). Attacking the bounds of cognition. *Philosophical Psychology, 19*(3), 329–344.
Menary, R. (2007). *Cognitive integration: Mind and cognition unbounded*. Basingstoke: Palgrave Macmillan.
Menary, R. (2009). Intentionality cognitive integration and the continuity thesis. *Topoi, 28*, 31–43.
Menary, R. (2010). Cognitive integration and the extended mind. In R. Menary (Ed.), *The extended mind*. Camb Mass: Bradford Books/MIT Press.
Millikan, R. (1984). *Language, thought, and other biological categories*. Camb Mass: Bradford Books/MIT Press.
Port, R., & Van Gelder, T. (Eds.). (1995). *Mind as motion: Dynamics, behavior, and cognition*. Cambridge Mass: MIT Press.
Ramsey, W. (2008). *Representation reconsidered*. Cambridge: Cambridge University Press.
Ross D. and Ladyman J. (2010) The alleged coupling-constitution fallacy and the mature sciences in Menary, R. (ed.) *The Extended Mind,* Camb Mass.: Bradford Books/MIT Press.
Rowlands, M. (1999). *The body in mind: Understanding cognitive processes*. Cambridge: Cambridge University Press.
Searle, J. (1980). Minds, brains and programs. *The Behavioral and Brain Sciences, 3*, 417–457.
Spivey, M. (2007). *Continuity in mind*. Oxford: Oxford University Press.
Sterelny, K. (2010). Minds: Extended or scaffolded? *Phenomenology and the Cognitive Sciences*. doi:10.1007/s11097-010-9174-y
Stotz, K. (2010). Human nature and cognitive–developmental niche construction. *Phenomenology and the Cognitive Sciences*. doi:10.1007/s11097-010-9178-7
Sutton, J., Harris, C. B., Keil, P. G., Barnier, A. J. (2010). The psychology of memory, extended cognition, and socially distributed remembering. *Phenomenology and the Cognitive Sciences*. doi:10.1007/s11097-010-9182-y
Thelen, E., & Smith, L. (1996). *A dynamic systems approach to the development of cognition and action*. Camb Mass: Bradford Books/MIT Press.