21

c00021

On the Grounds of (X)-Grounded Cognition

Michael L. Anderson

Department of Psychology, Franklin & Marshall College, Lancaster, PA; Institute for Advanced Computer Studies, Program in Neuroscience and Cognitive Science, University of Maryland, College Park, MD, USA

p0010

For the least the last 10 years, there has been growing interest in, and growing evidence for, the intimate relations between more abstract or higher order cognition—such as reasoning, planning, and language use—and the more concrete, immediate, or lower order operations of the perceptual and motor systems that support seeing, feeling, moving, and manipulating. A sub-field of the larger research program in *embodied* cognition (Clark, 1997, 1998; Wilson, 2001; Anderson, 2003, 2007d, 2008; Gibbs, 2006), this work has generally proceeded under the banner of *grounded* cognition, and works to support the claim that thinking is inherently tied to—grounded in—perceiving and acting. Thus, Glenberg and Kaschak (2002) discuss "grounding language in action"; Gallese and Lakoff (2005) argue that concepts are "grounded in the sensory–motor system;" and Barsalou (1999) at various times talks of "grounding cognition in perception," "grounding conceptual knowledge in modality-specific systems" (Barsalou et al., 2003), and most recently simply of "grounded cognition" (Barsalou, 2008).

p0020

Yet despite a great deal of terminological consensus, in fact there are nearly as many theories of grounding—what it is, and what it means—as there are theorists. Some, like Glenberg and Kaschak stress the origin (and continuing importance) of cognitive structures in controlling bodily action (Glenberg, 1997; Glenberg & Kaschak, 2002; see also Anderson & Rosenberg, 2008). The view seems to be: cognition is grounded in x (e.g., action) if some of x's elements are deployed in guiding it. Thus, their theory of language comprehension emphasizes the capacity to process the affordances associated with sentence elements to generate representations of possible actions. Insofar as these affordances both

guide action and guide comprehension, comprehension is action-grounded. Others, like Barsalou and his many collaborators, suggest that the relation is one of *simulation*: cognition is grounded in x if it requires, depends upon or otherwise involves simulations of x experiences. "As people represent *TREES*, for example, they [engage] … a partial reenactment of the perceptual, motor, and introspective states that occur as people actually experience trees."[1] (Barsalou et al., 2005, p. 251; see also Prinz, 2002) And finally, there are many advocates for *conceptual metaphor* theories, which hold that cognition is grounded in x just in case it inherits from domain x an inferential structure that limits and guides one's thinking (Lakoff & Johnson, 1980, 1999; Lakoff & Núñez, 2000; Fauconnier & Turner, 2002). For instance, there are some apparent similarities between our notions of a *purpose* and of a *destination*—we imagine a goal as being at some place ahead of us, we plan a route, we imagine obstacles, and we set benchmarks to track our progress. Thus, the argument goes, our thinking about purposes is grounded in our experience of moving through space.

p0030 One thing that is especially interesting about the current state of affairs is that despite significant differences in the underlying theoretical frameworks, there is little overt disagreement between the various camps. Each approvingly cites the work of the others, and often even casts their own theories in the others' terms. For instance, Glenberg and Kaschak (2002) cite Barsalou (1999), and explain that their theory also "proposes that language is made meaningful by cognitively simulating actions implied by sentences" (p. 559)—this even though their theory is entirely prospective, rooted in what actions are *possible*, and requires no re-enactment of any specific experience. Moreover, they *also* indicate that their findings are compatible with the more metaphorical-structuring-friendly framework of Talmy (1988), whereby causal sentences are understood by analogy with contrasts between agonists and antagonists. Similarly, in the course of developing a version of conceptual grounding that attempts to combine elements of both conceptual metaphor theory and simulation-based accounts, Gallese and Lakoff (2005) cite Glenberg and Kaschak (2002) as providing evidence for their view. Indeed, even theorists who are suspicious of grounded cognition—such as Lera Boroditsky, who is explicitly critical of the inferential move from evidence for

[1]In the essay being quoted, words in italics and all capital letters indicate concepts. Thus *TREES* is the categorical concept to which all and only trees belong. Given this, the quote should probably actually read: "When people represent trees with *TREES* …" but the meaning is nevertheless clear.

[2]See, Boroditsky and Ramscar (2002, p. 188): "[C]ontrary to the very strong, embodied view (that abstract thought is based directly on sensory motor representations), we found that actual motion was neither necessary nor sufficient to influence people's thinking about time. Rather it is thinking about spatial motion that seems to influence thinking about time. It appears that thinking about abstract domains is built on representations of more experience-based domains that are functionally separable from representations directly involved in sensorimotor experience itself." This is a significant dissent in the context of Boroditsky's other work, since she is arguing in part that this functional separation allows room for specifically cultural influences to shape cognitive structures. Abstract thought is not grounded in concrete, embodied experience, but in abstract, culturally influenced *thinking* about embodied experience. This line of reasoning appears to point in precisely the opposite direction from that advocated by the theorists of grounded cognition.

metaphorical structuring to evidence for experiential grounding[2]—are typically cited as providing further evidence in favor of grounded cognition, without mention of their dissent.[3] As supportive and collegial as this makes conferences on embodied and grounded cognition, I would like to submit that it is not the most scientifically productive situation. The theories noted earlier (and these represent only a small fraction of the work in this area) are different enough that they ought to make differentiating predictions. The failure to highlight and test for the different predictions made by competing theories of cognitive grounding represents a missed opportunity to challenge and improve those theories. And if indeed all the scientific evidence gathered so far supports all the various theories—and one can sometimes get this impression, reading the literature—that suggest there is something wrong with the evidence, with the theories, or with both.

p0040 One particular domain of evidence that has been over-generously interpreted in support of various theories of grounded cognition is neuro-imaging data. Consider the interesting fact that mental planning can activate higher motor areas even when the planning itself involves no motor activity (Dagher et al., 1999). Anderson (2007d) cites this finding as support for the view that "many, if not all, higher-level cognitive processes are body-based in the sense that they make use of (partial) simulations or emulations of sensorimotor processes through the re-activation of neural circuitry that is also active in bodily perception and action" (Svensson et al., 2007)—but then adds a footnote noting it also supports Lakoff and Johnson (1999). Similarly, each of the following findings has been cited in support of conceptual metaphor theory (Gallese & Lakoff, 2005), *and* in support of simulation-based views (Barsalou et al., 2003; Barsalou, 2008):

u0010 • Evidence that watching actions, imagining actions, and doing actions all activate similar networks of brain regions (Decety et al., 1990; Jeannerod, 1994; Decety et al., 1997; Decety & Grèzes, 1999).

u0020 • Evidence that brain areas involved in motor control functions are also activated in verb retrieval tasks, while naming colors and animals activated brain regions associated with visual processing (Damasio & Tranel, 1993; Martin et al., 1995, 1996).

u0030 • Evidence that perceiving objects and object names activates brain regions associated with grasping (Chao & Martin, 2000).

p0080 If all these findings do indeed support the various theories, it suggests that brain imaging data isn't all that useful a tool for distinguishing between theories of grounded cognition; and if they do not, it suggests that very many cognitive scientists—myself most certainly included—have been less than careful in evaluating the available evidence. So, what's going on? My own impression of the situation is that most tests of embodied/grounded cognition are not targeted

[3]For instance, Barsalou (2008, p. 621) writes: "Lakoff and Johnson (1980, 1999) proposed that abstract concepts are grounded metaphorically in embodied and situated knowledge. . . Increasing evidence suggests that these metaphors play central roles in thought (Boroditsky & Ramscar, 2002; Gibbs, 2006)."

tests of the specific theory under consideration; rather they are designed to differentiate predictions made by a generic theory of grounded cognition from a generic theory of abstract, computational, amodal, or otherwise largely cognitivist theory of cognition. At this, they are effective. Even critics of embodied cognition have largely ceded the cognitivist high ground, and are fighting rear guard actions against further extensions of the basic claims (see e.g., Rupert, 2004; Weiskopf, 2007). But we theorists of embodied cognition ought by now to be in a position where the embodied, situated, and distributed approaches to the study of the mind are seen not primarily as criticisms of some prevailing paradigm, but as established, vibrant and fruitful research programs in their own right, needing no justification other than their own success. The proliferation of presumably incompatible, but apparently equally well-supported models of grounded cognition is one concrete and (at least partly) harmful result of the general failure to move beyond the idea of providing an alternative to cognitivism, and toward building a general, unified, and specifically supported theory of cognition on embodied first principles.

p0090 Allow me to suggest that the best way forward is first to take a step back. For in the course of developing my own version of x-grounded cognition (where x = "action"), I came to a somewhat surprising—and certainly sobering—conclusion: most (if not all) of the brain imaging results, and much (if not most) of the behavioral data taken to support specific theories of grounded or embodied cognition can in fact be accounted for by a generic theory of the evolution of the cortex. It's called the massive redeployment hypothesis (MRH), and its fundamental tenet is that the human cortex evolved by neural exaptation, whereby circuits originally developed to serve some specific purpose are used for new purposes and combined to support new capacities, without disrupting their participation in existing programs.[4] Such redeployment of existing neural circuits is favored by straightforward considerations of efficiency, and MRH need not (and indeed, cannot) make any general assumptions about the functional implications of reuse. Whether a given instance of neuronal reuse results in a circuit that implements simulation, or supports metaphorical structuring, or reproduces the process of affordance meshing in another domain, or simply indicates that some low-level computational function is being borrowed is not something the mere fact of redeployment can adjudicate.

p0100 All this will be made clearer by a more detailed account of MRH and the evidence that supports it—something I will briefly offer in the next section—but I want to say at the outset that I do *not* believe that MRH undermines any theory of embodied or grounded cognition. In fact, I think MRH has three very important things to offer: first, it identifies what might be characterized as the generic grounds for x-grounded cognition. The discovery of frequent redeployment of

[4]MRH is related to (but more general and, I think, more empirically supported than) both the theory of neural exploitation proposed by Gallese and Lakoff (2005) and the neuronal recycling hypothesis developed by Dehaene and Cohen (2007).

neural circuits across many different domains—action, perception, language, reasoning, and the like—is indeed an important fact that needs to be more fully understood and assimilated into prevailing theories of cognition.[5] Second—and precisely because MRH makes clear that the mere fact of redeployment is not by itself evidence for embodied cognition, or for any more specific relation of grounding, but is only a starting point for further inference and investigation—it focuses attention on uncovering the precise nature of the inheritance that redeployment enables. That is, if a later-developing cognitive function—say, verb retrieval—reuses cortical circuits originally developed for motor control, then clearly there will be some sort of functional inheritance as a result. But *what* sort, exactly? Only answering this question, and not just noticing the overlap, will provide specific evidence for any more substantive account of grounded cognition (or tell us whether grounding is even the best metaphor to be using in a particular case). Third, MRH suggests a method for actually answering such questions, by leveraging cross-domain cognitive modeling to attribute functional roles to redeployed circuits. This has the potential to tell us interesting things not just about the newer cognitive function(s) in which a given circuit was redeployed, but about the older function in which it was originally developed. That is, discovery of specific functional inheritances between language and motor control, or categorization and perception, or any such similar x and y, will tell us something interesting about *both* domains.

s0010

THE MASSIVE REDEPLOYMENT HYPOTHESIS

p0110

As indicated earlier, MRH is both a theory about the functional topography of the cortex, and an account of how it got that way. According to MRH, neural circuits evolved for one use are frequently *exapted* for later uses, while retaining their original functional role. That is, the process of cognitive evolution is analogous to component reuse in software engineering (Heineman & Councill, 2001). Components originally developed to serve a specific purpose are frequently reused in later software packages. The new software may serve a purpose very different from the software for which the component was originally designed, but may nevertheless require some of the same low-level computational functions (e.g., sorting). Thus, efficient development dictates reuse of existing components where possible. Note that in such reuse, the component just does whatever it does (e.g., sorts lists) for all the software packages into which it has been integrated, even if that computational function serves a very different high-level purpose in each individual case. This important aspect of component reuse in software engineering is also part of the hypothesis as it applies to neural redeployment.

p0120

The end result of such reuse in the brain is a structure in which brain areas are typically recruited to support many different functions across cognitive domains.

[5]See Stewart and West (2007) for some work along these lines.

Such a story about the organization and development of the cortex has some interesting implications for its overall functional architecture. For instance, on this theory, brain areas are not domain-restricted entities. If this were not the case, if a typical brain region in fact served a very limited set of cognitive functions, then this would suggest instead a localization-based development, whereby the brain evolved by generating new, dedicated regions for each new purpose. Moreover, according to MRH we should expect that differences in domain functions will be accounted for primarily by differences in the way brain areas cooperate with one another, rather than by differences in which brain areas are used in each domain. If neural circuits do not change their roles when they are exapted, and if they are used in many different cognitive domains, then the only way to get different functions while using the same components is to put them together in different ways. Another straightforward consequence of MRH is that more recently evolved cognitive functions will utilize more, and more widely scattered brain areas than phlyogenetically older functions. Again, the reason is simple: the more established neural components there are when a given cognitive capacity is evolving, the more likely that one of them will already serve some purpose useful for the emerging capacity, and there is little reason to suppose that the most useful areas will be grouped together (and less reason to suppose this as evolutionary time passes, making available more functions supported by more areas). Finally, MRH predicts that evolutionarily older brain areas will be deployed in more cognitive functions. This is presumably because the longer an area has been around, the more likely it will have proved useful to some evolving cognitive capacity, and be incorporated into the functional network of brain regions supporting the new task.

p0130 Preliminary investigations have uncovered evidence for all four of these predictions (see Anderson, 2007a, b, c; Anderson et al., in press for details of the methods and results). For instance, in my lab we recently performed a coactivation analysis of 472 brain imaging experiments (representing about 10 years worth of studies from *Journal of Cognitive Neuroscience*) from 8 different cognitive domains (action, attention, emotion, language, memory, mental imagery, reasoning, and visual perception). We coded the results of each experiment in terms of which Brodmann areas were activated (using only post-subtraction activations), and determined the baseline chance of activation for each area (and for each possible pairing) by dividing the number of experiments in which it was activated by the total number of experiments in the database. Then, for each pair of Brodmann areas, we used a chi-squared measure to see if their observed degree of coactivation (in a given domain) was significantly different from what would be predicted by chance. We also performed a binomial analysis, since a binomial measure can provide directional information. [It is sometimes the case that, while area A and area B are coactive more (or less) often than would be predicted by chance, the effect is asymmetric, such that area B is more active when area A is active, but not the reverse.]

p0140 The idea is that if co-activation indicates functional cooperation, such an analysis should reveal the cortical networks supporting cognitive functions in

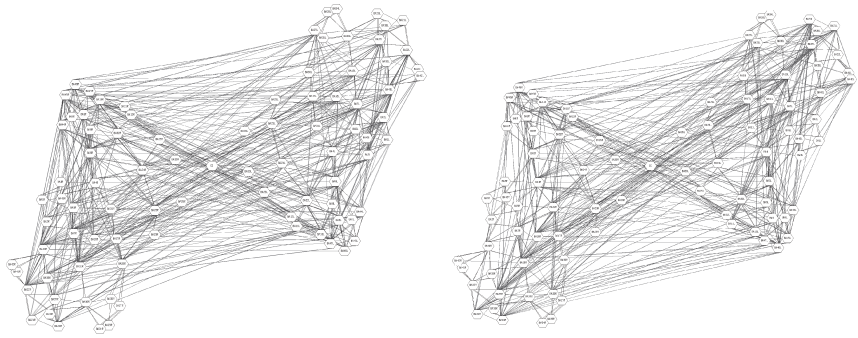

f0010 FIGURE 21.1 Cortex represented as adjacency + co-activation graphs. Here the Brodmann areas are nodes, with black lines between adjacent areas and grey lines between areas showing significant co-activation. The graph on the left shows co-activations from 56 action tasks, and the graph on the right shows co-activations from 57 visual perception tasks. Graphs rendered with aiSee v. 2.2.

the different domains. A graph offers a very nice representational format for the results. In the graphs we build, the nodes represent Brodmann areas, and the lines between nodes indicate significant co-activation (i.e., apparent functional cooperation). Figure 21.1 represents the co-activation graphs for the action and visual perception domains. The co-activation graph has been superimposed on an adjacency graph (where lines indicate that the connected areas share a physical border in the cortex) for ease of cross-domain visual comparison.

p0150 I mentioned earlier that MRH predicts that brain regions should support tasks across many different domains, but that the pattern of cooperation between the areas should be different in different domains. There is an obvious analog for these features in our co-activation graphs: comparing the graphs from different domains, node overlaps indicate Brodmann areas that support tasks in both domains, whereas edge overlaps would indicate a similar pattern of cooperation between Brodmann areas. Thus, MRH predicts a great deal of node overlap between co-activation graphs, but little edge overlap.

p0160 Using Dice's coefficient as our measure [$D = 2(o_{1,2})/(n_1 + n_2)$ where o is the number of overlapping elements, and n is the total number of elements in each set] and doing a pairwise comparison between all eight domains bears this prediction out. On average, there is very little edge overlap between the domains (Mean(D) = 0.15, SD 0.04), but a great deal of node overlap (Mean(D) = 0.81, SD 0.04); the difference is significant (2-tailed t-test, $P << 0.001$). Other specific results from the study include the fact that between action and visual perception the node overlap is 0.85 and the edge overlap is 0.14; between action and language we found a similarly high node overlap of 0.82 with an edge overlap of 0.06; and between visual perception and language the node overlap is 0.77 with an edge overlap of 0.17.

s0020

IMPLICATIONS OF MRH FOR X-GROUNDED COGNITION

p0170 The results reported above are interesting on their own, and certainly offer some support for MRH. But consider the following: insofar as MRH predicts high node overlap in this case, it *also* (by the same token) predicts a great deal of overlap of activations as seen in fMRI images of cognitive functions from different domains. That is to say, the fact that both language and action, and both language and perception, activate some of the same regions of the brain is an unsurprising consequence of the way the brain evolved; it does not represent, in and of itself, an anomaly that can only be explained by some version of x-grounded or embodied cognition.

p0180 And yet, there *is* a great deal of overlap in the networks supporting these functions—nearly as much as there is between perception and action, which one would expect to be intimately entwined indeed. This is by no means an insignificant fact. The question is, what shall we make of it? MRH itself has no answer to this question; it does, however, suggest a method for answering it. As noted earlier, the question of whether and how a given cognitive domain is x-grounded comes down to figuring out what that domain inherited from x; that is, it means knowing both what the shared neural components do in and of themselves (identifying their role), *and* what they are being redeployed to do (identifying their cognitive use). Given widespread overlap in the networks supporting different cognitive domains, MRH suggests that to determine the functional role of a given brain region it is better to focus on the brain region and consider its participation across multiple task categories. This is roughly the opposite of current practice, which generally involves choosing a given cognitive task (or task category) and identifying the various brain regions implicated in those tasks. Thus, rather than thinking about and modeling language functions in isolation from perception, attention, motor control, and other high-level cognitive domains, instead one needs to consider what sorts of components (and/or sub-functions) could serve functionality across domains. Finding the role of a given brain area will be something like finding the right letter to go into a box on a (multidimensional) crossword puzzle, determined not just by the answer to a single clue, but by all the clues whose answers cross that box. This makes the task both harder, because it is multiply constrained, but also easier, because it offers the possibility of leveraging information from several sources to make the attribution. For instance, the overlaps should suggest more fine-grained predictions about such matters as priming and cognitive interference, and this opens the possibility of designing experiments leveraging these overlaps, for example in further imaging, cross-domain priming, and interference studies.[6]

[6]Note the other implication, however: where there are shared neural resources, there will be interference. Thus, just as with MRH's implications for fMRI images, so too with its implications for cognitive interference; the mere existence of interference between cognitive domains like language and motor control is not an anomalous finding explainable only in terms of grounded or embodied cognition. Any such claims must rest instead on the exact nature of the interference in question.

p0190 This approach will have the effect of not just encouraging more creative hypotheses for the roles of brain areas, but also result in more integrated models of cognitive uses and domains. I have mentioned earlier the various theories supporting the idea that language is grounded in motor control. One common theme shared by all the accounts is that the action grounding of language implies that language is somehow *like* motor control—on Glenberg and Kaschak's affordance-mesh account, for instance, putting linguistic elements together in a meaningful sentence is like putting motor primitives together in an executable motor plan. But since this hypothesized functional inheritance is the result of shared neural circuits, and also these neural circuits are presumably playing the same role for each domain, this relation helps highlight a reverse implication that is worth considering: motor control should be in some way *like* language understanding. An affordance-based account suggests the following intriguing possibility: since affordances, the perceived availability of objects for certain kinds of interaction, aren't just motor programs, but features of the environment with specific significance for the organism, this opens the possibility that the motor control system is also, already, a primitive meaning processor (Gorniak & Roy, 2006). This would offer one explanation of how it is even possible to leverage motor control to support and constrain higher order processes like language understanding. After all, on a more mechanistic understanding of the nature of motor control, it would be hard to say why a motor control system would have *any* of the right basic elements for building a language-understanding system.

p0200 Of course, whether this is indeed the nature of the functional inheritance between language and motor control remains to be established; and the mere fact of overlapping neural circuits between the domains in no way does so. It nevertheless serves as a good example of the kinds of re-thinking that become possible when taking a more integrative approach. Much of this re-thinking will generate models that bear out the tenets of the embodied/grounded view of cognition—but not all of it. A case in point is some recent work on the relationship between finger gnosis and mathematical ability. Finger gnosis—the awareness of one's fingers—is commonly assessed via the ability to distinguish, without visual feedback, which fingers have been lightly touched. Developmentally, finger gnosis has been found to predict children's mathematics performance (for a review see Penner-Wilger et al., 2007), and studies have suggested that these two capacities are supported by some shared brain regions. For instance, Zago et al. (2001) found that a region associated with the representation of fingers (left parieto-premotor circuit) was activated during adults' arithmetic performance, and rTMS applied to the left angular gyrus that has been found to disrupt adults' performance on both finger gnosis and number magnitude tasks (Rusconi et al., 2005).

p0210 Now, any theorist with sympathies for embodied accounts of abstract cognition will be inclined to interpret this relation as an evidence for the grounding of mathematics in embodied experience, perhaps following Butterworth (1999) in the claim that using one's fingers to count causes one's finger representations

and number representations to become intertwined. But in this case, considerations of cross-domain modeling seem to point in a different direction. As Anderson and Penner-Wilger (2007; Penner-Wilger & Anderson, 2008) note, one foundational element in any calculating circuit is a register for storing the number(s) to be manipulated. Such a register is typically implemented as a series of switches that can be independently activated. Likewise, at least one way to implement the ability to know whether and which fingers have been touched (and other aspects of a general "finger sense") would be with such a register of independent switches. Such a finger register—one part of the functional complex supporting finger gnosis—would be a candidate for redeployment in any later developing complex with functional elements, able to take advantage of a component with this abstract functional structure. This, Anderson and Penner-Wilger suggest, is exactly what the number representation complex did.

p0220 There is some interesting evidence that seems to favor the view. Note, for instance, that Butterworth's position makes the experience of using the fingers to count a necessary condition for the observed intertwining, whereas the redeployment requires only that there be an intact finger register (i.e., intact finger gnosis) that can be put to various uses. Thus, the fact that children with Spina Bifida have poor finger agility co-morbid with significant mathematical difficulties (Banister & Tew, 1991; Barnes et al., 2005) has generally been taken as support for Butterworth's position. But since these children also have finger agnosia, the finding is equally compatible with the redeployment view. In contrast, children with developmental coordination disorder (DCD) have poor finger agility, but most have preserved finger gnosis, and do not generally evidence significant mathematical difficulties (Cermak & Larkin, 2001; Hamilton, 2002). This finding is consistent with the redeployment view, but appears to present some difficulties for Butterworth's position. Of course, these are just preliminary results, and whether the suggestion will be borne out by future investigations is an open question (for a discussion see Penner-Wilger & Anderson, 2008). But for current purposes, the point is twofold: (1) Such a proposal for one of the components of finger gnosis is unlikely to have occurred to researchers focusing only on results from their own domain; this may suggest the fruitfulness of the approach to modeling advocated here. (2) Not every overlap between cognitive domains is evidence that one is grounded in the other—at least not in the robust sense required by the various theories of embodied cognition.

p0230 Does any of this mean that cognition is not embodied, is not grounded? That mathematical understanding does not involve sensory–motor simulation? No. But it does mean that in many cases much more work needs to be done to establish the claim, whether gathering new and better empirical evidence, or reworking existing evidence to more clearly support a specific position. We must ask: what is it about *this* overlap of neural circuitry that suggests simulation in particular, rather than metaphorical mapping, or something else entirely? Which details in the general finding of cognitive interference indicate more than just a resource bottleneck? Is there some directionality, some selectivity to the interference

that may give us insight into the nature of the functional inheritance that the overlap enables? A few researchers have started to focus on such specifics, as a way to decide *not* between amodal and modal theories of cognition (for which such details are often beside the point), but precisely between competing theories of grounded and embodied cognition (Casasanto & Lozano, 2007). It is time for more of us to follow their lead.

s0030

CONCLUSION

p0240

In this chapter I have laid down a challenge to the field of embodied and grounded cognition. Should the field see fit to pick it up, there is good reason to believe that the results will be very positive. To do that, we must move beyond the too simple task of finding evidence against abstract, amodal, and cognitivist theories of cognition and focus on detailing and supporting specific accounts of the functional inheritances that abstract higher order cognition has received from the substrates on which it is built. This will mean being a bit more critical of each other's work—though hopefully not less friendly toward one another. And indeed, I hope that this chapter is taken in just such a friendly, but constructively critical, spirit.

s0040

ACKNOWLEDGMENTS

p0250

Discussions with Daniel Casasanto helped me to clarify certain aspects of the current state of the field. Thanks to Tony Chemero for her helpful comments on an earlier draft. Thanks are also due to Paco Calvo and Toni Gomila for inviting this contribution, and for their patience while it was prepared.

REFERENCES

Anderson, M. L. (2003). Embodied cognition: A field guide. *Artificial Intelligence, 149*(1), 91–103.

Anderson, M. L. (2007a). Evolution of cognitive function via redeployment of brain areas. *The Neuroscientist, 13*(1), 13–21.

Anderson, M. L. (2007b). Massive redeployment, exaptation, and the functional integration of cognitive operations. *Synthese, 159*(3), 329–345.

Anderson, M. L. (2007c). The massive redeployment hypothesis and the functional topography of the brain. *Philosophical Psychology, 21*(2), 143–174.

Anderson, M. L. (2007d). How to study the mind: An introduction to embodied cognition. In F. Santoianni & C. Sabatano (Eds.), *Brain Development in Learning Environments: Embodied and Perceptual Advancements*. Cambridge Scholars Publishing. Newcastle, UK; 65–82.

Anderson, M. L. (2008). Embodiment, evolution, and the nature of the mind. In B. Hardy-Vallée & N. Payette (Eds.), *Beyond the Brain: Embodied, Situated and Distributed Cognition*. Cambridge Scholars Publishing. Newcastle, UK; 15–28.

Anderson, M. L. & Penner-Wilger, M. (2007). Do redeployed finger representations underlie math ability? In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society* (p. 1703). Austin, TX: Cognitive Science Society.

Anderson, M. L. & Rosenberg, G. (in press). Content and action: The guidance theory of representation. *Journal of Mind and Behavior*.

Anderson, M. L., Brumbaugh, J., & Suben, A. (in press). Investigating functional cooperation in the human brain using simple graph-theoretic methods. In: P. M. Pardalos, V. Boginski, and P. Xanthopoulos, (eds.). Computational Neuroscience. Springer.

Banister, C. M. & Tew, B. (1991). *Current Concepts in Spina Bifida and Hydrocephalus*. Cambridge: Cambridge Press.

Barnes, M. A., Smith-Chant, B. L., & Landry, S. (2005). Number processing in neurodevelopmental disorders: Spina bifida myelomenigocele. In J. I. D. Campbell (Ed.), *Handbook of Mathematical Cognition*. New York, NY: Psychology Press.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences, 22,* 577–660.

Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology, 59,* 617–645.

Barsalou, L. W., Simmons, W. K., Barbey, A. K., & Wilson, C. D. (2003). Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences, 7*(2), 84–91.

Barsalou, L. W., Pecher, D., Zeelenberg, R., Simmons, W. K., & Hamann, S. B. (2005). Multimodal simulation in conceptual processing. In W. Ahn, R. Goldstone, B. Love, A. Markman, & P. Wolff (Eds.), *Categorization Inside and Outside the Lab: Essays in Honor of Douglas L. Medin* (pp. 249–270). Washington, DC: American Psychological Association.

Boroditsky, L. & Ramscar, M. (2002). The roles of body and mind in abstract thought. *Psychological Science, 13*(2), 185–188.

Butterworth, B. (1999). *The Mathematical Brain*. London: Nelson.

Casasanto, D. & Lozano, S. (2007). Meaning and motor action, In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society* (p. 149–154). Austin, TX: Cognitive Science Society.

Cermak, S. A. & Larkin, D. (2001). *Developmental Coordination Disorder*. Albany, NY: Delmar.

Chao, L. L. & Martin, A. (2000). Representation of manipulable man-made objects in the dorsal stream. *Neuroimage, 12,* 478–484.

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

Clark, A. (1998). Embodied, situated, and distributed cognition. In W. Bechtel & G. Graham (Eds.), *A Companion to Cognitive Science*. Malden, MA: Blackwell Publishers. 506–517.

Dagher, A., Owen, A., Boecker, H., & Brooks, D. (1999). Mapping the network for planning. *Brain, 122,* 1973–1987.

Damasio, A. & Tranel, D. (1993). Nouns and verbs are retrieved with differently distributed neural systems. *Proceedings of the National Academy of Science USA, 90,* 4957–4960.

Decety, J. & Grèzes, J. (1999). Neural mechanisms subserving the perception of human actions. *Trends in Cognitive Sciences, 3,* 172–178.

Decety, J., Sjoholm, H., Ryding, E., Stenberg, G., & Ingvar, D. (1990). The cerebellum participates in cognitive activity: Tomographic measurements of regional cerebral blood flow. *Brain Research, 535,* 313–317.

Decety, J., Grezes, J., Costes, N., Perani, D., Jeannerod, M., Procyk, E., Grassi, F., & Fazio, F. (1997). Brain activity during observation of actions. Influence of action content and subject's strategy. *Brain, 120,* 1763–1777.

Dehaene, S. & Cohen, L. (2007). Cultural recycling of cortical maps. *Neuron, 56,* 384–398.

Fauconnier, G. & Turner, M. (2002). *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. New York: Basic Books.

Gallese, V. & Lakoff, G. (2005). The brain's concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology, 22*(3–4), 455–479.

Gibbs, R. W. Jr., (2006). *Embodiment and Cognitive Science*. New York: Cambridge University Press.

Glenberg, A. (1997). What memory is for. *Behavior and Brain Sciences, 20,* 1–55.
Glenberg, A. & Kaschak, M. (2002). Grounding language in action. *Psychonomic Bulletin and Review, 9,* 558–565.
Gorniak, P. & Roy, D. (2006). Perceived affordances as a substrate for linguistic concepts. In R. Sun (Ed.), *Proceedings of the 29th Annual Cognitive Science Society* (p. 279–284). Austin, TX: Cognitive Science Society.
Hamilton, S. S. (2002). Evaluation of clumsiness in children. *American Family Physician, 66,* 1435–1440.
Heineman, G. T. & Councill, W. T. (2001). *Component-Based Software Engineering: Putting the Pieces Together*. New York: Addison-Wesley.
Jeannerod, M. (1994). The representing brain: Neural correlates of motor intention and imagery. *Behavioral Brain Science, 17,* 187–245.
Lakoff, G. & Johnson, M. (1980). *Metaphors We Live by*. Chicago, IL: University of Chicago Press.
Lakoff, G. & Johnson, M. (1999). *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*. New York: Basic Books.
Lakoff, G. & Núñez, R. (2000). *Where Mathematics Comes from: How the Embodied Mind Brings Mathematics into Being*. New York: Basic Books.
Martin, A., Haxby, J. V., Lalonde, F. M., Wiggs, C. L., & Ungerleider, L. G. (1995). Discrete cortical regions associated with knowledge of color and knowledge of action. *Science, 270,* 102–105.
Martin, A., Wiggs, C. L., Ungerleider, L. G., & Haxby, J. V. (1996). Neural correlates of category-specific knowledge. *Nature, 379,* 649–652.
Penner-Wilger, M., Fast, L., LeFevre, J., Smith-Chant, B. L., Skwarchuk, S., Kamawar, D., & Bisanz, J. (2007). The foundations of numeracy: Subitizing, finger gnosia, and fine-motor ability. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society* (p. 1385–1390. Austin, TX: Cognitive Science Society.
Penner-Wilger, M. & Anderson, M. L. (2008). An alternative view of the relation between finger gnosis and math ability: Redeployment of finger representations for the representation of number. In V. Sloutsky, B. Love, & K. McRae (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society.* Austin, TX: Cognitive Science Society.
Prinz, J. (2002). *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge, MA: MIT Press.
Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy, 101,* 389–428.
Rusconi, E., Walsh, V., & Butterworth, B. (2005). Dexterity with numbers: rTMS over left angular gyrus disrupts finger gnosis and number processing. *Neuropsychologica, 43,* 1609–1624.
Stewart, T. & West, R. (2007). Cognitive redeployment in ACT-R: Salience, vision, and memory, *Proceedings of ICCM-2007-Eighth International Conference on Cognitive Modeling*, 313–318. Ann Arbor, MI.
Svensson, H., Lindblom, J., & Ziemke, T. (2007). Making sense of embodied cognition: Simulation theories of shared neural mechanisms for sensorimotor and cognitive processes. In J. Zlatev, T. Ziemke, & R. Frank (Eds.), *Body, Language and Mind* (vol. 1). Berlin: Mouton de Gruyter.
Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science, 12,* 49–100.
Weiskopf, D. (2007). Concept empiricism and the vehicles of thought. *The Journal of Consciousness Studies, 14,* 156–183.
Wilson, M. (2001). The case for sensorimotor coding in working memory. *Psychonomic Bulletin and Review, 8,* 44–57.
Zago, L., Pesenti, M., Mellet, E., Crivello, F., Mazoyer, B., & Tzourio-Mazoyer, N. (2001). Neural correlates of simple and complex mental calculation. *NeuroImage, 13,* 314–327.