Philosophical Psychology

ISSN: 0951-5089 (Print) 1465-394X (Online) Journal homepage: http://www.tandfonline.com/loi/cphp20

Why Heideggerian AI Failed and How Fixing it Would Require Making it More Heideggerian

Hubert L. Dreyfus

To cite this article: Hubert L. Dreyfus (2007) Why Heideggerian AI Failed and How Fixing it Would Require Making it More Heideggerian, Philosophical Psychology, 20:2, 247-268, DOI: 10.1080/09515080701239510

To link to this article: http://dx.doi.org/10.1080/09515080701239510

Published online: 23 Apr 2007.

Philosophical Psychology
Vol. 20, No. 2, April 2007, pp. 247–268

Review Essay

Why Heideggerian AI Failed and How Fixing it Would Require Making it More Heideggerian

Hubert L. Dreyfus

Reconstructing the Cognitive World: The Next Step

MICHAEL WHEELER
Cambridge, MA: MIT Press, 2005
432 pages, ISBN: 0262232405 (hbk); $35.00

1. Symbolic AI as a Degenerating Research Program

When I was teaching at MIT in the 1960s, students from the Artificial Intelligence Laboratory would come to my Heidegger course and say in effect: "You philosophers have been reflecting in your armchairs for over 2000 years and you still don't understand intelligence. We in the AI Lab have taken over and are succeeding where you philosophers have failed." But in 1963, when I was invited to evaluate the work of Alan Newell and Herbert Simon on physical symbol systems, I found to my surprise that, far from replacing philosophy, these pioneering researchers had learned a lot, directly and indirectly, from us philosophers: e.g., Hobbes' claim that reasoning was calculating, Descartes' mental representations, Leibniz's idea of a 'universal characteristic' (a set of primitives in which all knowledge could be expressed), Kant's claim that concepts were rules, Frege's formalization of such rules, and Wittgenstein's postulation of logical atoms in his *Tractatus*. In short, without realizing it, AI researchers were hard at work turning rationalist philosophy into a research program.

But I began to suspect that the insights formulated in existentialist armchairs, especially Heidegger's and Merleau-Ponty's, were bad news for those working in AI laboratories—that, by combining representationalism, conceptualism, formalism,

Correspondence to: Hubert L. Dreyfus, Department of Philosophy, 314 Moses Hall 2390, University of California, Berkeley, CA 94720-2390, USA. Email: dreyfus@berkeley.edu

ISSN 0951-5089 (print)/ISSN 1465-394X (online)/07/020247-22 © 2007 Taylor & Francis
DOI: 10.1080/09515080701239510

and logical atomism into a research program, AI researchers had condemned their enterprise to reenact a failure. Using Heidegger as a guide, I began looking for signs that the whole AI research program was degenerating. I was particularly struck by the fact that, among other troubles, researchers were running up against the problem of representing significance and relevance—a problem that Heidegger saw was implicit in Descartes' understanding of the world as a set of meaningless facts to which the mind assigned values, which John Searle now calls function predicates.

Heidegger warned that values are just more meaningless facts. To say a hammer has the function, hammering, leaves out the defining relation of hammers to nails and other equipment, to the point of building things, to the skill required in actually using a hammer, etc.—all of which Heidegger called "readiness-to-hand"— so attributing functions to brute facts couldn't capture the meaningful organization of the everyday world and so missed the way of being of equipment. "By taking refuge in 'value'-characteristics," Heidegger said, "we are . . . far from even catching a glimpse of being as readiness-to-hand" (Heidegger, 1962, pp. 132–133).

Head of MIT's AI Lab, Marvin Minsky, unaware of Heidegger's critique, was convinced that representing a few million facts about objects including their functions, would solve what had come to be called the commonsense knowledge problem. It seemed to me, however, that the real problem wasn't storing millions of facts; it was knowing which facts were relevant in any given situation. One version of this relevance problem is called the 'frame problem.' If the computer is running a representation of the current state of the world and something in the world changes, how does the program determine which of its represented facts can be assumed to have stayed the same, and which might have to be updated?

As Michael Wheeler puts it in *Reconstructing the Cognitive World*:

> Given a dynamically changing world, how is a nonmagical system . . . to take account of those state changes in that world . . . that matter, and those unchanged states in that world that matter, while ignoring those that do not? And how is that system to retrieve and (if necessary) to revise, out of all the beliefs that it possesses, just those beliefs that are relevant in some particular context of action? (Wheeler, 2005, p. 179)

Minsky suggested as a solution that AI programmers could use descriptions of typical situations like going to a birthday party to list and organize those, and only those, facts that were normally relevant. He suggested a structure of essential features and default assignments—a structure Husserl had already proposed and called a "frame" (Husserl, 1973, p. 38).[1]

But a system of frames isn't *in* a situation, so in order to identify the possibly relevant facts in the current situation one would need a frame for recognizing that situation, etc. It thus seemed to me obvious that any AI program using frames was going to be caught in a regress of frames for recognizing relevant frames for recognizing relevant facts, and that, therefore, the commonsense knowledge storage and retrieval problem wasn't just a *problem*; it was a sign that something was seriously wrong with the whole approach.

Unfortunately, what has always distinguished AI research from a science is its failure to face up to, and learn from, its failures. To avoid the relevance problem, AI programmers in the 1960s and early 1970s limited their programs to what they called 'micro-worlds'—artificial situations in which the small number of features that were possibly relevant was determined beforehand. It was assumed that the techniques used to construct these micro-worlds could be made more realistic and generalized to cover commonsense knowledge—but there were no successful follow-ups, and the frame problem remains unsolved.

John Haugeland argues that symbolic AI has failed and refers to it as "Good Old Fashioned AI" (GOFAI). That name has been widely accepted as capturing symbolic AI's current status. Michael Wheeler goes further, arguing that a new paradigm is already taking shape: "A Heideggerian cognitive science is . . . emerging right now, in the laboratories and offices around the world where embodied-embedded thinking is under active investigation and development" (Wheeler, 2005, p. 285).

Wheeler's well informed book could not have been more timely since there are now at least three versions of supposedly Heideggerian AI that might be thought of as articulating a new paradigm for the field: Rodney Brooks' behaviorist approach at MIT, Phil Agre's pragmatist model, and Walter Freeman's dynamic neural model. All three approaches accept Heidegger's critique of Cartesian internalist representationalism, and, instead, embrace John Haugeland's slogan that cognition is "embedded and embodied" (Haugeland, 1998).

2. Heideggerian AI, Stage One: Eliminating Representations by Building Behavior-Based Robots

Winograd (1989) notes the irony in the MIT AI Lab's becoming a cradle of "Heideggerian AI" after its initial hostility to my presentation of these ideas (as cited in Dreyfus, 1992, p. xxxi). Here's how it happened. In March 1986, the MIT AI Lab under its new director, Patrick Winston, reversed Minsky's attitude toward me and allowed, if not encouraged, several graduate students to invite me to give a talk I called "Why AI Researchers should study *Being and Time*." There I repeated the Heideggerian message of my *What Computers Can't Do*: "The meaningful objects . . . among which we live are not a *model* of the world stored in our mind or brain; *they are the world itself*" (Dreyfus, 1972, pp. 265–266).

The year of my talk, Rodney Brooks published a paper criticizing the GOFAI robots that used representations of the world and problem solving techniques to plan their movements. He reported that, based on the idea that "the best model of the world is the world itself," he had "developed a different approach in which a mobile robot uses the world itself as is own representation—continually referring to its sensors rather than to an internal world model" (Brooks, 1997b, p. 416). Looking back at the frame problem, he says: "And why could my simulated robot handle it? Because it was using the world as its own model. It never referred to an internal description of the world that would quickly get out of date if anything in the real world moved" (Brooks, 2002, p. 42). Although he doesn't acknowledge the influence of Heidegger directly (and even denies it in 1997b, p. 415), Brooks gives me credit for "being right about many issues such as the way in which people operate in the world is intimately coupled to the existence of their body" (Brooks, 2002, p. 168).

Brooks' approach is an important advance, but his robots respond only to *fixed features* of the environment, not to context or changing significance. They are like ants, and Brooks aptly calls them "animats." Brooks thinks he does not need to worry about learning, putting it off as a concern of possible future research. But by operating in a fixed world and responding only to the small set of possibly relevant features that their receptors can pick up, Brooks' animats beg the question of changing relevance and so finesse rather than solve the frame problem.

Merleau-Ponty's work, on the contrary, offers a nonrepresentational account of the way the body and the world are coupled that suggests a way of avoiding the frame problem. What the learner acquires through experience is not *represented* at all but is *presented* to the learner as more and more finely discriminated situations, and, if the situation does not clearly solicit a single response or if the response does not produce a satisfactory result, the learner is led to further refine his discriminations, which, in turn, solicit more refined responses. For example, what we have learned from our experience of finding our way around in a city is sedimented in how that city *looks* to us. Merleau-Ponty calls this feedback loop between the embodied agent and the perceptual world the "intentional arc" (Merleau-Ponty, 1962, p. 136).

Brooks comes close to a basic insight spelled out by Merleau-Ponty (1966), namely that intelligence is founded on and presupposes the more basic way of coping we share with animals: "The 'simple' things concerning perception and mobility in a dynamic environment . . . are a necessary basis for 'higher-level' intellect. . . . Therefore, I proposed looking at simpler animals as a bottom-up model for building intelligence" (Brooks, 1997b, p. 418). Surprisingly, the modesty Brooks exhibited in choosing to first construct simple insect-like devices did not deter Brooks and Daniel Dennett from deciding to leap ahead and "[embark] on a long-term project to design and build a humanoid robot, Cog, whose cognitive talents will include speech, eye-coordinated manipulation of objects, and a host of self-protective, self-regulatory and self-exploring activities" (Dennett, 1994, p. 133).

Of course, the "long term project" was short lived. Cog failed to achieve any of its goals and is already in a museum.[2] But, as far as I know, neither Dennett nor anyone connected with the project has published an account of the failure and asked what mistaken assumptions underlay their absurd optimism. In a personal communication, Dennett blamed the failure on a lack of graduate students and claimed "Progress was being made on all the goals, but slower than had been anticipated." If progress was actually being made the graduate students wouldn't have left, or others would have continued to work on the project. Clearly some specific assumptions must have been mistaken, but all we find in Dennett's assessment is the implicit assumption that human intelligence is on a continuum with insect intelligence, and that therefore adding a bit of complexity to what has already been done with animats counts as progress toward humanoid intelligence. At the beginning of AI research, Yehoshua Bar-Hillel called this way of thinking the first-step fallacy, and my brother at RAND quipped, "it's like claiming that the first monkey that climbed a tree was making progress towards flight to the moon."

In contrast to Dennett's assessment, Brooks is prepared to entertain the possibility that he is barking up the wrong tree, making the sober comment that:

> Perhaps there is a way of looking at biological systems that will illuminate an inherent necessity in some aspect of the interactions of their parts that is completely missing from our artificial systems . . . perhaps at this point we simply do not *get it*, and that there is some fundamental change necessary in our thinking in order that we might build artificial systems that have the levels of intelligence, emotional interactions, long term stability and autonomy, and general robustness that we might expect of biological systems. (Brooks, 1997a, p. 301)

Heidegger and Merleau-Ponty would say that, in spite of the breakthrough of giving up internal symbolic representations, Brooks, indeed, doesn't get it—that what AI researchers have to face and understand is not only why our everyday coping couldn't be understood in terms of inferences from symbolic representations, as Minsky's intellectualist approach assumed, but also why it can't be understood in terms of responses caused by fixed features of the environment, as in Brooks' empiricist approach. AI researchers need to consider the possibility that embodied beings like us take as input energy from the physical universe and respond in such a way as to open them to a world organized in terms of their needs, interests, and bodily capacities, without their *minds* needing to impose meaning on a meaningless given, as Minsky's frames require, nor their *brains* converting stimulus input into reflex responses, as in Brooks' animats.

Later I'll suggest that Walter Freeman's neurodynamics offers a radically new basis for a Heideggerian/Merleau-Pontian approach to human intelligence—an approach compatible with physics and grounded in the neuroscience of perception and action. But first we need to examine another approach to AI contemporaneous with Brooks' that actually calls itself Heideggerian.

3. Heideggerian AI, Stage 2: Programming the Ready-To-Hand

In my talk at the MIT AI Lab, I introduced Heidegger's nonrepresentational account of the relation of Dasein (human being) and the world. I also explained that Heidegger distinguished two modes of being: the *readiness-to-hand* of equipment when we are involved in using it, and the *presence-at-hand* of objects when we contemplate them. Out of that explanation and the lively discussion that followed, grew the second type of Heideggerian AI—the first to acknowledge its lineage.

This new approach took the form of Phil Agre's and David Chapman's program, *Pengi*, which guided a virtual agent playing a computer game called *Pengo*, in which the player and penguins kick large and deadly blocks of ice at each other (Agre, 1988, Ch. A1, part A1a, p. 9). Agre's approach, called "interactionism," was more self-consciously Heideggerian than Brooks' in that Agre tried explicitly to capture "Heidegger's phenomenological analysis of routine activity" (Agre, 1997, p. 5).

Agre's interesting new idea is that the world of the game in which the Pengi agent acts is made up, not of present-at-hand objects with properties, but of possibilities for action that trigger appropriate responses from the agent. To program this situated approach Agre used what he called "deictic representations." He tells us: "This proposal is based on a rough analogy with Heidegger's analysis of everyday intentionality in Division I of *Being and Time*, with objective intentionality corresponding to the present-at-hand and deictic intentionality corresponding to the ready-to-hand" (Agre, 1997, p. 332). And he explains: "[Deictic representations] designate, not a particular object in the world, but rather a role that an object might play in a certain time-extended pattern of interaction between an agent and its environment. Different objects might occupy this role at different times, but the agent will treat all of them in the same way" (Agre, 1997, p. 251).

Looking back on my talk at MIT and rereading Agre's book, I now see that, in a way, Agre understood Heidegger's account of readiness-to-hand better than I did at the time. I thought of the ready-to-hand as a special class of *entities*, viz. equipment, whereas the Pengi program treats what the agent responds to purely as *functions*. For Heidegger and Agre the ready-to-hand is not a *what* but a *for-what*.[3]

As Agre saw, Heidegger wants to get at something more basic than simply a class of objects defined by their use. At his best, Heidegger would, I think, deny that a hammer in a drawer has readiness-to-hand as its way of being. Rather, he sees that, *for the user*, equipment is a solicitation to act, not an entity with a function feature. Heidegger also notes that to observe a hammer or to observe ourselves hammering undermines our skillful coping. We can and do observe our surroundings while we cope, and sometimes, if we are learning, monitoring our performance as we learn improves our performance in the long run, but, in the short run, such attention interferes with our performance.

Heidegger struggles to describe the special, and he claims, basic, way of being he calls the ready-to-hand. The Gestaltists would later talk of "solicitations." In *Phenomenology of Perception*, Merleau-Ponty speaks of "motivations," and later, of "the flesh." All these terms point at what is not objectifiable—a situation's way of drawing one into it. Indeed, in his 1925 course, *Logic: The Question of Truth*, Heidegger describes our most basic experience of what he later calls "pressing into possibilities" not as dealing with the desk, the door, the lamp, the chair and so forth, but as directly responding to a "what for":

> What is first of all "given" . . . is the "for writing," the "for going in and out," the "for illuminating," the "for sitting." That is, writing, going-in-and-out, sitting, and the like are what we are a priori involved with. What we know when we "know our way around" and what we learn are these "for-what's." (1976, p. 144; for the translation, see Heidegger, in press.)

It's clear here, unlike what some people take Heidegger to suggest in *Being and Time*, that this basic experience has no as-structure. That is, when absorbed in coping, I can be described *objectively* as using the door *as* a door, but I'm not *experiencing* the door *as* a door. In coping at my best, I'm not experiencing *the door* at all but simply pressing into the possibility of going out. The important thing to realize is that, when we are pressing into possibilities, there is no experience of an *entity* doing the soliciting; just the solicitation. Such solicitations disclose the world on the basis of which we sometimes do step back and perceive things *as* things.

But Agre's Heideggerian AI did not try to program this experiential aspect of being drawn in by an affordance. Rather, with his deictic representations, Agre *objectified* both the functions and their situational relevance for the agent. In *Pengi*, when a virtual ice cube defined by its function is close to the virtual player, a rule dictates the response (e.g., kick it). No skill is involved and no learning takes place.

So Agre had something right that I was missing—the transparency of the ready-to-hand—but he nonetheless fell short of being fully Heideggerian. For Heidegger the ready-to-hand is not a fixed function, encountered in a predefined type of situation that triggers a predetermined response that either succeeds or fails. Rather, as we have begun to see and will soon see further, readiness-to-hand is experienced as a solicitation that calls forth a flexible response to the *significance* of the current situation—a response which is experienced as either improving the situation or making it worse.

Moreover, although he proposed to program Heidegger's account of everyday routine activities, Agre doesn't even try to account for how our experience feeds back and changes our sense of the significance of the next situation and what is relevant in *it*. In putting his virtual agent in a virtual world where all possibly relevance is determined beforehand, Agre doesn't account for how we learn to respond to new relevancies, and so, like Brooks, he finesses rather than solves the frame problem. Thus, sadly, his Heideggerian AI turned out to be a dead end. Happily, however, Agre never claimed he was making progress towards building a human being.

4. Pseudo Heideggerian AI: Situated Cognition and the Embedded, Embodied, Extended Mind

Wheeler (2005) praises me for putting the confrontation between Cartesian and Heideggerian ontologies to the test in the empirical realm. Wheeler claims, however, that I only made negative predictions about the viability of GOFAI and cognitive science research programs. The time has come, he says, for a positive Heideggerian approach and that the emerging embodied-embedded paradigm in the field is a thoroughly Heideggerian one.

As if taking up from where Agre left off with his objectified version of the ready-to-hand, Wheeler tells us: "Our global project requires a defense of action-oriented representation. ... action-oriented representation may be interpreted as the subagential reflection of online practical problem solving, as conceived by the Heideggerian phenomenologist. Embodied-embedded cognitive science is implicitly a Heideggerian venture" (Wheeler, 2005, pp. 222–223). He further notes: "As part of its promise, this nascent, Heideggerian paradigm would need to indicate that it might plausibly be able either to solve or to dissolve the frame problem" (p. 187). And he suggests: "The good news for the reoriented Heideggerian is that the kind of evidence called for here may already exist in the work of recent *embodied-embedded cognitive science*" (Wheeler, 2005, p. 188). He concludes:

> Let's be clear about the general relationships at work here. Dreyfus is right that the philosophical impasse between a Cartesian and a Heideggerian metaphysics can be resolved empirically via cognitive science. However, he looks for resolution in the wrong place. For it is not any alleged empirical failure on the part of orthodox cognitive science, but rather the concrete empirical success of a cognitive science with Heideggerian credentials, that, if sustained and deepened, would ultimately vindicate a Heideggerian position in cognitive theory. (Wheeler, 2005, pp. 188–189)

I agree it is time for a positive account of how Heideggerian AI and an underlying Heideggerian neuroscience could solve the frame problem, but I think Wheeler is the one looking in the wrong place. Merely in supposing that Heidegger is concerned with subagential *problem solving* and action oriented *representations*, Wheeler's project reflects not a step beyond Agre but a regression to pre-Brooks GOFAI. Heidegger, indeed, claims that that skillful coping is basic, but he is also clear that, all coping takes place on the background coping he calls "being-in-the world" which doesn't involve any form of representation at all.[4]

Wheeler's cognitivist misreading of Heidegger leads to an overestimation of the importance of Clark and Chalmers' (1998) attempt to free us from the Cartesian idea that the mind is essentially inner by pointing out that in thinking we sometimes make use of external artifacts like pencil, paper, and computers. Unfortunately, this argument for the extended mind preserves the Cartesian assumption that our basic way of relating to the world is by using representations such as beliefs and memories, be they in the mind or in notebooks in the world. In effect, while Brooks and Agre dispense with representations where coping is concerned, all Clark, Chalmers, and Wheeler give us as a supposedly radical new Heideggerian approach to the human way of being in the world is the observation that memories and beliefs are not necessarily *inner* entities and that, therefore, *thinking bridges the distinction between inner and outer representations*.[5]

Heidegger's important insight is not that, when we solve problems, we sometimes make use of representational equipment outside our bodies, but that *being-in-the-world* is more basic than *thinking* and solving problems; it is not representational at all. That is, when we are coping at our best, we are drawn in by solicitations and respond directly to them, so that the distinction between us and our equipment vanishes. As Heidegger sums it up: "I *live* in the understanding of writing, illuminating, going-in-and-out, and the like. More precisely: as Dasein I am—in speaking, going, and understanding—an act of understanding dealing-with. My being in the world *is* nothing other than this already-operating-with-understanding in this mode of being" (Heidegger, 1976, p. 146).[6]

Heidegger's and Merleau-Ponty's understanding of embedded-embodied coping, therefore, is not that the *mind* is sometimes *extended into the world* but rather that, in our most basic way of being—i.e., as skillful copers—we are not minds at all but *one with the world*. Heidegger sticks to the phenomenon, when he says that, in its most basic way of being, "Dasein is its world existingly" (Heidegger, 1962, p. 416). (To make sense of this slogan, it's important to be clear that Heidegger distinguishes the human *world* from the physical *universe*.)

When you stop thinking that mind is what characterizes us most basically but, rather, that most basically we are absorbed copers, the inner–outer distinction becomes problematic. There's no easily askable question about where the absorbed coping is—in me or in the world. Thus, for a Heideggerian all forms of *cognitivist* externalism presuppose a more basic *existentialist* externalism where even to speak of "externalism" is misleading since such talk presupposes a contrast with the internal. Compared to this genuinely Heideggerian view, extended-mind externalism is contrived, trivial, and irrelevant.

5. What Motivates Embedded-Embodied Coping?

But why is Dasein called to cope at all? According to Heidegger, we are constantly solicited to improve our familiarity with the world: "Caring takes the form of a looking around and seeing, and as this circumspective caring it is at the same time anxiously concerned about developing its circumspection, that is, about *securing and expanding* its *familiarity* with the objects of its dealings" (2002, p. 115).[7] This pragmatic perspective is developed by Merleau-Ponty, and by Samuel Todes (2001). These heirs to Heidegger's account of familiarity and coping describe how an organism, animal or human, interacts with what is objectively speaking the meaningless physical universe in such a way as to experience it as an environment organized in terms of what that organism needs in order to find its way around. All such coping beings are motivated to get a more and more refined and secure sense of their environment and of the specific objects of their dealings. According to Merleau-Ponty (1962): "My body is geared into the world when my perception presents me with a spectacle as varied and as clearly articulated as possible. . . . (p. 250, translation modified).

In short, in our skilled activity we are drawn to move so as to achieve a better and better grip on our situation. For this movement towards maximal grip to take place, one doesn't need a mental representation of one's goal nor any subagential problem solving as would a GOFAI robot. Rather, acting is experienced as a steady flow of skillful activity in response to one's sense of the situation. When one's situation deviates from some optimal body-environment gestalt, one's activity takes one closer to that optimum and thereby relieves the "tension" of the deviation. One does not need to know what that optimum is in order to move towards it. One's body is simply solicited by the situation to lower the tension. As Merleau-Ponty puts it: "Our body is not an object for an 'I think', it is a grouping of lived-through meanings that moves towards its equilibrium" (Merleau-Ponty, 1962, p. 153).

6. Modeling Situated Coping as a Dynamical System

Describing the phenomenon of everyday coping as being "geared into" the world and moving towards "equilibrium" suggests a *dynamic* relation between the coper and the environment. Timothy van Gelder calls this dynamic relation "coupling," explaining its importance as follows:

> The post-Cartesian agent manages to cope with the world without necessarily representing it. A dynamical approach suggests how this might be possible by showing how the internal operation of a system interacting with an external world can be so subtle and complex as to *defy* description in representational terms—how, in other words, cognition can *transcend* representation. (van Gelder, 1997, p. 448)

Van Gelder shares with Brooks the idea that thought is grounded in a more basic relation of agent and world. As van Gelder puts it: "Cognition can, in sophisticated cases, involve representation and sequential processing; but such phenomena are best understood as emerging from a dynamical substrate, rather than as constituting the basic level of cognitive performance" (van Gelder, 1997, p. 439). This dynamical substrate is precisely the subagential causal basis of the skillful coping first described by Heidegger and worked out in detail by Merleau-Ponty.

Van Gelder importantly contrasts the rich interactive temporality of real-time online coupling of coper and world with the austere step by step temporality of thought. Wheeler helpfully explains:

> Whilst the computational architectures proposed within computational cognitive science require that inner events happen in the right order, and (in theory) fast enough to get a job done, there are, in general, no constraints on how long each operation within the overall cognitive process takes, or on how long the gaps between the individual operations are. Moreover, the transition events that characterize those inner operations are not related in any systematic way to the real-time dynamics of either neural biochemical processes, non-neural bodily events, or environmental phenomena (dynamics which surely involve rates and rhythms). (Wheeler, 2002, p. 345)

Computation is thus paradigmatically austere. Wheeler adds: "Turing machine computing is digital, deterministic, discrete, effective (in the technical sense that behavior is always the result of an algorithmically specified finite number of operations), and temporally austere (in that time is reduced to mere sequence)" (Wheeler, 2002, pp. 344–345).

Ironically, Wheeler's highlighting the contrast between rich and austere temporality enables us to see clearly that his appeal to extended minds as a Heideggerian response to Cartesianism leaves out the essential embodied embedding. Clark and Chalmers' examples of extended minds dealing with representations are clearly a case of computational austerity—no rates and rhythms are involved. Wheeler is aware of this possible objection to his backing of both the *dynamical systems* model and the *extended mind* approach. He asks: "What about the apparent clash between continuous reciprocal causation and action orientated representations?

On the face of it this clash is a worry for our emerging cognitive science" (Wheeler, 2005, p. 280). But instead of facing up to the incompatibility of these two opposed models of ground level intelligence, Wheeler suggests that we must somehow combine them and that "this is the biggest of the many challenges that lie ahead" (p. 280).

Wheeler's ambivalence as to which model is more basic, the representational or the dynamic, undermines his Heideggerian approach. For, as Wheeler himself sees, the Heideggerian claim is that action-oriented coping, as long as it is involved (online, Wheeler would say), is not representational at all and does not involve any problem solving, and that all representational problem solving takes place later offline.[8] Showing in detail how the representational un-ready-to-hand in all its forms is derivate from the nonrepresentational ready-to-hand is exactly the Heideggerian project. It requires a basic choice of ontology, phenomenology, and brain model, between a cognitivist model that gives a basic role to representations, and a dynamical model like Merleau-Ponty's and van Gelder's that denies a basic role to any sort of representation—even action oriented ones—and gives a primordial place to equilibrium and in general to rich coupling.

Ultimately, we have to choose which sort of AI and which sort of neuroscience to back, and so we are led to the questions: could the brain in its causal support of our active coping instantiate a richly coupled dynamical system, and is there any evidence it actually does so? If so, could this coupling be modeled on a digital computer to give us Heideggerian AI?

7. Walter Freeman's Heideggerian/Merleau-Pontian Neurodynamics

We have seen that our experience of the everyday world is organized in terms of significance and relevance and that this significance can't be constructed by giving meaning to brute facts—both because we don't experience brute facts and, even if we did, no value predicate could do the job of giving them situational significance. Yet, all that the organism can receive as input is mere physical energy. How can such senseless physical stimulation be experienced directly as significant? If we can't answer this question, the phenomenological observation that the world is its own best representation, and that the significance we find in our world is constantly enriched by our experience in it, seems to require that the brain be what Dennett derisively calls "wonder tissue."

Fortunately, there is at least one model of how the brain could provide the causal basis for the intentional arc. Walter Freeman, a founding figure in neuroscience and the first to take seriously the idea of the brain as a nonlinear dynamical system, has worked out an account of how the brain of an active animal can find and augment significance in its world. On the basis of years of work on olfaction, vision, touch, and hearing in alert and moving rabbits, Freeman proposes a model of rabbit learning based on the coupling of the brain and the environment. To bring out the relevance of Freeman's account to our phenomenological investigation, I propose to map Freeman's neurodynamic model onto the phenomena we have already noted in the work of Merleau-Ponty.

7.1. Involved Action–Perception (Merleau-Ponty's Being-Absorbed-In-The-World [Être Au Monde]—His Version of Heidegger's In-Der-Welt-Sein)

The organism normally actively seeks to improve its current situation. Thus, according to Freeman's model, when hungry, frightened, disoriented, etc., the rabbit sniffs around until it falls upon food, a hiding place, or whatever else it senses it needs. The animal's neural connections are then strengthened to the extent that reflects the extent to which the result satisfied the animal's current need. In Freeman's neurodynamic model, the input to the rabbit's olfactory bulb modifies the bulb's neuron connections according to the Hebbian rule that neurons that fire together wire together.

7.2. Holism

The change is much more radical than adding a new mechanical response. The next time the rabbit is in a similar state of seeking and encounters a similar smell, the entire bulb goes into a state of global chaotic activity. Freeman tells us:

> Experiments show clearly that every neuron in the [olfactory] bulb participates in generating each olfactory perception. In other words, the salient information about the stimulus is carried in some distinctive pattern of *bulb wide activity*, not in a small subset of feature-detecting neurons that are excited only by, say, foxlike scents. (Freeman, 1991, p. 79, italics added)

7.3. Direct Perception of Significance

After each sniff, the rabbit's bulb exhibits a distribution of energy states. The bulb then tends toward minimum energy the way a ball tends to roll towards the bottom of a container, no matter where it starts from within the container. Each possible minimal energy state is called an "attractor." The brain states that tend towards a particular attractor are called its "basin of attraction."[9]

The rabbit's brain forms a new basin of attraction for each new significant input. Thus, the significance of past experience is preserved in the set of basins of attraction. The set of basins of attraction that an animal has learned forms what is called an "attractor landscape," and "The state space of the cortex can therefore be said to comprise an attractor landscape with several adjoining basins of attraction, one for each class of learned stimuli" (Freeman, 2000, p. 62).[10]

Freeman argues that each new attractor does not *represent*, say, a carrot, or the smell of carrot, or even what to do with a carrot. Rather, the brain's current state is the result of the sum of the animal's past experiences with carrots, and this state is directly coupled with or resonates to the affordance offered by the current carrot. What in the physical input is directly picked up and resonated to when the rabbit sniffs, then, is the affords-eating. Freeman tells us "The macroscopic bulbar patterns [do] not relate to the stimulus directly but instead to the significance of the stimulus" (Freeman, 1995, p. 59).

7.4. The Stimulus is Not Further Processed or Acted Upon (Merleau-Ponty: We Normally Have No Experience of Sense Data)

Since on Freeman's account the bulb responds directly to the contextual significance of the current input, after activating a specific attractor landscape the stimulus has no further job to perform. So the stimulus need not be processed into a representation of the current situation, on the basis of which the brain then has to infer what to do. As Freeman explains:

> The new pattern is selected, not imposed, by the stimulus. It is determined by prior experience with this class of stimulus. The pattern expresses the nature of the class and its significance for the subject rather than the particular event. The identities of the particular neurons in the receptor class that are activated are irrelevant and are not retained. . . . Having played its role in setting the initial conditions, the sense-dependent activity is washed away. (Freeman, 1995, pp. 66–67)

7.5. The Perception–Action Loop

The brain's movement towards the bottom of a particular perceptual basin of attraction underlies the perceiver's perception of the significance of a particular experience. For example, if a carrot affords eating the rabbit is directly readied to eat the carrot, or perhaps readied to carry off the carrot depending on which attractor is currently activated. Freeman (2000) tells us: "The same global states that embody the significance provide . . . the patterns that make choices between available options and that guide the motor systems into sequential movements of intentional behavior" (Freeman, 2000, p. 114).

The animal must take account of how things are going and either continue on a promising path, or, if the overall action is not going as well as anticipated, the brain self-organizes so the attractor system jumps to another attractor. This either causes the animal to act in such a way as to increase its sense of impending reward, or the brain will shift attractors again, until it lands in one that improves its sense of impending reward.[11] The attractors can change like switching from frame to frame in a movie film with each further sniff or with each shift of attention. If the rabbit achieves what it is seeking, a report of its success is fed back to reset the sensitivity of the olfactory bulb. And the cycle is repeated.

7.6. Optimal Grip

The animal's movements are presumably experienced by the animal as its being drawn toward getting an optimal perceptual take on what is currently significant, and, where appropriate, an actual optimal bodily grip on it. Minimum tension is correlated with achieving an optimal grip.[12]

Freeman sees his account of the brain dynamics underlying perception and action as structurally isomorphic with Merleau-Ponty's:

> Merleau-Ponty concludes that we are moved to action by disequilibrium between the self and the world. In dynamic terms, the disequilibrium . . . puts the brain

> onto . . . a pathway through a chain of preferred states, which are learned basins of attraction. The penultimate result is not an equilibrium in the chemical sense, which is a dead state, but a descent for a time into the basin of an attractor, giving an awareness of closure. (Freeman, 2000, p. 121)

Thus, according to Freeman, in governing action the brain normally moves from one basin of attraction to another, descending into each basin for a time without coming to rest in any of them. If so, Merleau-Ponty's talk of *reaching* equilibrium or maximal grip is misleading. But Merleau-Pontians should be happy to improve their phenomenological description on the basis of Freeman's model. Normally, the coper moves *towards* a maximal grip but, instead of coming to rest when the maximal grip is achieved, the coupled coper, without coming to rest, is drawn to move on in response to the call of another affordance that solicits her to take up the same task from another angle, or to turn to the next task that grows out of the current one.

7.7. Experience Feeds Back Into the Look of the World (Merleau-Ponty's Intentional Arc)

Freeman (2000) claims his read out from the rabbit's brain shows that each learning experience that is significant in a new way sets up a new attractor and rearranges all the other attractor basins in the landscape:

> I have observed that brain activity patterns are constantly dissolving, reforming and changing, particularly in relation to one another. When an animal learns to respond to a new odor, there is a shift in all other patterns, even if they are not directly involved with the learning. There are no fixed representations, as there are in [GOFAI] computers; there are only significances. (Freeman, 2000, p. 22)

Freeman (1995) adds:

> I conclude that context dependence is an essential property of the cerebral memory system, in which each new experience must change all of the existing store by some small amount, in order that a new entry be incorporated and fully deployed in the existing body of experience. This property contrasts with memory stores in computers . . . in which each item is positioned by an address or a branch of a search tree. There, each item has a compartment, and new items don't change the old ones. Our data indicate that in brains the store has no boundaries or compartments . . . . Each new state transition . . . initiates the construction of a local pattern that impinges on and modifies the whole intentional structure. (Freeman, 1995, p. 99)

The whole constantly updated pattern of attractors is correlated with the agent's experience of the changing significance of things in the world. Merleau-Ponty likewise concludes that, thanks to the intentional arc, no two experiences of the world are ever exactly alike.

The important point is that Freeman offers a model of learning which is not an associationist model according to which, as one learns, one adds more and more fixed connections, nor a cognitivist model based on offline representations of objective facts about the world that enable offline inferences as to which facts to expect next, and what they mean. Rather, Freeman's model instantiates a genuine intentional arc according to which there are no linear casual connections nor a fixed library of data, but where, each time a new significance is encountered, the whole perceptual world of the animal changes so that significance as directly displayed is contextual, global, and continually enriched.

7.8. Circular Causality

Such systems are self-organizing. Freeman explains:

> Macroscopic ensembles exist in many materials, at many scales in space and time, ranging from . . . weather systems such as hurricanes and tornadoes, even to galaxies. In each case, the behavior of the microscopic elements or particles is constrained by the embedding ensemble, and microscopic behavior cannot be understood except with reference to the macroscopic patterns of activity. (Freeman, 2000, p. 52)

Thus, the cortical field controls the neurons that create the field. In Freeman's terms, in this sort of circular causality the overall activity ''enslaves'' the elements. As he emphasizes:

> Having attained through dendritic and axonal growth a certain density of anatomical connections, the neurons cease to act individually and start participating as part of a group, to which each contributes and from which each accepts direction . . . . The activity level is now determined by the population, not by the individuals. This is the first building block of neurodynamics. (Freeman, 2000, p. 53)

Given the way the whole brain can be tuned by past experience to influence individual neuron activity, Freeman can claim: ''Measurements of the electrical activity of brains show that dynamical states of Neuroactivity emerge like vortices in a weather system, triggered by physical energies impinging onto sensory receptors. . . . These dynamical states determine the structures of intentional actions'' (Freeman, 1995, p. 111). Merleau-Ponty seems to anticipate Freeman's neurodynamics when he says: ''It is necessary only to accept the fact that the physico-chemical actions of which the organism is in a certain manner composed, instead of unfolding in parallel and independent sequences, are constituted . . . in relatively stable 'vortices''' (Merleau-Ponty, 1966, p. 153).

In its dynamic coupling with the environment the brain tends towards equilibrium but continually (discontinuously) switching from one attractor basin to another. The discreteness of these global state transitions from one attractor basin to another makes it possible to model the brain's activity on a computer. Freeman notes that: ''At macroscopic levels each perceptual pattern of Neuroactivity is discrete, because it is marked by state transitions when it is formed and ended . . . . I conclude that brains don't use numbers as symbols, but they do use discrete events in time and space, so we can represent them . . . by numbers in order to model brain states with digital computers'' (Freeman, 1995, p. 105). That is, the computer can model the input and the series of discrete transitions from basin to basin they trigger in the brain, thereby modeling how, on the basis of past experiences of success or failure, physical input acquires significance for the organism. When one actually programs such a model of the brain as a dynamic physical system, one has an explanation of how the brain does what Merleau-Ponty thinks the brain must be doing, and, since Merleau-Ponty is working out of Heidegger's ontology, one has developed Freeman's neurodynamics into Heideggerian AI.

Freeman has actually programmed his model of the brain as a dynamic physical system, and so claims to have shown what the brain is doing to provide the material substrate for Heidegger's and Merleau-Ponty's phenomenological account of everyday perception and action. This may well be the new paradigm for the cognitive sciences that Wheeler proposes to present in his book but which he fails to find. It would show how the emerging embodied-embedded approach could be a step towards a genuinely existential AI. Although, as we shall see, it would still be a very long way from programming human intelligence. Meanwhile, the job of phenomenologists is to get clear concerning the phenomena that must to be explained. That would include an account of how human beings, unlike the so-called Heideggerian computer models we have discussed, don't just ignore the frame problem nor solve it, but show why it doesn't occur.

Time will tell whether Freeman's Merleau-Pontian model is on the right track for explaining how the brain finds and feeds back significance into the meaningless physical universe. Only then will we find out if one can actually produce intelligent behavior by programming a model of the physical state transitions taking place in the brain. That would be the positive Heideggerian contribution to the cognitive sciences that Wheeler proposes to present but fails to find. It would show how the emerging embodied-embedded approach, when fully understood, could, indeed, be a step towards a genuinely Heideggerian AI. Meanwhile, the job of phenomenologists is to get clear concerning the phenomena that need to be explained. That includes an account of how *we*, unlike classical representational computer models, avoid the frame problem.

8. How Would Heideggerian AI Dissolve the Frame Problem?

As we have seen, Wheeler rightly thinks that the simplest test of the viability of any proposed AI program is whether it can solve the frame problem. We've also seen that the two current supposedly Heideggerian approaches to AI avoid the frame problem. Brooks' empiricist/behaviorist approach in which the environment directly causes responses avoids it by leaving out significance and learning altogether, while Agre's action-oriented approach, which includes only a small fixed set of possibly relevant responses, fails to face the problem of changing relevance.

Wheeler's approach, however, by introducing flexible action-oriented *representations*, like any representational approach, has to face the frame problem head on. To see why, we need only slightly revise his statement of the problem (quoted earlier), substituting 'representation' for 'belief': "Given a dynamically changing world, how is a nonmagical system . . . to retrieve and (if necessary) to revise, out of all the *representations* that it possesses, just those *representations* that are relevant in some particular context of action?'' (Wheeler, 2005, p. 179). Wheeler's frame problem, then, is to explain how his allegedly Heideggerian system can determine in some systematic way which of the action-oriented representations it contains or can generate are relevant in any current situation, and to keep track of how this relevance changes with changes in the situation. Not surprisingly, in the concluding chapter of his book where Wheeler returns to the frame problem to test his proposed Heideggerian AI, he offers no solution or dissolution of the problem. Rather he asks us to ''give some credence to [his] informed intuitions'' (Wheeler, 2005, p. 279).

I agree with Wheeler's general intuition, which I take to be on the scent of Freeman's account of rabbit olfaction, viz., that nonrepresentational causal coupling must play a crucial role. But I take issue with his conclusion that:

> In extreme cases the neural contribution will be nonrepresentational in character. In other cases, representations will be active partners alongside certain additional factors, but those representations will be action oriented in character, and so will realize the same content-sparse, action-specific, egocentric, context-dependent profile that Heideggerian phenomenology reveals to be distinctive of online representational states at the agential level. (Wheeler, 2005, p. 276)

All representational states are part of the problem. Therefore, Wheeler as I understand him, cannot give an explanation of how online dynamic coupling will dissolve the online frame problem. Nor does it help to bring in, as Wheeler does, action-oriented representations and the extended mind. Any attempt to solve the frame problem by giving any role to any sort of representational states, even online ones, has so far proved to be a dead end. It looks like nonrepresentational neural activity can't be understood to be the ''extreme case'' as Wheeler claims it is. Rather, such activity must be, as Heidegger, Merleau-Ponty and Freeman claim, our basic way of responding directly to relevance in the everyday world so that the frame problem does not arise.

Heidegger and Merleau-Ponty argue that, thanks to our embodied coping and the intentional arc it makes possible, our skill in directly sensing and responding to relevant changes in the world is constantly improved. In coping in a particular context, say a classroom, we learn to ignore most of what is in the room, but, if it gets too warm, the windows solicit us to open them. We ignore the chalk dust in the corners and chalk marks on the desks but we attend to the chalk marks on the blackboard. We take for granted that what we write on the board doesn't affect the windows, even if we write 'open windows', and what we do with the windows doesn't affect what's on the board. And as we constantly refine this background know-how, the things in the room and its layout take on more and more significance. In general, given our experience in the world, whenever there is a change in the current context we respond to it only if in the past it has turned out to be significant, and when we sense a significant change we treat everything else as unchanged except what our familiarity with the world suggests might also have changed and so needs to be checked out. Thus a local version of the frame problem does not arise.

But the frame problem reasserts itself when we need to change contexts. How do we sense when a situation on the margin of our current activity has become relevant to our current tasks? Merleau-Ponty has a suggestion. When speaking of one's attention being drawn by an affordance on the margin of one's current experience, Merleau-Ponty uses the term 'summons' to describe the influence of the affordance on the perceiver: "To see an object is either to have it on the fringe of the visual field and be able to concentrate on it, or else respond to this *summons* by actually concentrating on it" (Merleau-Ponty, 1962, p. 67, italics added). Thus, for example, as one faces the front of a house, one's body is already being *summoned* (not just *prepared*) to go around the house to get a better look at its back (Kelly, 2005).

Merleau-Ponty's treatment of what Husserl calls the *inner* horizon of the perceptual object, e.g., its insides and back, applies equally to our experience of a situation's *outer* horizon of other potential situations. As I cope with a familiar task in a specific situation, other situations that have in the past been relevant are *right now* present on the horizon of my experience summoning my attention as potentially (not merely possibly) relevant to the current situation. If Freeman is right, the attraction of familiar-but-not-currently-fully-present aspects of what is currently ready-to-hand (inner horizon) as well as the attraction of potentially relevant other familiar situations on the outer horizon of the current situation might well be correlated with the fact that our brains are not simply in one attractor basin at a time but are influenced by other attractor basins in the same landscape, and by other attractor landscapes.

According to Freeman, what makes us open to the horizonal influences of other attractors instead of our being stuck in the current attractor is that the whole system of attractor landscapes collapses and is rebuilt with each new rabbit sniff, or in our case, presumably with each shift in our attention. And once one correlates Freeman's neurodynamic account with Merleau-Ponty's description of the way the intentional arc feeds back our past experience into the way the world appears ever more familiar to us and solicits from us ever more appropriate responses to its changing significance, the frame problem of how we can deal with changing relevance by seeing what will change and what will stay the same no longer seems unsolvable.

But there is a generalization of the problem of relevance, and thus of the frame problem, that seems intractable. In *What Computer's Can't Do* I gave as an example how, in placing a racing bet, we can usually restrict ourselves to such facts as the horse's age, jockey, past performance, and competition, but there are always other factors such as whether the horse is allergic to goldenrod or whether the jockey has just had a fight with the owner, which may in some cases be decisive. Human handicappers are capable of recognizing the relevance of such facts when they come across them (Dreyfus, 1997, p. 258). But since anything in experience can be relevant to anything else, such an ability seems magical.

Jerry Fodor follows up on my pessimistic remark:

> The problem is to get the structure of an entire belief system to bear on individual occasions of belief fixation. We have, to put it bluntly, no computational formalisms that show us how to do this, and we have no idea how such formalisms might be developed . . . If someone—a Dreyfus, for example—were to ask us why

> we should even suppose that the digital computer is a plausible mechanism for the simulation of global cognitive processes, the answering silence would be deafening. (Fodor, 1983, pp. 128–129)

However, once we give up computational cognitivism, and see ourselves instead as basically coupled copers, we can see how the frame problem can be dissolved by an appeal to existential phenomenology and neurodynamics. In the light of how learning our way around in the world modifies our brain so that relevance is directly experienced in the way tasks summon us, even the general problem raised by the fact that anything in our experience could in principle be related to anything else no longer seems a mystery.

9. Conclusion

It would be satisfying if we could now conclude that, with the help of Merleau-Ponty and Freeman, we can fix what is wrong with current allegedly Heideggerian AI by making it more Heideggerian. There is, however, a big remaining problem. Merleau-Ponty's and Freeman's account of how we directly pick up significance and improve our sensitivity to relevance depends on our responding to what is significant for *us* given our needs, body size, ways of moving, and so forth, not to mention our personal and cultural self-interpretation. Thus, to program Heideggerian AI, we would not only need a model of the brain functioning underlying coupled coping such as Freeman's, but we would also need—and here's the rub—a model of our particular way of being embedded and embodied such that what we experience is significant for us in the particular way that it is. That is, we would have to include in our program a model of a body very much like ours with our needs, desires, pleasures, pains, ways of moving, cultural background, etc. If we can't make our brain model responsive to the significance in the environment as it shows up specifically for human beings, the project of developing an embedded and embodied Heideggerian AI can't get off the ground.

So, according to the view I have been presenting, even if the Heideggerian/ Merleau-Pontian approach to AI suggested by Freeman is ontologically sound in a way that GOFAI and the subsequent supposedly Heideggerian models proposed by Brooks, Agre, and Wheeler are not, a neurodynamic computer model would still have to be given a detailed description of our body and motivations like ours if things were to count as significant for it so that it could learn to act intelligently in our world. The idea of super-computers containing detailed models of human bodies and brains may seem to make sense in the wild imaginations of a Ray Kurzweil or Bill Joy, but they haven't a chance of being realized in the real world.

Acknowledgments

This review essay is a revised version of my Barwise Prize talk delivered at the 2006 Pacific Division meeting of the American Philosophical Association, which will be

published by MIT Press in Philip Husbands, Owen Holland and Michael Wheeler (Eds.), *The Mechanisation of Mind in History* (Forthcoming, title to be confirmed).

Notes

[1] Roger Schank proposed what he called "scripts." He tells us: "A script is a structure that describes appropriate sequences of events in a particular context. A script is made up of slots and requirements about what can fill those slots. The structure is an interconnected whole, and what is in one slot affects what can be in another. A script is a predetermined, stereotyped sequence of actions that defines a well-known situation" (Schank & Abelson, 1977, p. 41; as cited in Preston & Bishop, 2002, p. 17).

[2] Although you couldn't tell it from the Cog web page: http://www.ai.mit.edu/projects/humanoid-robotics-group/cog/cog.html

[3] Heidegger himself is unclear about the status of the ready-to-hand. When he is stressing the holism of equipmental relations, he thinks of the ready-to-hand as equipment, and of equipment as things like lamps, tables, doors, and rooms that have a place in a whole nexus of other equipment. Furthermore, he holds that breakdown reveals that these interdefined pieces of equipment are made of present-at-hand stuff that was there all along (1962, p. 97). At one point Heidegger even goes so far as to include the ready-to-hand under the same categories that characterize the present-at-hand: "We call '*categories*'—characteristics of being for entities whose character is not that of Dasein. . . . Any entity is either a '*who*' (existence) or a *what* (present-at-hand in the broadest sense)" (p. 71, italics added).

[4] Merleau-Ponty (1962) says the same: "[T]o move one's body is to aim at things through it; it is to allow oneself to respond to their call, which is made upon it independently of any representation" (p. 139).

[5] According to Heidegger, intentional content isn't in the mind, nor in some third realm (as it is for Husserl), *nor in the world*; it isn't anywhere. It's a way of being-towards.

[6] It's important to realize that when he introduces the term 'understanding', Heidegger (1982, p. 276) explains that he means a kind of know-how.

[7] This way of putting the source of *significance* covers both animals and people. By the time he published *Being and Time*, however, Heidegger was interested exclusively in the special kind of significance found in the world opened up by human beings who are defined by the stand they take on their own being. We might call this *meaning*. In this paper I'm putting the question of uniquely human meaning aside to concentrate on the sort of *significance* we share with animals.

[8] I'm over simplifying here. Wheeler does note that Heidegger has an account of online, involved problem solving that Heidegger calls dealing with the "un-ready-to-hand." But while for Heidegger and for Wheeler coping at its best deals directly with the ready-to-hand with no place for representations of any sort, for Heidegger but not for Wheeler *all* un-ready-to-hand coping takes place on the background of an even more basic nonrepresentational holistic coping which allows copers to orient themselves in the world. As we shall see, it is this basic coping, not any kind of problem solving, agential or subagential, that enables Heideggerian AI to avoid the frame problem.

[9] Just how Hebbian learning is translated into an attractor is not something Freeman claims to know in detail. He simply notes: "The attractors are not shaped by the stimuli directly, but by previous experience with those stimuli . . . and neuromodulators as well as sensory input. Together these modify the synaptic connectivity within the neuropil and thereby also the attractor landscape" (2000, p. 62).

[10] Quotations from Freeman's books have been reviewed by him to correspond to his latest vocabulary and way of thinking about the phenomena.

[11] In this connection Freeman speaks of "expectations" and a brain function he calls "preafference" but I suspect that this is bad phenomenology leading to dubious neuro-speculation. Once the stimulus has been classified by selecting an attractor that says eat this now, the problem for the brain is just how this eating is to be done. Online coping needs a stimuli-driven feedback policy dictating how to move rapidly over the terrain and approach and eat the carrot. Here, an actor-critic version of temporal difference reinforcement learning (TDRL) is needed to augment the Freeman model. According to TDRL, learning the appropriate movements in the current situation requires learning the expected final award as well as the movements. These two functions are learned slowly through repeated experiences. Then the brain can monitor directly whether the expectation of reward is being met as the rabbit approaches the carrot to eat it. There need be no expectation of a goal state.

If the expected final reward suddenly decreases due, for example, to the current inaccessibility of the carrot, the relevant part of the brain prompts the olfactory bulb to switch to a new attractor or perspective on the situation that dictates a different learned action, say dragging the carrot with *its* expected reward. Only after a skill is thus acquired can the current stimuli, plus the past history of responding to related stimuli now wired into cell assemblies, produce the rapid responses required for on-going skillful coping.

[12] If it seems that much of the time we don't experience any such pull toward the optimal, Merleau-Ponty would no doubt respond that the sensitivity to deviation is nonetheless guiding one's coping, just as an airport radio beacon doesn't give a warning signal unless the plane strays off course, and then, let us suppose, the plane gets a signal whose intensity corresponds to how far off course it is and the intensity of the signal diminishes as it approaches getting back on course. The silence that accompanies being on course doesn't mean the beacon isn't continually guiding the plane, likewise in the case of the absence of tension felt in perception.

References

Agre, P. E. (1988). *The dynamic structure of everyday life* (Technical Report 1085). Cambridge, MA: MIT Artificial Intelligence Lab.

Agre, P. E. (1997). *Computation and human experience*. Cambridge, England: Cambridge University Press.

Brooks, R. A. (1997a). From earwigs to humans. *Robotics and Autonomous Systems*, *20*, 291–304.

Brooks, R. A. (1997b). Intelligence without representation. In J. Haugeland (Ed.), *Mind design II: Philosophy, psychology, and artificial intelligence* (pp. 395–420). Cambridge, MA: MIT Press, (Original work published 1991).

Brooks, R. A. (2002). *Flesh and machines: How robots will change us*. New York: Vintage Books.

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, *58*, 7–19.

Dennett, D. (1994). The practical requirements for making a conscious robot. *Philosophical Transactions of the Royal Society of London, Series A*, *349*, 133–146.

Dennett, D. (1997). Consciousness in human and robot minds. In M. Ito, Y. Miyashita & E. T. Rolls (Eds.), *Cognition, computation and consciousness* (pp. 17–30). Oxford, England: Oxford University Press.

Dreyfus, H. (1972). *What computers can't do: A critique of artificial reason*. New York: Harper & Row.

Dreyfus, H. (1992). *What computers still can't do: A critique of artificial reason*. Cambridge, MA: MIT Press.

Dreyfus, H. L. (1997). *What computers can't do*. New York: Harper and Row.

Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.

Freeman, W. J. (1991). The physiology of perception. *Scientific American*, *264*, 78–85.

Freeman, W. J. (1995). *Societies of brains: A study in the neuroscience of love and hate*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Freeman, W. J. (2000). *How brains make up their minds*. New York: Columbia University Press.

Freeman, W. J., & Grajski, K. A. (1987). Relation of olfactory EEG to behavior: Factor analysis. *Behavioral Neuroscience*, *101*, 766–777.

Haugeland, J. (1998). Mind embodied and embedded. *Having thought: Essays in the metaphysics of mind* (pp. 207–237). Cambridge, MA: Harvard University Press.

Heidegger, M. (1962). *Being and time* (J. Macquarrie & E. Robinson, Trans.). New York: Harper & Row.

Heidegger, M. (1976). *Gesamtausgabe: Vol. 21. Logik: die Frage nach der Wahrheit*. Frankfurt am Main: Vittorio Klostermann.

Heidegger, M. (1982). *The basic problems of phenomenology* (A. Hofstadter, Trans.). Bloomington, IN: Indian University Press. 276.

Heidegger, M. (2002). Phenomenological interpretations in connection with Aristotle. In J. Van Buren (Ed.), *Supplements: From the earliest essays to being and time and beyond* (pp. 111–146). Albany, NY: State University of New York Press.

Heidegger, M. (in press). *Logic: The question of truth* (Thomas Sheehan, Trans.).

Husserl, E. (1973). *Experience and judgment*. Evanston, IL: Northwestern University Press.

Kelly, S. D. (2005). Seeing things in Merleau-Ponty. In T. Carman & M. B. N. Hansen (Eds.), *The Cambridge companion to Merleau-Ponty* (pp. 74–110). Cambridge, England: Cambridge University Press.

Kozma, R., & Freeman, W. J. (2003). Basic principles of the KIV model and its application to the navigation problem. *Journal of Integrative Neuroscience*, *2*, 125–145.

Merleau-Ponty, M. (1962). *Phenomenology of perception* (C. Smith, Trans.). London: Routledge & Kegan Paul.

Merleau-Ponty, M. (1966). *The structure of behavior* (A. L. Fisher, Trans., 2nd ed.). Boston: Beacon Press.

Preston, J., & Biship, M. (Eds.) (2002). *Views into the Chinese room: New essays on Searle and artificial intelligence*. Oxford, England: Clarendon Press.

Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals and understanding: An inquiry into human knowledge structures*. Hillsdale, NJ: Lawrence Erlbaum.

Todes, S. (2001). *Body and world*. Cambridge, MA: MIT Press.

van Gelder, T. (1997). Dynamics and cognition. In J. Haugeland (Ed.), *Mind design II* (pp. 421–450). Cambridge, MA: MIT Press.

Wheeler, M. (2002). Change in the rules: Computers, dynamical systems, and Searle. In J. Preston & M. Bishop (Eds.), *Views into the Chinese room: New essays on Searle and artificial intelligence* (pp. 338–359). Oxford, England: Clarendon Press.

Wheeler, M. (2005). *Reconstructing the cognitive world: The next step*. Cambridge, MA: MIT Press.

Winograd, T. (1976). Artificial intelligence and language comprehension. *Artificial intelligence and language comprehension* (pp. 1–26). Washington, DC: National Institute of Education.

Winograd, T. (1989, September). Heidegger and the design of computer systems. Paper delivered at Applied Heidegger Conference, Berkeley, CA.

Winograd, T. (2000). Foreward. In M. Wrathall (Ed.), *Essays in honor of Hubert L. Dreyfus: Vol. 2. Heidegger, coping, and cognitive science* (pp. vii–x). Cambridge, MA: MIT Press.