RoPP 3.1., Special Issue:

The Body Represented/Embodied Representation

Contents

Environmental Representation of the Body

Adrian Cussins

Review of Philosophy and Psychology 3 (1):15-32 (2012)

Adrian Cussins

adriancussins@mac.com

Universidad Nacional de Colombia

Much recent cognitive neuroscientific work on body knowledge is representationalist: "body schema" and "body images", for example, are cerebral representations of the body (de Vignemont 2009). A framework assumption is that representation of the body plays an important role in cognition. The question is whether this representationalist assumption is compatible with the variety of broadly situated or embodied approaches recently popular in the cognitive neurosciences: approaches in which cognition is taken to have a 'direct' relation to the body and to the environment. A "direct" relation is one where the

boundaries between the body and the head, or between the environment and the animal are not theoretically important in the understanding of cognition. These boundaries do not play a theory privileged role in cognitive explanations of behavior. But representationalism appears to put a representational veil between the locus of cognition and that which is represented, making cognitive relations to the body and to the environment be indirect, with a high associated computational load. For this reason, direct approaches have tried to minimize the use of internal representations (Suchman, 1987, Barwise, 1987; Agre and Chapman, 1987; Brooks, 1992; Thelen and Smith, 1994; van Gelder, 1995; Port and van Gelder, 1995; Clark 1997, 1999; Rupert, 2009, p. 180). Does a cognitive neuroscience committed to direct relations rule out a representationalist approach to body knowledge? Or is direct representationalism possible?[1]

---

1. The Formality Constraint and Extended Mind

Traditional Representationalism generates a theoretical division of labor, as Fodor (1980) suggested long ago: between an internalist computational and syntactic psychology, and an externalist semantic and naturalistic psychology.[2] This is a divide between a theory of the

---

[1] Many thanks to Frederique de Vignemont and Carlos Márquez for comments; to the members of the CCP research group, especially Guillo Gomez, Carlos Márquez, Camilo Ordoñez, Diana Acosta, Anderson Pinzon, Juan Camilo Espejo, Manuel Amado, Alfonso Conde and Juan Diego Morales Audióptica for joint investigation of these themes; to the Facultad de las Ciencias Humanas, el DIB and the Departamento de Filosofía of the Universidad Nacional de Colombia, for financial and other support of CCP.

[2] It is salutary, given the role of fashion, to remember that Fodor saw himself in 1980 as arguing against the tradition of 'embedded psychology': "I have in mind the tradition ... that psychology is a branch of biology, hence that one must view the organism as embedded in a physical environment" (p. 229 in the RePresentations (1981) reprint)]

processes, e.g. computational processes, that act on representations which are instantiated, or have their 'vehicles', in the nervous system of the animal; and a theory of the 'wide' or 'broad' processes that have to do with relations to the environment. A syntactic psychology is a theory of the representers, and of the computational processes defined over these representing vehicles and their syntax, which ignores, or has only an indirect relation to, the structures and processes of the body and the environment that is represented. The division between the cognizing mind and its world would be drawn theoretically as the division between the representer and the represented. The formality constraint (Fodor 1980) entails that computational psychology - as opposed to semantics - has direct access only to the representing vehicles. Psychology has access to the environment only in the sense that it has access to representational vehicles whose contents may present the environment. This is indirect access in that psychological processes are defined over the local, causally potent, vehicular, syntactic properties, and not over the causally distal content properties; and the former are not sufficient to determine the latter. Hence cognitive theory, subject to Fodor's formality constraint, implies a theory privileged divide between cognition on the one hand, and the body and the environment on the other. That's methodological solipsism, or the contrary of the thesis that cognition has a direct relation to body and to the environment.[3]

---

[3] In philosophy most work on 'broad' cognition has been ontological in character: do cognitive states supervene on the environment, or only on the brain? But the more interesting claims have to do with explanation; with which boundaries are theoretically important, and what kinds of primitives are required by a good theory of cognition. It is this second, explanatory, form of the thesis that has far more relevance for empirical work in the cognitive neurosciences.

There is a modest response to this worry about the compatibility of direct cognition with representationalism. The modest response is that the use of representations does not entail the rejection of environmental cognition because representational vehicles of cognition exist not only in the head but also in the environment, and that cognitive processes may be defined, also, over environmental representations. If representational vehicles important for cognition are also instantiated in the environment, then the Fodorian representationalist thesis that the theoretical divide between mind and world is the divide between representer and represented would not entail the thesis that a theory of the environment is not a part of the theory of mind. Insofar as a theory of the environment is, or includes, a theory of its representations, it would belong to a theory of cognition, and there would be no privileged cognitive - theoretical boundary between head and environment. (Or, at least, there would be no privileged boundary in virtue of the theory's commitment to vehicles of representation).

We may call this (relatively) modest response to the threat of incompatibility, "Extended Mind". There have been many versions of representationally Extended Mind. One of the earliest, and still one of the best, was Hutchins (1995) in his book "Cognition in the Wild". Hutchins gave us a beautiful, cognitive anthropological account of ship navigation in which effective and intelligent management of the task of navigation involved computational processes distributed throughout the physical and social and technological spaces of the ship, and defined over representations in those spaces. As we will see, Hutchins' work is of particular importance here, because it is a fine example of how to combine representationalism with the importance for cognition of the body and of the environment. More recently, Clark and Chalmers (1998) have discussed "Extended Mind" using David

Kirsh's (1995) work on the cognitive use of space, and the example of the use of a notebook, that is always carried around by an Alzheimer's patient (See also more recent discussions in Menary, 2010).

Part of what interests me here is a contrast between Hutchins's discussion and the discussion in Clark & Chalmers. The latter allows that a cognitive theory of processes defined over representational vehicles is environmental, and to that extent involves direct relations between cognition and the environment. The former, however, goes beyond relations involving the vehicles, in that it points us towards an understanding of how representational content relations can also be direct. The discussion in Clark and Chalmers is about *where* we find cognitively useful representations, whereas the ambition of Hutchins (1995) is to change how we understand computation and representation.[4]

Hutchins, and others that share his anthropological orientation, or an orientation that has been influenced by the tradition of Soviet Psychology (Leontiev, Bakhtin, Vygotsky, etc.), are searching for a deeper sense in which cognition is environmentally distributed; a sense that is more explanatorily basic than that which comes across from the idea of extensions of individual, head-based cognition through the use of external representations, that might be found, for example, in a notebook, or in a video game. Hutchins, too, makes much play of external representations such as charts, compasses, logs and alidades. But there is a difference. In Clark & Chalmers, the extension of representation into the environment is shown as cognitively useful, but not explanatorily essential. The Kirsh and Maglio (1994) Tetris rotations can be done in the head, it just takes a few hundred milliseconds longer.

---

[4] Hutchins (1995) has been an important influence on Extended Mind, although I am suggesting that it should be read as not being theoretically modest.

Writing things down in a notebook is one strategy; another is to learn better memory skills and discipline: in terms of cognitive performance, they might have the same result (indeed, the latter has numerous advantages). The fact that one strategy involves environmental tools, and the other is employed 'in the head', is a difference that is not theoretically essential: what is theoretically important is the nature of memory, and its use of representations. It's not theoretically important where those representations are. The examples that Clark & Chalmers point to are examples of environmental methods of improving cognition; they are not examples that show how the environment and its structures are an essential part of cognition. Their language is of cognitive processes "being extended into the world" or "leaking out into the world". Their picture is of a non-environmental cognitive home, from which cognition may make, undoubtedly useful, forays. We should, perhaps, call this picture "Contingently Extended Mind".

At one point Clark and Chalmers say,

> in seeing cognition as extended one is not merely making a terminological decision; it makes a significant difference to the methodology of scientific investigation. In effect, explanatory methods that might once have been thought appropriate only for the analysis of "inner" processes are now being adapted for the study of the outer, and there is promise that our understanding of cognition will become richer for it (Clark and Chalmers, 1998, p. 10).

But what the more broadly anthropological cognitive scientists (such as Hutchins, Suchman, Lave, or Agre) have been doing is the reverse of Clark & Chalmers' suggestion. Not to extend inner explanation outwards, but, rather, that we can learn something about how to explain brain processes through investigations of, and the use of explanations

appropriate to, the environment: to understand the inner at least partly on the basis of the outer. Of course there are purely anatomical-physiological specifications of the brain that are entirely 'inner'; the 'anthropological' point here, however, is that specifications of the brain *that show its role in cognition* are not entirely inner, and should be modelled, in part, on specifications of the environment. Extended Mind makes cognitive use of representational vehicles that are environmental, but it still accepts the formality constraint.

The environment contains spatial resources that are cognitively useful, and whose use is cognitive, but is the environment also a model for what we take to be our explanatory primitives? Going forward, should research in the cognitive neurosciences use notions that fit happily within the individual head paradigm - although they can be extended beyond it - such as consciousness, syntax, or sensory-motor contingencies, or should research use a different kind of intrinsically environmental notion? Would such an environmental notion be representationalist, or not? And what would an environmentalist conception of body knowledge look like?

Hutchins is pointing us towards senses in which cognition is necessarily environmental, because - unlike psychology governed by a formality constraint - an anthropological psychology employs explanatory primitives that are environmental. To better grasp these senses, we need to understand an aspect of the logic of representational content, an aspect that is, in general, little understood, yet which is of importance both for questions about the possibility of cognition in animals, about how to go beyond the formality constraint, and for more specific questions about the cognitive neuroscience of body knowledge. I will explain how it is that this fundamental logical feature, missing in Clark and Chalmers, and indicated

in Hutchins, is the key to understanding how to develop direct representationalism of the body and of the environment.

---

2. Against referentialism in the theory of representational content: the argument from the logic of conditions.

Most work in the cognitive neurosciences on representation tends to be on the nature and form of the representational vehicles involved in some cognitive process: imagistic or symbolic, procedural or declarative, local or distributed, implicit or explicit, analog or digital, narrow or broad. There has been much less attention paid to the nature of the content carried by the representational vehicle. No doubt this is because there is a general, implicit, assumption in the community that there is only one kind of content: what we may call "truth-conditional" or "referential" content. But the assumption is incorrect, and, as we will see, it is what generates the deeper incompatibility between representationalism and direct approaches to cognition. The cognitive neurosciences should pay far more attention to representational content, both because it would allow us to consider a variety of distinct kinds of representational relation between brain and body, and between brain and environment, and also because it would give us a way to think about a variety of kinds of computation. Computational explanations depend not only on the nature of representational vehicles, but also on their contents. So if it turns out that there is more than one kind of content, then that result affects our understanding of the space of possible computation.

"Truth-conditional" or "referential" content is the kind of content involved in Clark & Chalmers' (1998) story about how the Alzheimer's patient gets to the exhibition:

> Otto carries a notebook around with him everywhere he goes. When he learns new information, he writes it down. When he needs some old information, he looks it up. For Otto, his notebook plays the role usually played by a biological memory. Today, Otto hears about the exhibition at the Museum of Modern Art, and decides to go see it. He consults the notebook, which says that the museum is on 53rd Street, so he walks to 53rd Street and goes into the museum.

<The museum is on 53rd Street> is the content carried by the written words "The museum is on 53rd Street". <The museum is on 53rd Street> is truth-evaluable, it is a condition on the world, it is true if and only if the condition is satisfied (i.e. that the museum *is* on 53rd Street), and functions effectively (in helping Otto get to the museum) only if it is true. The parts of the content <The museum is on 53rd Street> have a referential function to compositionally determine the truth condition of the whole content. The general form of the example content <The museum is on 53rd Street> is that an object is at a location, and the example content's constituents fix the object, the spatial relation between object and location, and the location. These constituents are truth condition determiners. I will call them "referential contents". Included in the extension of 'referential content' are propositions, or truth-evaluable contents with referential constituents, because these are the determiners of the truth conditions of more complex propositions that contain them as constituents. Referential contents are the contents of judgements, and their characteristic normativity is the normativity of guidance in judgement. This is generally the normativity of truth guidance: the point of the activity of judgement is to judge truly. But related

normativities of judgement also fix referential content: correctness conditions, conditions of accuracy or verification conditions. A referential content is a content that presents a referent, which is its contribution to the determination of the truth condition of the content of which it is a part. <The museum> presents as its referent the object which is the museum, because the truth condition which is <The museum is on 53rd Street> is that that object - the museum - is on 53rd Street. That is, the building is the semantically presented referent of <the museum> because the building is the content's contribution to the truth condition to which <the museum> belongs.

In summary, a representational content has a presentational property (its mode of presentation of something in the world), and a normative property, which - in the case of referential content - is either a truth value or a truth condition or a correctness condition or a satisfaction condition. The compositional role of a constituent in determining the governing normativity of the whole propositional content to which it belongs is that which fixes what is presented by the content; its ontology.

It is possible to tell this story a little differently, in terms of the determination of truth values, rather than truth conditions, but that difference won't make a difference in the present context. What is important for the notion of referential content is the constitutive role of truth. Apart from details of this sort, all this sounds truistic, because it forms the heart of the logical tradition, a tradition that is so dominant that we often lose any sense that there are alternative ways in which content can function. If our interest is in giving an account of perception and action or activity (rather than the foundations of mathematics, or the nature of some entirely formal domain), then either such an account should not be

representational, or we need another notion of representational content, distinct from referential content.

Why? Recall the anxiety about the compatibility of representationalism with direct approaches to cognition, with which we started. There were two parts to that anxiety, only one of which depended on the assumption that representational vehicles are exclusively located in the heads of cognizers; the assumption that was rejected by, for example, proponents of Contingently Extended Mind. The second worry had to do with the idea that where there is representation, there is a veil between representer and represented, that renders indirect the presence of that which is represented. Representationalism's veil between cognitive process and environment contrasts with the direct role of the environment in cognition: as Rodney Brooks likes to say, the world is its own model (Brooks, 1992). This is an argument from the nature of representational content: it makes no difference whether the representational vehicles, like Clark & Chalmers' notebook, are in the environment rather than in the head.

Now we are in a position to make this second part of the worry less metaphoric. A specification of referential content is the specification of a condition on the world. If the content of an intention to act is a referential content, then an intention fixes satisfaction conditions for the action. But if the content of the intention doesn't do more than that (if its content is exclusively referential), it will always be a further question, of some possible way of acting, whether that way satisfies, or does not satisfy, the condition. One can grasp the intention (and so grasp its conditions of satisfaction) without knowing of any particular action whether or not it satisfies the intention. So having an intention - under the assumption that there is only referential content ("referentialism") - is not cognitively

sufficient for action. But having an intention is cognitively sufficient for action, so we had better reject the referentialist conception of content: there is more to content than referential content.

This argument makes use of a general logical point about conditions, and their relation to the world of things that may or may not satisfy the condition. At the beginning of a criminal investigation, the detective, who is acquainted with John, knows very well that the person she is seeking is the unique satisfier of the condition <the assassin of John>, while having no idea of any particular person she meets, whether that person satisfies the condition. If this weren't so, then a detective's investigation would be completed as soon as it began.

<Do PHI>, where <PHI> is a referential content, leaves open, of any particular intervention in environmental activity, whether or not that intervention is a way of PHI-ing, just as the condition <the unique assassin of John> leaves open of any particular person, whether or not that person satisfies the condition. If someone acts rationally, they act based on their knowledge of what to do: given my knowledge of what to do, <this> was the thing to do, and not <that>. But if <this> and <that> are interventions in environmental activity, and my knowledge of what to do consists exclusively of referential content, then my knowledge would not be sufficient to rule in <this> and rule out <that>. My knowledge would establish a condition on the action, but what I need in addition is to know of some particular intervention in activity, that that intervention satisfies the condition. I would need a further inferential step whose conclusion is that <this>-ing is a way to PHI. Yet I can't get exactly that. What I can achieve through standard inference (i.e. inference on referential contents) is that whatever satisfies some condition, PSI, also satisfies the condition PHI.

<this>-ing is not a condition, but an intervention in activity, and so it can only figure in inference through some conceptualisation of it, that is through its 'condition-al' presentation in referential content. But then I am no further forward because, once again, it is an open question whether or not the intervention satisfies the condition-al presentation. That's a further thing that I need to know in order to act.

Someone might quite reasonably respond to the argument of the previous three paragraphs ("the argument from the logic of conditions") by saying that what links intentions to specific, concrete actions is -- not inference! -- but experience. I just *see* that this is the way to PHI. Or, I have *body knowledge* of how to PHI. True enough. But now what is the content of the experience, or of the perception, or of the body knowledge? If the only kind of content is referential content, then we are back with the same problem. Worries of this kind have motivated abandoning representation in a theory of perception in favor of either non-cognitive causal relations to the environment, or the use of notions like consciousness or direct acquaintance (Campbell, 2002). This is not the place to explain the problems with those theoretical avenues, but fortunately we don't need to: our goal was to account for the compatibility of representationalism with direct cognitive relations to the body and to the environment. Abandoning representationalism won't further us in that goal.

Note that I am not denying that there is particular or singular referential content. On the contrary, referential content that picks out particulars as particulars is of fundamental cognitive importance. The question is whether an account that countenances *only* the theoretical resources used in the semantics of referential content can make room for singularity. If we have assumed that perceptual experience makes available singular

contents, then those singular contents can enter into inferences that have conclusions about a singular object or action. The difficulty is whether a referentialist is entitled to such an assumption. <Do <this>> is not a command to do whatever satisfies a satisfaction condition. Rather, it is a command to do what is demonstrated. Satisfaction conditions are not demonstrations. You can't get demonstrative contents out of conditional contents, unless the demonstrative contents were already in there. So where do the demonstrative contents come from? Experience! We're back in the same argumentative circle. There are two ways out of the circle. Either you deny that characteristic experience is content-bearing, claiming instead that it is a matter of direct, non-representational, e.g. conscious, relations to the world, and then go on to show how content bearing states can be singular in virtue of their relations to non-representational, conscious states (Campbell, 2002; Brewer, 2011). Or you recognize that the theoretical nature of experiential content – content which is *characteristic* of experience – is not referential content, but some other kind of content. The first option is a rejection of *representationalism*. A compatibilist must adopt the second option. The second option is a rejection of *referentialism*.

Another kind of response to the argument from the logic of conditions is to say that the connection between cognitive content, e.g. an intention, and action is neither representational nor experiential nor inferential, but sub-personal and neuro-physiological. It is a matter of the machinery in the brain, and there is no cognitive correlate for the operation of that machinery; that is, there is no correlate that can play an explanatory role in a cognitive theory. Cognitive theory just breaks down at this point, and we have to shift explanatory levels, to, as it were, pure neurophysiology, uncontaminated by cognitive notions such as representation or inference. The sub-personal response to the argument

from the logic of conditions holds that intentions cause actions not because of a causally explanatory law at the level of cognitive description, but because of a causally explanatory law at the level of neurophysiology. Cognition can't make direct the relation between intention and action, but neurophysiology can.

Maybe, but this should be an option of last resort, once all alternative cognitive options have been shown not to function. Especially in the present context, where we are concerned with the coherence of the neuroscience of body knowledge, which as a form of knowledge, belongs to cognitive theory. Cognitive theory had better not break down here, because the relation between an intention and an action has cognitive consequences. Perhaps it matters for the epistemic responsibility of the actor, which of the ways of doing PHI is selected. Or perhaps doing <this> is a way of doing PHI only if <this> is performed with great care or precision or skill. The epistemic virtues of care and precision and skill are properties of the relation between the intention (or other cognitive progenitor of action) and action. So this relation had better belong to the explanatory space of the person. If the relation were nomologically governed not at the level of cognitive theory, but only at the level of neurophysiology, then the relation would not belong to the space of the person. So we should reject the sub-personal response.

Perhaps if sufficient care or precision or skill is not forthcoming in the person's trying to act <this>, then the person should shift to doing <that>, to some other way of trying to PHI. But readiness to make the shift requires that the relation between intention and action is cognitively monitored. If the theory of the relation has been handed off to non-cognitive, 'pure' physiological theory, then the relation is not cognitively accessible: the relation could not be cognitively monitored. So the possibility of acting with epistemic virtue requires a

cognitive explanation of the relation between cognition and action. *How* we carry out our intentions is a manifestation of our skills and virtues, and so it is something for which we should take cognitive responsibility. Going sub-personal – an option perhaps when robotic control is appropriate – is a matter of abandoning cognitive responsibility (Acosta, 2012). It is one thing not to take cognitive responsibility for everything for which we could in principle take cognitive responsibility (that's not a practical possibility), but it is quite another that the theory of representational content should make it impossible to take cognitive responsibility for any relation between intention and action.

So, for person-level reasons to do with the epistemic virtues, the relation between action and the cognitive progenitors of action must be subsumable under cognitive level explanation. (We require a cognitive level account of the relation between, for example, intention and action). And there are also internal difficulties in giving a purely neurophysiological account of this relation. Perhaps most importantly, there is no adequate pure neurophysiological explanation waiting in the wings. Shorn of all cognitive notions, there is very little that we know how to say about how the brain controls intelligent behavior. We do know something about cognition-free explanation of adaptive behavior, if we presuppose that behavior occurs within an animal niche. If we presuppose a fairly fixed animal niche, then it becomes possible to treat the relation between a progenitor of activity and activity in a non-cognitive way: certain stimulus conditions, or other progenitors, can have fixed relations to effector responses. This is common in the animal kingdom. But personal epistemic virtues, such as intelligence, depend on not treating the relation to the environment as a fixed niche relation: intelligence depends on managing change in the niche. Niche-dependence is a sign of lack of intelligence. So, as one might expect, the

argument for not going non-personal in behavioral explanations depends on our need to explain person-level characteristics, such as the epistemic virtues.

These more philosophical points about the role of epistemic virtue in the possibility of successful action are also connected with a motivation for a science of direct cognition: a way of doing the science of behavior in which cognition consists in the constant monitoring of, and intervening in, environmental relations and feedbacks between perception and activity. Direct cognition is dissatisfied with traditional accounts, in which perception is an input module, and action is an output module each with distinct theories; and each divided from the other by ... the world. It is no better if perception and action are divided by the brain. Going non-cognitive, and sub-personal, as a response to the argument from the logic of conditions, entails a causal explanatory discontinuity in cognition and in cognitive theory between action and the cognitive progenitors of action, such as perception. So direct cognition rules out the going sub-personal response to the argument that I gave against referentialism. Whether perception and action are divided by the world, or by the brain, that division would entail the necessity of separate theories for perception and action. Whereas a direct account entails active perception (Ballard et al., 1997) in which a theory of perception is also a theory of activity. According to such an account, cognition is the management of a cycle, which is, at each point, both experiential and active. Cognition is active guidance, and the very notion of guidance involves both perception and action. In section §4, I show, briefly, how we might develop a notion of representational but non-referential content given that cognition is active guidance.

---

3. Against referentialism: the cognitive content cold argument

But before moving on to a positive account of non-referential content, I want to mention - but not develop - another argument against referentialism: the referentialism makes cognitive content cold argument (RCCC). The importance of thinking through these negative arguments is that they demonstrate some of the characteristics that non-referential content must have. Non-referential content must not be a condition, but rather something that has the kind of relation to an object or to an action that a demonstration has; and so can directly yield action without requiring further knowledge or inference. And non-referential content must be hot.

RCCC notes that the same referential content can be the content of very different attitudes: I can hate, or be indifferent to, what is given by the very same referential content as what I love, or adore or desire. When I desire something, I am drawn towards it. When I hate something I am drawn away from it. Given that 'it' is given by the same referential content in these two cases, the referential content must be neutral with respect to whether I am drawn towards or away from the referent of the content. A referential content has no inherent activity-valence (it does not dispose us - has no intrinsic connection to a disposition - to being drawn towards or away), and no inherent motivational force. So, within referentialism, motivational force can come only from the attitude to the referential content; we need something like <desire that p>. But if one of the elements in a cognitive generation of action is the desire that p, then that process will also require a belief about the best or the right way or a good enough way to bring about that p. It is the fact that referential content is free of activity-valence (that it is cold) that generates the belief -

desire - inference - intention structure of practical reasoning. We get something like: S desires that the window be open, has no competing desires, believes that pressing the <open> button is the way to open the window, and so forms the conclusion that pressing the <open> button should be done. But that conclusion is just another belief, albeit a normative belief. And like any other referentialist belief content, it is motivationally cold. A subject could use the belief in further reasoning, or could simply 'appreciate' it, valorizing its truth. No action will result.

By contrast, if a kind of content is motivationally hot, it needs no inferential structure for it to have a relation to action. Hot content is intrinsically active; the content itself disposes the system to intervene actively in the environment. So one kind of non-referential experiential content is the content <*yummy*>, or the content <*yucky*>. I will return to these delicious / revolting hotties, but they help to make the contrast with the coldness of purely referential content. Experiencing <yummy>, the child is drawn to put more chocolate in its mouth. It is so drawn not as the result of some further cognitive or inferential step. None is needed. Of course, a cognitively complex system does not always, or, in general, do that which it is drawn to do; the felt disposition in experiencing <yummy> may be overcome by other cognitive factors. The present point is that no further cognitive step is necessary to mediate the connection between the content and the action in the case of <yummy / yucky> contents. Such further steps are always necessary for purely referential contents. Referentialism has no account of action.

An aspect, perhaps, of the coldness of pure referential content, is that it is subjectivity-free. This is a large and difficult topic that I cannot go into here, beyond noting its relevance. Strictly speaking, a full-blooded discussion of representational content should not

presuppose that the bearer of content is a subject, since the conditions on being a subject depend on the nature of the contents, and the content-involving activities, available to the bearer. So I will talk of a 'content using system' or System, for short.

The result of referentialist practical reasoning might be the formation of an intention: <System intends that pressing the <open> button be done>. Aware of the intention, System stares out at the world, secure in the appropriateness of its intention.

Nothing happens in the environment.

System could be causally rigged up so that any physical state that realizes an intention of System has causal, physical consequences. But that way of generating activity from cognition is the sub-personal route that we have already had reason to reject, in the case of systems that are persons.

Note also that System intends that pressing the <open> button be done does not fix which source-of-activity is to effect the pressing. So the practical reasoning could be jigged so that its conclusion is something like: System intends that pressing the <open> button be done by System. But, within a referentialist framework in which no content contains intrinsic me-ness — what we might call a "subjectivity-valence" — System has no special cognitive relation to System. System, for System, is just one object amongst all the objects in the world. Being me-free — being without subjectivity-valence — is another kind of coldness, another kind of disconnection from action. From its lofty referentialist position, System refers to one object in the world, and assigns an intention to be carried out by that object. The database is updated. System maintains its immobile state.

There is, of course, a special relation — identity — between System, which is the referent of the grammatical-subject of: System intends that pressing the <open> button be done by System and System which is the referent of the grammatical-instrument. This allows a reformulation of the representational process within System: <I intend that pressing the <open> button be done by me>. From the premise of identity, and the previous database entry, there is a valid inference to this new practical conclusion. And, in a real Subject, having an intention with this form results, normally, in action. That's what we want to explain. We are not entitled to presuppose that. The problem is that in referentialism, the content <I intend that pressing the <open> button be done by me> is the very same content as <System intends that pressing the <open> button be done by System>. Purely referential content is exhausted by its reference, and so two co-extensive contents have the same referential content. But <System intends that pressing the <open> button be done by System> has no intrinsic cognitive connection to action; hence <I intend that pressing the <open> button be done by me> has no intrinsic connection to action within a referentialist System.

Nothing happens in the environment.

Something happens in the theorist, who accepts the need for an account of non-referential content, which is not a condition on the world, but more like a demonstration of something in the environment; which has intrinsic activity-valence; and which has intrinsic subjectivity-valence.

RCCC and the argument from the logic of conditions show that when representational content is restricted to referential content there is a gap – unbridgeable by cognition – between intention or perception on the one side and action or the environment on the

other. This gap is the veil of representation; an indirectness with respect both to action and to the singularity of what is presented in perception.[5] The possibility of direct representationalism depends on the possibility of a kind of content which is not referential, and which intrinsically disposes the bearer to intervene in the environmental flow of activity.

---

4. Cognition as active guidance

The normativity of guidance in truth-judgement is the constitutive normativity for referential content. Understanding the possibility of a distinct kind of content requires substituting a distinct kind of normativity for truth normativity in a theory of content. Suppose we develop a theory of content that takes seriously the motto of direct cognition that we met with above: cognition is active guidance. Then the kind of normativity that we need is a normativity of getting about well in an environment; that which permits being well-guided, or having efficacy, in the environment; *correctings* in activity rather than conditions of correctness, or that which permits affective sensitivity to 'how things normally go on around here'. I call this "mundane normativity" because its management is a basic capacity of all normal, mature animals: animals are good at finding their way around in their particular environments, or niche. Direct representationalism requires a kind of

[5] I have focused here on the relation between intention and action because the difficulty faced by a referentialist theory shows the importance of the role of body knowledge in mediating this relation. The argument can be adapted for the parallel relation of perception and environment.

representational content grounded in our animal nature, whose structures are determined by the mundane normativities of active environmental guidance.[6]

Good guidance depends on the active monitoring of the relation between perception and action, and in general between cognition and action. Its locus is precisely in those relations over which referentialism has difficulty. And actively monitored guidance is cognitively basic in human cognition, both with respect to capturing what is characteristic of animal cognition, and with respect to ordinary, mundane getting about.

Knowing how to get around well — being environmentally well guided, or having environmental efficacy — is a kind of cognition that we exercise irrespective of what intentions to act we may have, and even if we have no intentions with referential content. We are always in an environment; we are not always acting with intentions or plans. During any ordinary active hour of our lives, we are doing uncountably (but accountably) many things (moving around without bumping into things; avoiding what is less pleasant; getting right the appropriate social distance for conversation; comporting one's body parts in socially appropriate, or identity expressing, ways; maintaining a string of words; entering a conversation at a pause point; stopping talking when it is appropriate; being funny or not being; looking warm or stern; moving around in harmony with the rhythm of the movements of those around one), things of which we have not thought, and may not even be consciously aware, and for which we formed no intention, or goal or plan. Yet they form part of our experience. All of these bodily manifestations *can* become objects of our consciousness or thought, but they do not need to do so in order to accountably and

---

[6] Earlier but more complete expositions, with more examples, of mediational content are in Cussins (2003) and Cussins (1992).

experientially manifest one's skills, identity, or membership in a society: our body is cognitively mediating these activities. This is the experiential substratum of our activity; it is normative; it manifests skill; it is embodied; and it does not depend on practical reasoning. Intervening in substratum activity is a normative cognitive capacity — knowing how to get around well — which subserves any of our environmental intentions with referential content, if we have them. It is the cognitive means by which an intention yields action, and it guides and motivates activity even when there are no intentions.

Representational content has a normative property, but also a presentational property: content presents the environment or the world. Referential content presents the world in terms of truth structure: the determiners of truth conditions and truth values. An ontology of objects, properties and states of affairs is an ontology of truth structures: for Otto, the museum is an object in this ontology, because it is a determiner of the truth conditions of <the museum is at 53rd street>. What a representational content presents is a part of the structure of the content's governing normativity; in the case of referential content, the governing normativity is truth, and its structures are truth condition determiners such as objects and properties. Replace truth normativity with the mundane normativity of active guidance, and thereby replace referential content with *mediational content*. Then mediational representational content presents the environmental structures that guide the flow of activity through the environment. We have already encountered a very simple example of these structures: <yummies> and <yuckies>. At this simple level we experience the environment as a spatial and temporal arrangement of <yummies> and <yuckies>; guiding us by drawing us towards them or away from them. From these simple elements

more complex ontological structures are formed: trails and rhythms and territory[7]. The mundane normativity involved in getting around well involves an active sensitivity to the mediational ontology of trails and territory that are presented in mediational contents.

Trails are in the environment, but their characterization depends on an animal's sensitive – active relation to the world. An ant trail is an environmental material structure, but that material structure is an ant trail only in relation to the capacities of ants: a trail for an ant is not a trail for a monkey. It is the same with human cognitive trails, that structure our capacities for active guidance. So a theoretical specification of a trail depends on the theory of those capacities. An animal's capacities, to follow and to make its trails, is in part experiential: they depend on the contents of perception. There is a subjective element to these contents, and so there is a subjective element in the trails themselves.

Trails have subjectivity-valence in that if an environmental structure is a trail, it is so *for-the-animal*, and this aspect of being for-the-animal shows up in the experience of more advanced animals as a *for-me* element in the experience, or — less conceptually — the experiential form of the subjectivity-valence is that the experience is intrinsically motivational / dispositional. An environmental <yucky> repels *me* or *us*, and the trail solicits and affords *my* or *our* activity, where the use of these first-person indexicals in an account of experience should not be replaced by co-extensive non-indexical terms: yuckiness and trailishness are environmental but also subjective. A yucky is experienced as

---

[7] Each of these notions — yummy / yuckies, trails, rhythms and territory — , and the developmental relations between them, are quite intricate. In the CCP research group ('Communication, Coordination and Perception') in the Universidad Nacional de Colombia we have been investigating them, both philosophically and through modelling. In the present context, an intuitive understanding of trails and territory, etc., should be sufficient.

a *subjectively felt, environmentally warranted, pushing*. A trail is experienced as a *subjectively felt, environmentally warranted, drawing-onwards*.

Mediational content has intrinsic activity-valence: it solicits activity and is motivationally hot, and so disposes its subject to activity.[8] Cognitive rhythms and cognitive trails are passionate: trails call to us, beckoning us, or rejecting us. The rhythm and trail structures are laden with affect: they, or their parts, are appealing or repulsive; alluring or horrifying, enticing or nauseating, intriguing or scary, captivating or revolting, engrossing or annoying, pleasing or disturbing… When animals follow their trails, they are not neutral or dispassionate; it's a kind of affective compulsion. And that's how it is with us: we find ourselves in a force-field of affect, being pulled and pushed, yummied and yuckied, attracted and revolted. The animal acting within the flows of the realm of mediation has at least the freedom of animals: sniffing, and probing; listening and touching; changing the environment through explorations of it. The world calls out to the experientially absorbed animal, who feels how the environment calls to it, affording and attracting here, and repulsing there; yielding and resisting. The freedom of animals is not just receptivity, for it involves the skills by which these mediational structures can be reconfigured to support a more even flow of activity, a better coordination, or a more faithful reproduction of the environment's characteristic patterns of activity. Mediational content is not only being experientially disposed to follow trails, but also to trail-blaze. The management of mundane normativity is the means by which we may achieve efficacy in the environment: being actively related to the environment in such a way as to well manipulate the

[8] Activity-valence and subjectivity-valence are not distinct *parts* of a mediational content, but different aspects, or different focuses that a theorist may have in specifying a mediational content.

environment's "dynamic potential born of the disposition and propensity of things" (adapted from Jullien 1995).

When the only explanatory primitives are *syntactic*, content has been disconnected from its materiality, and has no computational or causal potency. Under the formality constraint, I may think of the moon, but the moon itself has no causal role in my psychology. By contrast, a theory of cognition as a theory of the relations between referential content and mediational content uses amongst its explanatory primitives the niche-environmental notions of trails, rhythms and territory. These are notions of representational content, but also of environmental materiality. Within such a conception of cognition, content can be computationally effective.

Referentialist content, as we have seen, would come between a person and their object, as conditions that a particular object or intervention in activity would or would not satisfy. Referentialist representation is indirect. But mediation of activity by the environmental forms of activity involves no indirection. These mediational forms are spatial and temporal instantiations of the mundane normativities of acting: they are inherently motivational, and are experienced as disposing one (me or us) towards a particular intervention in activity. This combination of having subjectivity-valence and activity-valence is what permits a direct relation to what they represent.

This direct relation has a distinctive phenomenology. When the workman drills a hole in wood, the drill, like the body of the workman, are not *objects* of the activity, but mediate the activity. When the activity flows well, the tool and the body are phenomenologically

transparent: mediating the drilling is what allows a direct relation to the object of the activity, which might be the hole in the wood. More generally, a skilled animal at home in its environment of activity, has a direct representational relation to its environment because that relation is mediated by the mediational content of its experience; content which is phenomenologically transparent when activity is flowing well. In human animals, perception carries both mediational and referential content. The objects which are the referents of perception's referential content are given directly, because they are mediated by perception's mediational content.

When a mediator becomes an object, that may be a sign of a rupture or breakdown in the activity, a situation in which the distancing of a representational veil is cognitively useful, because it may allow the reflexion necessary to fix the breakdown. When the drilling goes badly, the drill becomes again phenomenologically present as a referent. This is a simple example of a transformation from mediational content to referential content. (Hutchins' examination of navigation serves to illustrate a more complicated transformation from a referential content to mediational contents).

Pain is often treated as a conscious sensation that is supposed to carry no representational content. But we can also analyze pain as a simple example of mediational content. Pain is not a system of disapprobation; a bodily response that involves a negative judgement of what causes the pain. A negative judgement is not intrinsically motivational, as is shown by smoking and many eating behaviors, but pain is intrinsically motivational. But like disapprobation, pain involves negative normative feedback. Nor is pain experience simply an alarm system. An alarm, like pain, solicits the attention of a subject but does not fix the significance of the alarm; of what would be an appropriate response to the alarm.

Likewise, an alarm — experienced simply as a warning signal — does not have activity-valence. As Dennett (1991, p. 61) put the point: "But why do pains have to hurt so much? Why couldn't it just be a loud bell in the mind's ear, for instance?" Pain is a very simple example of a NASAS structure: it contains negative *Normative* feedback, indicating the need for some intervention and guiding the intervention, it *Affords* a certain kind of activity (such as withdrawal), it *Solicits* our attention (this is the alarm element), it carries *Affect* (the hot motivational element: a subject who is experiencing pain does not require any further cognitive process in order to be motivated to act), and it has *Subjectivity*-valence (experienced pain is felt as my pain; it presents <me> just as much, perhaps more, than it presents the world).[9] Each of these NASAS elements is unified within a pain experience. It is the unification of NASAS that is characteristic of, and which yields the great power of, mediational content. Mediational content is content type-NASAS, not conditional content.

Yucky experience is not unlike pain experience, except that it is typically more unpleasant than painful. <Yucky> also has a NASAS structure. In <yucky>, I experience myself and not just the worm in my food. And this is characteristic of mediational content; it has not achieved referential content's separation of the subjective and the objective elements in content. Mediational content is a mangle of what relates to the subject with what relates to the world[10]. Once again, in mediational content the subjective and the objective elements are not independent; they have been unified into what I call 'subobjectivity' (Cussins, 1992). Trail experience likewise: a trail calls to us, soliciting our

---

[9] There exist clinical breakdowns in each of these aspects of the experience of an ordinary person.

[10] This quasi-technical notion of a 'mangle' derives from Pickering (1995).

attention, affording movement along the trail, drawing us towards it (positive affect), and providing a spatial structure of correctings of activity, or small negative feedbacks (the trail boundaries) which keep us on the trail.

There are different forms of body knowledge. Some body knowledge involves referential content, for example when a subject attends conceptually to a part of their body. But a hypothesis, to be tested, is this. That the basic role of body knowledge is in mediating those relations that under referentialism generate the veil of representation: the relation between cognition and action, and the relation between the environment and perception. The governing normativity for body knowledge is mundane normativity, and its characteristic content is mediational content. In yummy experience I experience a mangle of the environment and my body through subjectively felt, environmentally warranted, pullings. A mediational content that presents a trail carries embodied subjectivity-valence: it is subjective knowledge of the body, given not as an object but as a mediation of activity. In trail experience I experience my body, and not just a route through the environment: environmental representation of the body.

This distinction, between two forms of being given, or two forms of presentation, is the logical point at which I promised to arrive, and which explains the apparent incompatibility with which we began. If the only form of presentation is being given as an object (or other referent of a referential content), then representation entails a veil, and representationalism is incompatible with direct cognition. If, on the other hand, representational theory recognizes mediational content, and so includes presentations of

the environment, and of the body as environmental mediations of activity, then direct cognition is compatible with representationalism. Direct representationalism is possible.

---

5. Cognition in the Wild [11]

An important part of Hutchins's (1995) account is his hypothesis that effective navigational cognition consists in the transformation of a referential, navigational question (where is the ship? where, locally, is the ship going?) through a series of active, environmentally distributed, representational steps, into questions which can be directly answered by trained and disciplined human bodies. The success of a ship's navigation consists in its reconfiguring of the referential task, through a repeating cycle of coordinations of mediational contents, into local, mediational tasks as things that the body-in-the-environment knows immediately how to do: the alignment of a hairline in the alidade sight with an image of an external landmark, for example (p. 120), or aligning the straightedge of the one-armed protractor (a 'hoey') with a digital marking on the hoey's compass scale. And then putting this configured hoey into coordination with a navigational chart by aligning the base of the hoey with the directional frame of the chart, and the edge of the hoey arm with a symbol on the chart. This coordination is simplified by:

> placing the point of a pencil on the symbol of the landmark on the chart, bringing the edge of the rule up against the pencil point, and then, [with gentle

[11] I end with Hutchins as a suggestion for the reader to use examples of computational ethnography, of which Hutchins (1995) is a paradigm, in order to think through an account of cognition in terms of transformations between mediational content and referential content.

> pressure] keeping the edge in contact with the pencil point, moving the base of the hoey until it is aligned with the directional frame of the chart (p. 144).

These are tasks that can be solved directly through environmental body knowledge: perceptually salient alignments, coordinations, translations, and gentle pressures. In the days before GPS, the referential question might have seemed impossible to answer. No person on the ship contains beliefs adequate to infer the current location. The distribution of cognition throughout a structured physical, social and technological space serves to change the impossible referential question into a repeating cycle ("the fix cycle") of mediational questions that can be answered through the direct application of body knowledge. The 'forms' of the formality constraint are replaced by body/environment mediations. The navigational task-environment or problem-space has been transformed, through repeated iterations of the fix-cycle, into a mediational environment directly navigable by body knowledge.

REFERENCES

Acosta, D. (2012) *Perception, Content and Action*. Master's Thesis in the Universidad Nacional de Colombia, sede Bogotá.

Agre, P., and D. Chapman. (1987). Pengi: An implementation of a theory of activity. In *The Proceedings of the Sixth National Conference on Artificial Intelligence, American Association for Artificial Intelligence*. Seattle: Kaufmann, pp. 268 - 272.

Barwise, J. (1987). Unburdening the language of thought. *Mind and Language*, 2: 82–96.

Ballard, D., Hayhoe, M., Pook, P., and Rao, R. (1997). Deictic Codes for the Embodiment of Cognition. *Behavioral and Brain Sciences* 20: 723–42.

Brewer, B. (2011). *Perception and its Objects*. Oxford: Oxford University Press

Brooks, R. (1992). Intelligence Without Representation. In D. Kirsh (ed), *Foundations of Artificial Intelligence*. Cambridge, MA: MIT Press, pp 139-159.

Campbell, J. (2002). *Reference and Consciousness*. Oxford: Clarendon Press

Clark, A. (1997). *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.

Clark, A. (1999). An embodied cognitive science? *Trends in Cognitive Sciences* 3 (9): 345-351.

Clark, A. and Chalmers, D. (1998). The Extended Mind. *Analysis* 58: 1: 7-19.

Cussins, A. (1992). Content, Embodiment and Objectivity: the Theory of Cognitive Trails. *Mind*, 101, pp. 651-688. Reprinted in A. Clark and J. Toribio (eds), *Artificial Intelligence and Cognitive Science: Conceptual Issues*, Garland, 1997.

Cussins, A. (2003). Postscript: Experience, Thought, and Activity (2002). In Gunther, Y. (ed) *Essays on Nonconceptual Content*, Cambridge, MA: MIT Press.

Dennett, D. (1991). *Consciousness Explained*. Boston: Little, Brown and Co.

de Vignemont, F. (2009). Body Schema and Body Image - Pros and Cons. *Neuropsychologia,* 48(3):669-80.

Fodor, J.A., (1980). Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology. *Behavioral and Brain Sciences* 3, 1, 63-73. Reprinted in *RePresentations: Philosophical Essays on the Foundations of Cognitive Science*, Harvester Press, 1981.

Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press.

Jullien, F. (1995). *The Propensity of Things: Towards a History of Efficacy in China*. Translated by Janet Lloyd. New York: Zone Books.

Kirsh, D. (1995). The intelligent use of space. *Artificial Intelligence* 73:31-68.

Kirsh, D. and Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science* 18:513-49.

Menary, R. (2010). *The Extended Mind*. Cambridge, MA: MIT Press.

Pickering, A. (1995). *The Mangle of Practice: Time, Agency, and Science*. Chicago: University of Chicago Press.

Port, R. F., and van Gelder, T. (eds.) (1995). *Mind as Motion: Exploration in the Dynamics of Cognition.* Cambridge, Cambridge, MA: MIT Press.

Rupert, R.D. (2009) *Cognitive Systems and the Extended Mind* Oxford: Oxford University Press

Suchman, L. A. (1987). *Plans and Situated Action*. New York: Cambridge University Press.

Thelen, E., and Smith, L. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press.

van Gelder, T (1995). What Might Cognition Be, If Not Computation? *Journal of Philosophy* 92: 345–81.