

The functionalist's body

Interview with Robert D Rupert⁵⁹

by Przemyslaw Nowakowski & Witold Wachowski
The interview was realized in August–October 2014

In the history of philosophy and science, we may find many works that have been forgotten, even though it would seem that they anticipated some of the ongoing research in a particular way. Do you have your favorite older works that you think deserve to be appreciated (especially in the context of embodied or embedded cognition)?

I think Ron McClamrock's *Existential Cognition: Computational Minds in the World* deserves more attention than it has gotten in the 4E literature. It was published in 1995 and, in my opinion, should be counted among the classics from the dawn of the current 4E movement.

As we know, you support the embodied mind theory, but you remain skeptical about the extended mind theory. What properties of the body of do you consider particularly important for shaping our cognition?

The bodily properties that influence cognition the most are, I think, structural properties of motor commands and of our sensory-based interactions with the world. Of particular importance are structural properties having their roots in the temporal, spatial, and more generally geometrical profiles of motor commands and routines and sensory-based bodily interaction with the environment.

I support a weak form of the embodied theory, according to which neural representations of, or neural simulations of, bodily activity play a central role in cognitive processing. I don't claim that *all* human cognitive processing is carried out entirely in a sensori-motor code, but I suspect that a surprising large amount is. To the extent that our cognitive processing does involve representations in an amodal format (a non-motoric code that is not specific to any

⁵⁹ Professor of Philosophy, Department of Philosophy, University of Colorado at Boulder. Professorial Fellow, School of Philosophy, Psychology, and Language Sciences, The University of Edinburgh.
Robert.Rupert@Colorado.edu

sensory modality), I think that sensori-motor interactions cause those amodal representational units to appear in the cognitive system; they cause what might have otherwise been little more than snippets of neural noise to take on functional roles in cognition, and this process also affects the processing profiles of the amodal concepts so acquired. This suggests one clear way in which the nonneural body shapes the mind, at least if we take “shape” to be used in a standard causal sense.

In addition, bodily experience or activity causes the formation and refinement of bodily representations—involved in motor control, bodily contact with the world, correlated sensory experience—and these representations (or their neural realizers) affect cognitive processing both via their causal connection to the formation of amodal representations (or to the realizers of amodal representations) and by the way in which they contribute to the control of behavior alongside co-contributing amodal representations that are directed toward the same objects as are the co-active, co-contributing bodily representations.

I’m not convinced, though, that the nonneural body does any actual cognitive processing or is properly part of the cognitive system. For example, when arguing for a strong embodiment thesis, Andy Clark appeals to Susan Goldin-Meadow’s work on the cognitive role of gestures, which seems to show that gesturing during problem-solving changes the nature of the cognitive process; on this basis, Clark argues that the gestures literally constitute part of the relevant cognitive processing. I’m just as inclined, however, to think that what enhances cognition in these cases is the efferent copy of the motor command to gesture, not the gesture itself. These are open empirical questions, though, and it might be that the hands (one of the most likely candidates) contribute extensively to actual cognitive processing.

Jerry Fodor said: "Who could doubt that the mind is embodied? And given that we are all clear that the mind is embodied, where does that get us? I mean, everybody knows the mind is embodied, unless you’ve got religious stuff or something, some metaphysical or ideological biases, but now, right, I agree, the mind’s embodied, and now what do I do? Search me!"

Do you think that the fact of embodied cognition modifies our view of cognitive processes (as in Shapiro’s replacement hypothesis), or when we describe cognition as embodied, we merely complement the classic cognitive research?

Yes, the embodied perspective substantially modifies our view of cognitive processes and, at the same time, it may still be that the embodied view merely complements classic cognitive research. Let me explain.

I think the embodied perspective modifies our view of cognitive processing in at least the three following ways:

First, with regard to historical interaction and theories of representation. Here the embodied perspective helps us to see how the nonsemantically individuated units of cognitive processing arise and how that history affects not just the processing profile of those units—which I emphasized in response to the preceding question—but the representational value of those units. For example, by understanding the conditions under which some bit of neural flux becomes a functional cognitive unit, we see simultaneously the sensorimotor interactions that establish the representational content of that unit, that is, we locate the property, kind, or individual in the environment (typically) that the unit represents. The nonsemantically individuated unit represents the thing (broadly speaking) interaction with which led to the firming up of that unit as a cognitively relevant unit.

Second, with regard to the processing profile of cognitive units. For any representing unit that comes to play a role in cognitive processing, there are conditions under which that happens, and those conditions affect what that unit ends up doing. Moreover, these conditions normally have a bodily component, the nature of which can then contribute to the processing profile of the unit the use of which is thereby reinforced. These might be conditions that reinforce certain movements over others or cause certain kinds of experiences or thoughts to be more closely associated than they would have been.

Third, with regard to architecture. The embodied literature draws our attention to a wealth of results that implicate bodily processes in cognition. I think these results are best interpreted as part of a larger body of evidence, emerging from all quarters in cognitive science; this body of evidence suggests something about the human cognitive architecture, that it contains a massive number of representations, many families of which co-represent (they represent the same object, individual, or property), and that much of the variation in behavior we associate with conscious attention, or lack of it, is to be explained in terms of variation in the sheer number of co-representing units active. (Imagine lots of sticks being pushed from slightly different angles in an effort to get a huge stone to move; the more contributing sticks, the higher the chance the stone will go where it's "meant" to go. Similarly with mental representations: the sticks correspond to co-referring representations, and the moving of the huge rock is the body's skillful or attentive interaction with the environment.) The embodiment-related data seem to show that what we think of as distinctively bodily stimuli (including actually induced movement) increase the number of active neural representations in bodily formats or in somatosensory cortex. This results in a stronger contribution to behavioral control from those areas, but in the typical case, these units work in conjunction with co-referring representations not in a bodily code in order to control behavior.

At the same time, I sympathize deeply with the sentiment behind Fodor's remarks, partly because none of the contributions listed above go so far as the replacement hypothesis. Even when modelers and experimentalists focus on the interface between the organism and the environment—and here we might begin to think that bodily-action-in-the-world replaces the need for computational processing—the typical cognitive model remains of the same sort as traditional computational and representational ones, at least with regard to its abstract characteristics of the sort Fodor has historically cared about.

Here's what I have in mind. So far as I can tell, most proponents of strong embodiment theses misunderstand the metaphysical relation of their views to classical cognitive research; as a result, they claim that their strongly embodied views stand at odds with functionalism and computationalism (and they typically disparage these classical views as "disembodied"). This seems to me to be simply wrong. According to the classical view, the physical body (the brain, in particular, but this applies to whatever matter realizes the computational system) determines, in the strict metaphysical sense, which functions the human cognitive system computes. That is the very nature of the realization-relation. Thus, on the classical view, there's a one-way determination relation running from body to mind; the body is very deeply "in charge" and is the root of our cognitive beings. Therefore, unless we construe embodied theories as type-type identity theories—which raises a host of problems—embodied views are a complement to orthodox views. Embodied approaches do make an important contributions; they entail that human computational (or otherwise functional) cognitive processes are best described in a fair degree of detail that, practically speaking, can be discovered only by examination of bodily structures and the ways in which humans interact with their environments. But, so long as the embodied theorist isn't offering type-type reduction (the identification of types of cognitive or mental states or processes with bodily types of states and processes), the embodied view is no less functionalist, and no less disembodied, than the orthodox view, at least metaphysically speaking.

Granted, epistemologically speaking, the embodied approach recommends a different methodology than was pursued by many classical computationally oriented cognitive scientists. The embodied functionalist lets bodily activity be her guide, epistemically, when attempting to figure out which algorithms govern human cognition and where they're realized (and how many shortcuts the embodied computing system can take given stabilities in the environment). But, this has nothing to do with the truth of functionalism but rather it stands in opposition to a certain empirical bet that many computationalist-functionalists made in the early days of cognitive science: that the relevant algorithms and the location of the machinery that executes them could be identified from the armchair.

Although there is a lot of research on embodied cognition, some authors still explore the relationship between these issues and the classical epistemological issues. Do you think that we should have a specific embodied epistemology? Or would that be a completely unnecessary effort, as we already have conceptions compatible with the notion of embodiment, and if so, which concepts may be used in this way?

I do think a situated perspective has some important implications for traditional epistemology, partly by bolstering the general case for externalism about justification (or warrant). But, to my mind, one of the most interesting ways in which the embodied perspective enters into epistemological debates is via its bearing on philosophical methodology, that is, on the epistemology of philosophy itself. Timothy Williamson has argued that the supposedly exhaustive dichotomy “*a priori* versus *a posteriori* justification” is sometimes inapt, that its application doesn’t shed light on philosophical matters of great interest; such matters involve possession and application of philosophical concepts (such as the concept of justification itself) that cannot be reduced to a sensory base and at the same time do not ground analytic truths known *a priori* by all who grasp these concepts. Williamson emphasizes the extent to which bodily experiences shape our intuitions about concept application (to hypothetical cases, for instance) and argues that two individuals who share the same concept might have had their intuitions shaped by different series of interactions with the environment (and the differences in these interactions is not a matter of having different stored sensory impressions, in the traditional empiricist’s sense). Perhaps without intending to, it seems to me that Williamson presents an embodied conception of the formation of philosophical intuition and the way it guides responses to thought experiments. People who share the same concepts, by many standard epistemological or semantic measures, have differing embodied experiences connected to that concept. As a result they make conflicting judgements about, or have contrasting intuitions in response to, philosophical thought experiments. These embodied experiences don’t constitute part of a “reduction base” of experiences out of which concepts are built in empiricist fashion, but neither are the judgments they lead subjects to make constitutive of the concepts in question and given *a priori*. I think Williamson’s exactly right (or at least, I think he is on my gloss of his work!).

Another important epistemological debate concerns the status of the embodied approach itself. On some more extreme versions of the embodied view (Lakoff’s, for example), the question of truth doesn’t arise. Folks who eschew truth and objective reality (and think instead that we enact our own worlds) should do more to explain the epistemological status of their own views or to develop a view of cognitive processing that gives an epistemological (and corresponding metaphysical) anchor to the embodied view itself. I would find the embodied view uninteresting if its strongest advocates could at best say “we believe in the embodied view because the fine-grained details of our body

made us do it"! So, I think there's pressure on philosophers to develop an embodied epistemology (and accompanying metaphysics) that underwrites their own theorizing and theoretical conclusions.

Which contemporary experimental works, in your opinion, are changing or will change our view of embodied (or embedded) cognition?

I think that much of the work on the neural basis of decision-making, planning, and abstract reasoning is of special importance, with regard to embodied cognition. It's helping us to gauge the extent, for example, of visceral contributions to these supposedly purely rational processes. We have to go beyond the mere observation of activation of motor cortex during so-called higher cognitive processing and discover the role (say, via articulatory encoding) that motor processes play in higher cognition.

Two threads of contemporary cognitive neuroscience seem especially interesting to me. The first involves the exploration of neural re-use and redeployment, the idea that what have been thought of as various functionally specific areas of the brain each participates in a fairly wide variety of forms of cognitive processing. The second involves decoding, the use of machine learning resources to extract signatures distinctive of different mental or cognitive processes (of, say, physical pain versus emotional pain, or of differing visual images). A combination of these kinds of experimental work has the potential to offer substantive embodiment-related insights. I'm inclined toward a hybrid view—to think that embodied representations and simulation make up only part of our cognitive resources, but resources that are constantly in use and contributing to the full range of cognitive processes. So, in my ideal world, decoding is used to identify the alphabet of neural computing (to identify the "workings" as Michael Anderson calls them), which I suspect will reveal that many of these fundamental components are in fact deeply connected to bodily experience and contribute their bodily content to a wide range of cognitive processes.

It seems that there are several versions of enactivism, or at least several approaches known as enactivism. Some advocates of enactivism emphasize the convergence, or even the conflation, of cognition and life⁶⁰. Others indicate that there are elements of constructionism in this approach. Still others see enactivism as based on law-like relations called sen-

⁶⁰ By the way, it is worth mentioning the words of G. Bateson: from an enactivist perspective "epistemology and theories of mind and theories of evolution are very close to being the same thing" [Bateson, G. 1987. *Men are grass: Metaphor and the world of mental process*. W.I. Thompson, ed. *Gaia, a Way of Knowing: Political Implications of the New Biology*. Hudson, NY: Lindisfarne Press: 37-47].

sensorimotor contingencies. Could you describe what valuable insights may be prompted by these perspectives in the studies on cognition?

In my opinion, the most instructive versions of enactivism emphasize the contingency of the range and kinds of properties in the environment that an individual human comes to interact with, exploit, and represent. It serves no good purpose, in my opinion, for the enactivist to claim that the subject's acting on the environment „brings a world into existence.” In order for the interaction to take place, the world must already be there! But, the enactivist is right about the following aspect of cognition: which of the many, many properties in the world the subject represents or becomes sensitive to depends on the subject's previous interactions and her biologically given capacities and needs. There's a kind of canalization that takes place—over the course of individual development, over the course of an entire life (say, in the gradual acquisition of expertise), and over the course of smaller-scale cognitive interactions (in which one kind of output from the agent renders salient to the agent a different aspect of the environment than would have been salient to her had she taken a slightly different action, with a „snowballing” effect).

“Alternative” approaches in cognitive sciences are undergoing a sort of a conceptual turmoil. Some authors propose a notion of “3E Cognition” or “4E Cognition” (Embodied, Embedded, Extended, Enacted) as a kind of a core, or a sheltering umbrella. Are your arguments—including, among others, in *Cognitive Systems and the Extended Mind*—insufficient?

In chapter 1 of the book, I challenge the utility of that umbrella. The various approaches are genuinely distinct; and in some ways, they are in tension with each other while in many other ways they are „mix and match.” For example, an extended cognitive system might or might not proceed with its cognitive tasks in a heavily embedded way. Having distinguished these various views, I proceed, in the first two major sections of the book, to raise various objections to the extended view in particular and to the arguments in support of it, especially when it's offered as a new paradigm in cognitive science or the basis of a scientific revolution. But, I also offered a positive theory of cognition and cognitive systems (in the book and in more recent papers), one that grounds various objections to the extended view (and of arguments in support of it) but that also allows, at least in principle, for the appearance of extended cognition. I think that's still where the situation stands.

I'm a staunch supporter of the embedded view, and combined with my view of cognition as the activity of a persisting collection of integrated mechanisms, it seems to entail that human cognition takes place inside the organism but depends on and exploits external (noncognitive) resources to a surprising extent. I'm quite impressed by much of the same empirical work proponents

of the extended mind are impressed by (Ballard's work on deictic pointers, for instance); in my opinion, though, most of that work supports the organismically oriented embedded view, as opposed to the extended one.

I'm also quite committed to a weak embodied view of the human cognitive system. Many interpreters of body-related experimental work claim that it reveals the bodily constituents of cognition or that it shows that linguistic meaning is inherently embodied. I draw a different conclusion. To my mind, the experimental results on embodiment show that the cognitive system contains an enormous number of representations that are being activated in parallel, and many of them represent the same thing (take the same object, or co-refer); moreover many of these appear in body-related areas of cortex or represent bodily processes, and their activation alongside amodal but co-representing units affects cognitive processing in measurable ways.

I think the lumping of all of the E-views together has the potential to cause a lot of confusion; after all, some of the most well-known enactivists reject representation, yet many proponents of other 4E views embrace representations. So, it might seem like these various approaches share too little to constitute a coherent view. Then again, some of the basic conceptual ideas behind many of these approaches may provide a dovetailing vision of cognition. And here I mean to go beyond just the idea that 4Es have a common enemy, a certain kind of traditional cognitive science (perhaps a straw man, in actuality). Rather, I'm claiming that, if wedded together, one might get a picture of human cognition that looks very different from what one might have extracted from such early models of cognition as the General Problem Solver. Here's the possibility I have in mind: There is an organismically centered cognitive system that relies heavily on simulations of bodily processes to guide the active exploitation of resources in the environment, most often taking up those resources in a fleeting way that does not render those resources truly cognitive, but sometimes in durable ways that add new external components (TVSS perhaps?) to the otherwise bodily bounded cognitive system.

How do you think, are you rather seen as a critic, or rather as an advocate of any approach?

I would guess that more philosophers know me as a critic of the extended view (and the view that groups of humans have minds or cognitive systems) than as anything else. But, I think some of my other positive work is reasonably well known, for example, the Best Test Theory of representational content, the component-forces theory of *ceteris paribus* laws, and my developmental views about the origins of representations. And, it's worth noting that some distinctively positive views have come out of my critical work on extended cognition, including a theory of cognition (that it is constituted only by activity

in a collection of mechanisms that contribute in overlapping subsets to the production of a wide range of forms of intelligent behavior). Grappling with embodied views has also led to the formulation of my view that the mind is massively representational and wholly subpersonal (in the sense that there is no metaphysically distinct personal level or even one of much epistemological import). Thinking about the overarching nature of cognition, as a generic kind, has also catalyzed a model-based theory of natural kinds, the „tweak-and-extend” theory, found in some of my most recent work.

How do you perceive the relationship between philosophy and science? And what do you think is the value of interdisciplinarity?

I think the relationship has many facets. Philosophers familiar with scientific results and theories can interpret that work for other philosophers and for the educated public, by bringing it to bear on issues of long-standing philosophical and popular interest. Some such work can be quite revolutionary, changing our everyday conception of selves or of social interaction, for example. At the same time philosophers can play the role of theorists in the sciences themselves, offering precise accounts of foundational concepts (such as that of representation) or attempting to help straighten out inconsistencies in the use of those concepts. To my mind, though, what’s most interesting, challenging, and rewarding results from a kind of freedom philosophers have to survey the landscape and attempt to provide an integrated and coherent vision of some large domain, identifying a pattern that appears only when one considers a wide variety of work being done on different sub-topics in a field or in different scientific fields altogether.

What sorts of gadgets do you usually use in order to “extend” your own mind? We mean both science-related extensions as well as the ones that you use solely for entertainment. Do you read a lot of printed books? Do you still write using pen and paper, and if so, then how often?

I do most of my writing on my laptop, and I use my iPhone quite a bit. I use pen and paper almost every day, mostly for doing quick calculations, making lists, or writing personal reflections and notes. I read only paper books, although I do often read journal articles on the computer screen. I’m even more of a Luddite when it comes to entertainment. In that domain, I don’t use any “mind-extending” technologies (although it depends on how liberal an approach we take—is a guitar mind-extending technology?).

What educational and academic choices guided you to the subjects you are currently interested in? Did you experience any dramatic change in research interests during your career? What did you imagine your future would turn out to be like when you were a teenager?

I was very much interested in psychology and the human condition when I was young, but largely through literature, history, movies, and songs—in other words, from a more humanistic perspective. But as I began taking college courses, analytic philosophy seemed to promise the most enduring or abstract truths, and it also drew on my longstanding fascination with analytical puzzles and games. I think those inclinations ultimately led me to a scientific approach to the exploration of content, consciousness, and cognition. The sort of philosophy that offered the deepest and less purely speculative insight into the mind was oriented toward the sciences. I might add that in my first philosophy of mind course, the professor (Charles Marks) assigned Nisbett and Wilson's 1977 paper, "Telling More Than We Can Know: Verbal Reports on Mental Processes," which instilled in me a healthy skepticism regarding our "pure" philosophical intuitions about the workings of the mind.

My current research interests grew initially out of my fascination with logic and philosophy of language. These latter interests turned me toward the mind at the very end of my undergraduate education when I encountered externalist semantics for mental states. Then, in the summer after I finished my bachelor's degree, I attended a seminar on mental content co-taught by Scott Soames—who was at that time a serial visitor in my department—and Charles Marks, my department's resident philosopher of mind. I had been thinking extensively about the indeterminacy of reference and the causal theory of reference, and the way some of these issues connected up skeptical views in epistemology and philosophy of science. Then, in the seminar with Soames and Marks, I saw how to cast many of these issues as questions about mental content and the mind's causal connection with the word. It was in this course, too, that I had my first substantive exposure to Fodor's work and began to think of philosophy of mind as deeply empirically informed.

When I was a teen-ager, I wanted to be a musician, a folk rocker like Bob Dylan or Neil Young—and I wanted play lead guitar like Jerry Garcia. But I was realistic with regard to my career options. I had intellectual leanings and abilities and assumed they would lead to stimulating remunerated work, probably to do with politics or history, that would in some way put me into contact with great texts of history and the community of human thinkers. But, to be honest, none of these thoughts took the form of concrete commitments or plans when I was a teenager. I was carefree and idealistic. I wanted to play music and read books and think deeps thoughts, and my attitude was, well, whatever happens happens. (I must admit that I encouraged my own children to give their adult futures more careful consideration!)

Bibliography

- Anderson, M. L. 2010. Neural reuse: A fundamental organizational principle of the brain. *Behavioral and Brain Sciences*, 33(04): 245-266.
- Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. 1997. Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, 20(04): 723-742.
- Clark, A. 2013. Gesture as Thought? Z. Radman, ed. *The Hand, an Organ of the Mind: What the Manual Tells the Mental*. MIT Press.
- Fodor, J. 2011. An Interview with Jerry Fodor. *Rutgers Philosophy Newsletter* (Spring): 8-10
[http://philosophy.rutgers.edu/dmdocuments/Rutgers%20philosophy%20newsletter%20\(Spring%202011\).pdf](http://philosophy.rutgers.edu/dmdocuments/Rutgers%20philosophy%20newsletter%20(Spring%202011).pdf) (10.11.2012)
- Lakoff, G., & Johnson, M. 1999. *Philosophy in the flesh: The embodied mind and its challenge to western thought*. Basic Books.
- McClamrock, R. 1995. *Existential cognition: Computational minds in the world*. University of Chicago Press.
- Nisbett, R. E., & Wilson, T. D. 1977. Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84(3): 231.
- Rupert, R.D. 2009. *Cognitive Systems and the Extended Mind*. Oxford University Press.
- Shapiro, L. A. 2011. *Embodied cognition*. London: Routledge.
- Williamson, T. 2013. How Deep is the Distinction between A Priori and A Posteriori Knowledge? Casullo, A., & Thurow, J. C. *The A Priori in Philosophy*. Oxford University Press.