A Phenomenology of Skill Acquisition as the basis for a Merleau-Pontian Non-representationalist Cognitive Science

Hubert L. Dreyfus

Department of Philosophy

University of California, Berkeley

I. Introduction

Existential phenomenologists hold that the two most basic forms of intelligent behavior, learning, and skillful action, can be described and explained without recourse to mind or brain representations. This claim is expressed in two central concepts in Merleau-Ponty's Phenomenology of Perception[1]: the intentional arc and the tendency toward achieving a maximal grip. The *intentional arc* names the tight connection between the agent and the world, viz. that, as the agent acquires skills, these skills are "stored", not as representations in the mind, but as more and more refined dispositions to respond to the solicitations of more and more refined perceptions of the current situation. *Maximum grip* names the body's tendency to respond to these solicitations in such a way as to bring the current situation closer to the agent's sense of an optimal gestalt.

I will argue that neither of these abilities requires mental or brain representations. Rather, simulated neural networks exhibit crucial structural features of the intentional arc, and Walter Freeman's account of the brain dynamics underlying perception is structurally isomorphic with Merleau-Ponty's account of the way a skilled agent moves towards the sense of equilibrium that signals a maximum grip.

II. Skill Acquisition: The Establishment of the Intentional Arc.

According to Merleau-Ponty our skills are acquired by dealing with things and situations, and in turn they determine how things and situations show up for us as requiring our responses. To appreciate this claim we need to lay out more fully than Merleau-Ponty does how our relation to the world is transformed as we acquire a skill.

Many of our skills are acquired at an early age by trial and error or by imitation, but to make the establishment of the intentional arc as perspicuous as possible, I will consider the case of an adult acquiring a skill by instruction.[2]

Stage 1: Novice

Normally, the instruction process begins with the instructor decomposing the task environment into context-free *features*, which the beginner can recognize without previous experience in the task domain. The beginner is then given *rules* for determining actions on the basis of these features, and so acts like a computer following a program.

The student automobile driver learns to recognize such domain-independent features as speed (indicated by the speedometer) and is given rules such as shift to second when the speedometer needle points to ten. The novice chess player learns a numerical value for each type of piece regardless of its position, and the rule: "Always exchange if the total value of pieces captured exceeds the value of pieces lost."

But merely following rules will produces poor performance. A car stalls if one shifts too soon on a hill or when the car is heavily loaded; a chess player who always exchanges to gain points is sure to be the victim of a sacrifice by the opponent who gives up valuable pieces to gain a tactical advantage.

The learner obviously not only needs the facts but also an understanding of the context in which the facts makes sense.

Stage 2: Advanced Beginner

As the novice gains experience actually coping with real situations, he begins to note, or an instructor points out, perspicuous examples of meaningful additional *aspects* of the situation. After seeing a sufficient number of examples, the student learns to recognize these new aspects. Instructional *maxims* now can refer to these new situational aspects, recognized on the basis of experience, as well as to the objectively defined non-situational features recognizable by the inexperienced novice.

The advanced beginner driver uses (situational) engine sounds as well as (non-situational) speed in deciding when to shift. He learns the maxim: Shift up when the motor sounds like it’s racing and down when it sounds like it’s straining. Engine sounds cannot be adequately captured by a list of features, so features cannot take the place of a few choice examples in learning the relevant distinctions.

With experience, the chess beginner learns to recognize overextended positions and how to avoid them. Similarly, he begins to recognize such situational aspects of positions as a weakened king's side or a strong pawn structure, despite the lack of precise and situation-free definitions. The player can then follow maxims such as: attack a weakened king’s side. Unlike a rule, a maxim requires that one already have some understanding of the domain to which the maxim applies.[3]

Still, at this stage, learning is carried on in a detached, analytic frame of mind, as the student follows instructions and is given examples. To progress further, however, requires a special sort of involvement.

Stage 3: Competence

With more experience, the number of potentially relevant elements that the learner is able to recognize becomes overwhelming. To cope with this overload and to achieve competence, people learn, through instruction or experience, to devise a plan, or choose a perspective, that then determines which elements of the situation are important and which ones can be ignored. As students learn to restrict themselves to only a few of the vast number of possibly relevant features and aspects, understanding and decision making becomes easier.

Naturally, to avoid mistakes, the competent performer seeks rules and reasoning procedures to decide which plan or perspective to adopt. But such rules are not as easy to come by, as are the rules and maxims given beginners in manuals and lectures. Indeed, in any skill domain the performer encounters a vast number of situations differing from each other in subtle ways. There are, in fact, more situations than can be named or precisely defined, so no one can prepare a list of types of possible situations and what to do or look for in each. Competent performers, therefore, must decide for themselves in each situation what plan or perspective to adopt without being sure that it will turn out to be appropriate.

Given this uncertainty, coping becomes frightening rather than merely exhausting. Prior to this stage, if the rules don't work, the performer, rather than feeling remorse for his mistakes, can rationalize that he hadn't been given adequate rules. But, since at this stage, the result depends on the learner's choice of perspective, the learner feels responsible for his or her choice. Often, the choice leads to confusion and failure. But sometimes things work out well, and the competent student then experiences a kind of elation unknown to the beginner.

A competent driver leaving the freeway on an off-ramp curve, learns to pay attention to the speed of the car, not whether to shift gears. After taking into account speed, surface condition, criticality of time, etc., he may decide he is going too fast. He then has to decide whether to let up on the accelerator, remove his foot altogether, or step on the brake, and precisely when to perform any of these actions. He is relieved if he gets through the curve without mishap, and shaken if he begins to go into a skid.

The class A chess player, here classed as competent, may decide after studying a position that her opponent has weakened his king's defenses so that an attack against the king is a viable goal. If he chooses to attack, she ignores weaknesses in her own position created by the attack, as well as the loss of pieces not essential to the attack. Pieces defending the enemy king become salient. Since pieces not involved in the attack are being lost, the timing of the attach is critical. If he attacks too soon or too late, his pieces will have been lost in vain and she will almost surely lose the game. Successful attacks induce euphoria, while mistakes are felt in the pit of the stomach.

So the learner is naturally frightened, elated, disappointed, or discouraged by the results of his or her choice of perspective. And, as the competent student become more and more emotionally involved in his task, it becomes increasingly difficult for him to draw back and adopt the detached maxim-following stance of the advanced beginner. Only at the level of competence is there an emotional investment in the *choice of action.* Then, emotional involvement seems to play an essential role in switching over from what one might roughly think of as a left-hemisphere analytic approach to a right-hemisphere holistic one.[4]

Stage 4: Proficient

Only if the detached, rule-following stance of the novice, advanced beginner is replaced by involvement, is the student set for further advancement. Then, the resulting positive and negative emotional experiences will strengthen successful responses and inhibit unsuccessful ones, and the performer's theory of the skill, represented in rules and principles, will gradually be replaced by situational discriminations, accompanied by associated responses. Only if experience is assimilated in this embodied, atheoretical way do intuitive reactions replace reasoned responses.[5]

To understand this stage of skill acquisition we must remember that the involved, experienced performer sees goals and salient aspects, but not what to do to achieve these goals. This is inevitable since there are far fewer ways of seeing what is going on than there are ways of reacting. The proficient performer simply has not yet had enough experience with the outcomes of the wide variety of possible responses to each of the situations he can now discriminate, to react automatically. Thus, the proficient performer, after spontaneously seeing the point and the important aspects of the current situation, must still *decide* what to do. And to decide, he must fall back on detached rule and maxim following.

The proficient driver, approaching a curve on a rainy day, may *feel in the seat of his pants* that he is going dangerously fast. He must then *decide* whether to apply the brakes or merely to reduce pressure by some specific amount on the accelerator. Valuable time may be lost while making a decision, but the proficient driver is certainly more likely to negotiate the curve safely than the competent driver who spends additional time *considering* the speed, angle of bank, and felt gravitational forces, in order to *decide* whether the car's speed is excessive.

The proficient chess player, who is classed a master, can recognize almost immediately, a large repertoire of types of positions. He then deliberates to determine which move will best achieve his goal. He may know, for example, that he should attack, but she must calculate how best to do so.

Stage 5: Expertise

The proficient performer, immersed in the world of his skillful activity, *sees* what needs to be done, but must *decide* how to do it. The expert not only sees what needs to be achieved; thanks to a vast repertoire of situational discriminations he sees immediately what to do. Thus, the ability to make more subtle and refined discriminations is what distinguishes the expert from the proficient performer. Among many situations, all seen as similar with respect to a plan or perspective, the expert has learned to distinguish those situations requiring one action from those demanding another. That is, with enough experience in a variety of situations, all seen from the same perspective but requiring different tactical decisions, the brain of the expert performer gradually decomposes this class of situations into subclasses, each of which shares the same action. This allows the immediate intuitive situational response that is characteristic of expertise.

The chess Grandmaster experiences a compelling sense of the issue and the best move. Excellent chess players can play at the rate of 5 to 10 seconds a move and even faster without any serious degradation in performance. At this speed they must depend almost entirely on intuition and hardly at all on analysis and comparison of alternatives. It has been estimated by cognitivists such as Herbert Simon that an expert chess player remembers roughly 50,000 types of positions. As we shall see, there is no need for the expert to *remember* any positions, but for expert performance, the number of classes of *descriminable* situations, built up on the basis of experience, must be comparatively large.

The expert driver, not only feels in the seat of his pants when speed is the issue; he knows how to perform the appropriate action without calculating and comparing alternatives. On the off-ramp, his foot simply lifts off the accelerator and applies the appropriate pressure to the brake. What must be done, simply is done. As Aristotle says, the expert “straightway” does “the appropriate thing, at the appropriate time, in the appropriate way.”

We can see now that skilled action does not require mental representations. A beginner calculates using rules and facts just like a programmed computer, but with talent and a great deal of involved experience, the beginner develops into an expert who intuitively sees what to do without recourse to rules nor to remembered cases. The tradition has given an accurate description of the mental representations used by beginners and of the experts facing unfamiliar situations, but normally an expert just immediately does what normally works and, of course, it normally works.

III. Learning without Brain Representations: Merleau-Ponty’s Intentional Arc and Feed-Forward Neural Networks

We can now see why, according to Merleau-Ponty, what the learner acquires through experience is not *represented* in the mind at all but is *presented* to the learner as more and more finely discriminated situations, and that, if the situation does not clearly solicit a single response or the response does not produce a satisfactory result, the learner is led to further refine his discriminations, which, in turn, solicit more refined responses. Merleau-Ponty calls this feedback loop between the embodied agent and the perceptual world the intentional arc. He says:

> Cognitive life, the life of desire or perceptual life – is subtended by an 'intentional arc' which projects round about us our past, our future, [and] our human setting.[6]

The agent does not merely receive input passively and then process it. Rather, the agent immediately sees things from some perspective and sees them as affording a certain action.

Merleau-Ponty holds that no mentalistic model, whether empiricist or idealist, can account for the way past learning is manifest in present experience, but, fortunately, there are now models of what might be going on in the brain that make no use of empiricist association nor of the sort of symbols, rules, and remembered cases presupposed in rationalist philosophy and classical Cognitive Science. Such models are called feed-forward simulated neural networks. They consist of an input layer of simulated neurons connected to an output layer of neurons in all possible combinations by way of one or more intermediate layers. When the net is trained, each time the net makes what the trainer considers a mistaken association of input with output, the weights on the connections between the neurons are changed according to a back propagation algorithm that adjusts the weight on each connection to the extent that that connection was responsible for the mistake.

After such training, similar inputs will produce the same or similar outputs. If the input corresponds to the experience of the current situation, the already given activation of the hidden nodes, determined by inputs leading up to the current situation, might be said to correspond to the expectations or perspective that the expert brings to the situation, in terms of which the situation solicits a specific response. It is precisely the advantage of simulated neural networks that past experience, rather than being stored as a memory, modifies the connection strengths between the simulated neurons. New input can then produce output based on past experience without the net having to, or even being able to, represent its past experience.

Still there are many important ways in which neural nets differ from embodied brains. Some of these ways are limitations that can be overcome by further research. Thus, nets now depend for their learning on people giving them examples by pairing input and output, but work is underway on reinforcement learning techniques that enable the nets to learn directly by feedback from their successes and failures in the target domain.[7]

A more fundamental difficulty, however, is endemic to neural net learning. Whether the net learns by being given appropriate situation-action pairs or by finding for itself which pairings work, in order to learn to recognize the sorts of situations and things we recognize and to respond appropriately, a network must respond to the same similarities human beings respond to. But everything is similar to everything else and different from everything else in an indefinitely large number of ways. We just do not notice it. This leads to the problem of generalization.

Neural-network modelers agree that an intelligent network must be able to generalize. For example, for a given classification task, given sufficient examples of inputs associated with one particular output, it should associate further inputs of the same type with that same output. But what counts as the same type? The network's designer usually has in mind a specific definition of type required for a reasonable generalization and counts it a success if the net generalizes to other instances of this type. But when the net produces an unexpected association, can one say it has failed to generalize? One could equally well say that the net has all along been acting on a different definition of type, based on different perceived similarities, and that that difference has just been revealed. One might think of this as an alien sort of intelligence, but if a neural-net does not respond to the same types of situations as similar that human beings do, it will not be able to learn our skills and so will fail to find its way about in our world.

But there seems to be a puzzle here. How do human beings learn to generalize like other human beings so that they acquire the skills required to get around in the human world? If everything is similar to everything else in an indefinitely large number of ways, what constrains the space of possible generalizations so that trial and error learning has a chance of succeeding?

Merleau-Ponty would no doubt hold that the fact that we have bodies is essential to understanding how we generalize. There are at least three ways the human body constrains the space of possible generalizations. The first is due to the brain; the other two are due to our actual body structure.

First, the possible responses to a given input must be constrained by brain architecture. This innate structure accounts for phenomena such as the perceptual constants the Gestaltists investigated. These are given from the start by the perceptual system as if they had always already been learned. Merleau-Ponty calls them “déjà monté.”[8]

But this alone would not be enough to constrain the generalization-space so that all human beings learn to respond to the same set of inputs as similar. It turns out, however, that the order and frequency of the inputs further constrains how a net will generalize. This order is determined by the trainer in what is called supervised learning, but if the net is to learn by itself, that is, if its connection strengths are to be allowed to adjust themselves on the basis of the input-output pairs it encounters, then the order and frequency of inputs will depend on the interaction of the structure of the embodied network and the structure of the world. For example, things nearby that afford reaching will be noticed early and often. Their various ways of being reachable and the kind of grip they provide will be an obvious source of shared similarities. Thus, body-dependent order and similarity of presentation provides the second constraint on generalization.[9]

The third constraint depends on what counts as success. In supervised learning, the researcher defines what counts as success in each specific domain. And, indeed, in everyday learning success is often understood as the achievement of a previously represented goal. But in reinforcement learning simulations, although the simulated learner is at first given a goal and told of successes and failures like a novice, the simulation gradually develops the equivalent of sense of how it is doing at each moment without needing to represent the goal.[10] We must, therefore, now turn to the way in everyday coping we take account of improvement without a representation of what would count as success.[11]

IV. Maximum Grip: Action without Representing a Goal

So far we have seen that simulated feed-forward neural networks exhibit crucial structural features of the intentional arc. We must now return to Merleau-Ponty's account of maximal grip. According to Merleau-Ponty, higher animals and human beings are always tending towards getting a *maximum grip* on their situation. To take Merleau-Ponty's example:

> For each object, as for each picture in an art gallery, there is an optimum distance from which it requires to be seen, a direction viewed from which it vouchsafes most of itself: at a shorter or greater distance we have merely a perception blurred through excess or deficiency. We therefore tend towards the maximum of visibility, and seek a better focus as with a microscope.[12]

More generally,

> my body is geared into the world when my perception presents me with a spectacle as varied and as clearly articulated as possible, and when my motor intentions, as they unfold, receive the responses they expect from the world.[13]

So, in our skilled activity we move to achieve a better and better grip on our situation. For this movement towards maximum grip to take place, one does not need a mental representation of one's goal. Rather, acting is experienced as a steady flow of skillful activity in response to one's sense of the situation. Part of that experience is a sense that when one's situation deviates from some optimal body-environment relationship, 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. One's body is simply solicited by the situation to get into equilibrium with it. As Merleau-Ponty puts it:

> Our body is not an object for an 'I think', it is a grouping of lived-through meanings which moves towards its equilibrium.[14]

To get the phenomenon in focus, consider a tennis swing. If one is a beginner or is off one's form one might find oneself making an effort to follow rules: to keep one's eye on the ball, keep the racket perpendicular to the court, hit the ball squarely, and so forth. But if one is an expert, what one experiences is more like one's arm going up and its being drawn to the appropriate position, the racket forming the optimal angle with the court -- an angle one need not even be aware of -- all this so as to complete the gestalt made up of the court, one's running opponent, and the oncoming ball. One feels that one's comportment was caused by the perceived conditions in such a way as to reduce a sense of deviation from some satisfactory gestalt. But that final gestalt need not be represented in one's mind. Indeed, it is not something one *could* represent. One only senses when one is getting closer or further away from the optimum.

Thus, such skillful coping does not require a mental representation of its goal. It can be *purposive* without the agent entertaining a *purpose.* As Merleau-Ponty puts it:

> [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.[15]

To distinguish this body-based intentionality from the representational intentionality studied by Husserl and Cognitive Science, Merleau-Ponty calls the body's response to the affordances of the situation, motor intentionality.

To help convince us that no representation of the final gestalt is needed for the skilled performer to move towards it, Merleau-Ponty uses the analogy of a soap bubble. The bubble starts as a deformed film. The bits of soap respond to local forces according to laws that happen to work so as to dispose the entire system to end up as a sphere, but the spherical result does not play a causal role in producing the bubble. The same holds for the final gestalt of body and racket in my example. Indeed, I cannot represent how I should turn my racket since I do not know what I do when I return the ball. I may once have been told to hold my racket perpendicular to the court, and I may have succeeded in doing so, but now experience has sculpted my swing to the situation in a far more subtle and appropriate way than I could ever achieve following a rule.

According to Merleau-Ponty, we not only move to complete a good gestalt in any skill domain, we also tend to improve what counts as a good gestalt in that domain. As we have seen, the involved performer tends to discriminate more and more refined situations and pair them with more and more appropriate actions. Thus, the intentional arc is steadily enriched and refined.

So we can now add that one additional way experiences can count as similar is that, in each situation, those actions will count as similar that reduce disequilibrium and thereby give the agent a sense of improvement. Thus, motor intentionality gives us an additional embodied sense of similarity that would not be available to disembodied neural nets.

V. A brain model that explains Merleau-Ponty's account of Maximal Grip

To understand how non-representational motor intentionality works, we can begin by considering a game in which one player guides the other's search for some hidden object by saying "hot" or "cold." In that case the performer is led by the clues without knowing where they are leading. Of course, in the hot/cold game, the player giving the clues needs to know where the hidden object is, and Merleau-Ponty admits that it seems impossible that an agent could have a sense of whether his grip was improving without sensing what would count as success. Since he was clear that no account of brain function conceivable in his day could account for this phenomenon, Merleau-Ponty called it magical.[16]

Fortunately, Walter Freeman has worked out a model of learning that can be adapted to show how the brain, operating as a dynamical system, could cause a series of movements that achieve a goal without the brain in any way representing that goal in advance of achieving it.[17] According to Freeman's model of learning, after an animal has repeatedly encountered a situation in which a particular response has produced results that are useful or harmful to the animal, it forms neuron connections which, when the animal encounters stimuli from a similar situation, causes the neurons to produce a burst of global activity whose energy state occupies a point in an energy landscape. An energy landscape is composed of several attractors. In Freeman's model of learning, the animal's brain forms a new attractor each time the animal learns to respond to a new type of situation.

Now we can explain the phenomenon of being drawn toward getting a maximum grip. Applying Freeman's model to action, we can suppose that, through exposure to satisfactions and frustrations brought about by specific actions in a number of similar situations, the sensory-motor system forms an energy landscape that is shaped by the possibilities for successful comportment in that type of situation. When a specific sensory input moves the system-state into the vicinity of a specific attractor, the organism is caused to move in a way that brings the system-state closer to the bottom of that basin of attraction. The tennis player's experience, in my example, of a tension drawing him to move towards a satisfactory gestalt would, on this account, be correlated with the tendency of his sensory-motor system to relax into a specific minimum energy state.

At any given moment, the system, like the player in the "hot" and "cold" game, is in a state that is near or far from the bottom of some specific basin. But, if that were all that was going on in the person's brain, the person would be like a player who could only guess where to look next, and so at best could find what he was seeking by trial and error.

Happily, the energy landscape gives more information than just "hot" or "cold." In our hypothetical case, as soon as the experienced tennis player's perception of the situation brings his sensory-motor system under the pull of a specific attractor, his brain's relaxing into a basin of attraction is correlated with his sense of which direction of movement would make him hotter, without his knowing where the hottest point is. The system thus underlines the player's being drawn to make those movements that result in his feeling a lowering of tension -- the same movements that result in his brain-state approaching the lowest accessible point in its current energy landscape. All this happens without the brain representing the lowest energy state in advance and without the player's needing to represent to himself what the final equilibrium state would be like or how to get there. As Merleau-Ponty already pictured it, the person's brain would simply be moving to lower a tension, like a soap bubble relaxing into a spherical shape without in any way representing the spherical shape toward which it was tending, while the player would simply feel drawn to lower a tension without knowing in advance the shape of the equilibrium to which he was tending.[18]

Obviously, the sort of knowledge such a system embodies could not be something one was conscious of and so could not be understood as a conscious or unconscious representation. The attractor could be called a representation only in the very weak sense that it does incorporate past experience and leads to appropriate action on the bases of that past experience. Thus, thanks to Freeman's work, Merleau-Ponty's claim that the representationalist account of our most basic and pervasive form of learning and skillful action is mistaken can be defended not only on the phenomenological grounds but on neurological grounds as well.[19]

---

[1] Maurice Merleau-Ponty, Phenomenology of Perception,, Routledge & Kegan Paul, 1962.

[2] For a detailed treatment of the phenomenology of skill acquisition, see H. Dreyfus and S. Dreyfus, Mind Over Machine, Free Press, 1982, and The Road to Mastery, MIT Press, forthcoming.

[3] See, Michael Polanyi, Personal Knowledge, Routledge and Kegan Paul, 1958.

[4] That amateur and expert chess playing uses different parts of the brain has been confirmed by recent MRI research. The researchers report:

"Here we use a new technique of magnetic imaging to compare focal bursts of y-band activity in amateur and professional chess players during matches. We find that this activity is most evident in the medial temporal lobe in amateur players, which is consistent with the interpretation that their mental acuity is focused on analysing unusual new moves during the game. In contrast, highly skilled chess grandmasters have more γ-bursts in the frontal and parietal cortices, indicating that they are retrieving chunks from expert memory by recruiting circuits outside the medial temporal lobe." [It should be noted that the cognitivist claim that the MRI results confirm that experts "are retrieving

---

chunk's (of typical chess positions) from memory" is in no way supported by this research. What does seem to be support is the weaker claim that:

"These marked differences in the distribution of focal brain activity during chess playing point to differences in the mechanisms of brain processing and functional brain organization between grandmasters and amateurs."

Ognjen Amidzic, Hartmut J. Riehle, Thorsten Fehr, Christian Weinbruch, Thomas Elbert, "Patterns of focal y-bursts in chess players: Grandmasters call on regions of the brain not used so much by less skilled amateurs," Nature**:** Brief Communications, Vol. 412, 9 August 2001, p. 603.

[5] Of course, not just any emotional involvement such as enthusiasm, or fear of making a fool of oneself, or the exultation of victory, will do. What matters is taking responsibility for one's successful and unsuccessful choices, even brooding over them; not just feeling good or bad about winning or losing, but replaying one's performance in one's mind step by step or move by move. The point, however, is not to *analyze* one's mistakes and insights, but just to *let them sink in*. Experience shows that only then will one become an expert.

[6] Ibid., 136.

[7] Richard S. Sutton and Andrew G. Barto, Reinforcement Learning, Bradford, MIT Press, Cambridge, MA., 1998.

[8] Find page reference.

[9] For a worked out account of human body-structure and how it is correlative with the structure of the human world, see Samuel Todes, Body and World, MIT Press, 2001.

Of course, the body dependence puts disembodied neural-networks at a serious disadvantage when it comes to learning to cope in the human world. Nothing is more alien to our life-form than a network with no up/down, front/back orientation, no

---

interior/exterior distinction, no preferred way of moving, such as moving forward more easily than backwards, and no emotional response to its failures and successes. The odds are overwhelming against such a net being able to respond to the similarities we do and so to learn to classify situations and affordances as we do. In our world the cards are stacked to enable entities that share our embodied form of life to learn to cope in a way we find intelligent, while leaving disembodied creatures looking to us hopelessly stupid.

[10] E. Mizutani and S. E Dreyfus "Totally Model-Free Reinforcement Learning by Actor-Critic Elman Networks in Non-Markovian Domains", Proceedings of the IEEE World Congress on Computational Intelligence (Wcci'98), Alaska USA, May 1998 pp.2016—2021.

[11] For a phenomenological account of the difference between propositional *success conditions* and nonpropositional *conditions of improvement* see Hubert L. Dreyfus, “The Primacy of Phenomenology over Logical Analysis,” Philosophical Topics, Mark A. Wrathall and Hubert L. Dreyfus. Eds, Vol. 27, No. 2, Fall 1999, (2001).

[12] Phenomenology of Perception, 302.

[13] Ibid., 250.

[14] Ibid., 153.

[15] Ibid., 139.

[16] Ibid., 103.

It’s important to note that Merleau-Ponty uses “magical” in two ways. In discussing how the mind can control movement he says, “We still need to understand by what magical process the representation of a movement causes precisely that movement to be made by the body.” And he adds, “The problem can be solved provided that we cease to draw a distinction between the body as a mechanism in itself and consciousness as being for itself,” (139) Here he is using the term magical pejoratively to mean that a

---

causal claim is based on an ontology that makes it *impossible* to account for how it could be implemented.

In the case just cited, however, Merleau-Ponty uses "magical" to mean that there is no *currently conceivable* way to cash out the causal claim that absorbed coping is directed towards a goal without representing that goal.

[17] Walter J. Freeman, "The Physiology of Perception", Scientific American, 264: 78-85, 1991a; and W. J. Freeman and K.A. Grajski, "Relation of olfactory EEG to behavior: Factor Analysis," Behavioral Neuroscience, 101: 766-777, 1987.

[18] It is important to bear in mind that this account is only supposed to cover skillful coping in flow; not cases of deliberate action. But even thinking of all absorbed coping as moving so as to reduce a felt tension and so *reach equilibrium* is obviously an oversimplification. Some absorbed coping, like carrying on a conversation, does not seem to be governed by a tendency to reduce tension. The Freeman/Merleau-Ponty model applies best to the basic skills we have for getting around in the world. And, even when one is being draw to reduce a tension and so is tending towards equilibrium, one usually finds oneself in a new situation before one arrives at stasis, and so one is drawn towards a new equilibrium before the first equilibrium is actually reached. Thus, the agent is continually drawn towards some equilibrium state or other but seldom arrives. In this way the Platonic/Freudian vision that people act so as to arrive at fulfillment and rest might be reconciled with the Aristotelian/Deweyian idea that the good life consists in on-going coping. Everyday, comportment might well consist in a tendency to achieve equilibrium that never arrives at equilibrium but, instead, produces on-going activity.

[19] A further Merleau-Pontian feature of Freeman's model of perception is that the brain does not form conditioned responses to specific stimuli but, on the basis of experience, produces its own attractors that are evoked and modified on the basis of further

experience. Once the stimulus from the current situation has triggered a burst of neuronal activity that forms a specific attractor landscape, the attractor landscape takes over and draws the system to relax into a specific attractor. Thus, once the sensory input has put the system into a specific attractor landscape, it has no further job and so can, as Freeman puts it, be “thrown away.”

As Freeman says in his paper presented this morning: “The patterns generated by cortex are not representations of stimuli. They are neural discharges that give the meanings of stimuli for individuals. They depend on experiences that have been embedded in the synapses in cortical networks, which are different for everyone. The sensory cortices broadcast these spatial patterns, while the raw sense data, the phantasms, having done their work, are removed.”