"It is far too little recognized how entirely the intellect is built up of practical interests" (William James, 1897)

Excerpt from Chapter 1--
The Homeostatic Mind: A Developmental Study of Object Cognition
Russ McBride

I. The Project

Every day you manipulate objects. You grab a doorknob, press a gas pedal, or hold a glass of water. Every day you think about objects. The amazing fact is that somehow from amid the onslaught of stimuli that you continuously experience you are able to organize it into objects so you can do those things. You are able to interact with a world that presents itself as pre-divided into objects. "And of course it is", we say. There are chairs and sneakers and rabbits and bicycles. But at one point in your early past you did not understand the world as flush with objects; it was at best an undifferentiated sensory onslaught. What changed to allow you to discover objects in the world, to make it a world filled with objects *for you*? Well, you did, of course. But how? This project is an exploration into the birth of those everyday physical objects of perceptual experience.

This seems like such a simple problem. Objects are everywhere, after all. Surely we just interact with and think about what's already out there in the world—tables, dogs, glasses of water, etc. The problem of how we develop the ability to discern objects in our world is fundamental, but not obviously easy. We currently have little idea how it happens. One might, with only a bit of exaggeration, argue that this problem is the most fundamental problem in cognitive science and the philosophy of mind. Learning, perception, memory, reasoning, interacting in the world, planning, and talking about any thing rely on our ability to see the world as consisting of things, i.e., of objects. All of these abilities are abilities that are exercised over objects.

Here's one theory: we evolved to 'mirror' (in our cognitive systems) what is already out there in the real world. Only once we matured enough to represent internally the objects out in the world were we then able to successfully interact with those objects via their copies. This is arguably the dominant picture in cognitive science today (cf. Wheeler 2005). It's grounded in a strong realism about the external world, a representational approach to cognition, and a computational account about how those representations are manipulated. The external world has an objective existence apart from us, we make representational copies of objective objects from that world, and manipulate those copies by applying standard algorithms to them, rules well known from logic, statistics, and computer science. This view has the effect of making it easy to gloss over the core issue, the process by which those representational copies are formed—"Oh, we just pick up what nature has already laid out in front of us."

My project is a sustained effort to take up the torch of American pragmatists like William James and John Dewey and shine light into the recess of that gloss, that dark basement which we rush past with phrases like 'pick up' or 'copy', and to ask the simple question, before we get comfortable with talk of 'representations'. How do we get to the point sometime during the early years of our life when we experience the world as consisting of physical objects, objects not unlike those we experience now? How does our world come to be inhabited by trees and cats and tennis balls and cars?

Despites the efforts of Dreyfus (1979) and (1992), Brooks (1987) and (1991), Johnson (2007), Varela, et al (1991), Nöe (2001a, b), Wheeler (2005), and others, based on work by Heidegger, Dewey, and Merleau-Ponty, contemporary cognitive scientists are still, by and large, 'representation-philes'. A sore thumb on the claim that objects are mere representational "givens" is that it fails to take seriously a simple fact: our cognitive system has to do a lot of work for a long time under the tent of a fascinating and elaborate biological circus of activity, without which you would have never been able to see a cat or pick up a tennis ball. In short, our cognitive system develops in such a way that we effortlessly and regularly experience a world rich in objects. An object is *not* a "given" for an embryo or a pre-born or, for the majority of objects, even a newborn. And the standard account fails to take seriously the central question—*how* does our cognitive system do it? How can we explain the discernment of objects in our world as we mature? It seems like pulling a rabbit out of a hat. Restating the position of realism—saying that the tennis ball exists objectively in a real external world—doesn't get us any closer to answering these questions. In fact, in some sense, it makes it harder as we lull ourselves into complacency with the idea that the problem is "more or less" solved with talk about 'representations' and 'copies' of objects.

This is a work about those middle-sized physical objects sometimes called 'physical goods', objects that fall under Rosch's 'basic-level categories', i.e., cars and apples and rabbits. You might still think at this point that the problem isn't so hard. You might think that if we can get a handle on how we *differentiate* objects from one another that we're halfway to a solution, and we differentiate a tennis ball from a baseball by means of the simple fact that a tennis ball just looks different from a tennis ball. The qualitative differences in our different experiences do the work. In chapter two I try to show why this approach doesn't work. I also work through the problems facing two other approaches, a behaviorist approach and a Dretske-inspired theory about reliable co-activation between perceived objects and neural patterns.

The failures of these approaches are instructive as a set of avoidance markers. It'll turn out that the 'standard theory' is precisely wrong. It gets it backwards. The infant does not start out with objects in its perceptual world with which it then learns to interact. *It interacts with what at first appears to be a completely chaotic world and, as a result of those interactions, eventually comes to understand the world as consisting of objects*. On the ‘standard’ radically objective account things are straightforward. Our *experienced* objects—that is, the objects that we experience as existing in the world—are, rather simply, just those objects that exist, objectively, in the real world. This seems plausible. But in the way it’s typically formulated, it is wrong, and it’s wrong in two ways. First, it’s wrong insofar as it implies that there is some kind of simple *copying* process by means of which we come to build up a perceptual world of objects. We do not start out by simply making mental copies of real objects which have already been given to us “fully-formed” from the external world.

Second, a radically objective account gets it wrong insofar as it preemptively rules out the contribution our cognitive system makes to our understanding of any ‘objective’ object. That our cognitive system makes a contribution to how we understand an object should be obvious. We can see this by first acknowledging that how we experience an object is limited to which features of it our sensory systems can detect. Migratory birds can perceive geo-magnetic flux lines, bats perceive sonar, snakes see infrared light, and bees can see into the ultraviolet range. You and I can’t. I will never *experience* an object as oriented along the geo-magnetic current as a goose does or sonar-sense a mosquito as a bat can. So, the objects that we experience are at least to a minimal degree subjective in the banal sense that without an advanced cognitive system in place that is tuned to certain ranges of input we would never have the opportunity to detect anything. What character and kind of object we experience depends on what kinds of features we’re capable of detecting; it depends on abilities that we bring to the table. An object, e.g., a mosquito, cannot be delivered fully-formed into my experienced world and the bat’s experienced world as type-identical experiences. The difference in the character of the experience of the mosquito for me vs. the bat illustrates that *the experience* of an object is not a universal objective phenomenon that makes itself available to just any organism. Therefore, a radically objective account that fails to acknowledge the cognitive contributions of the organism, contributions that in part determine how an object is experienced, is an account that’s missing an important piece of the puzzle.

The objectivists are of course correct in saying that, *from our third person perspective*, infants *are* born into a world of objects. Equally as obvious is that they are *not* born into such a world, *from the perspective of the infants*. And, it should also be just as obvious (though some of the more radical relativists state otherwise) that objects are not *merely* subjective manifestations. I can conjure up an internal image of a 4-headed, fire-breathing dragon, but I cannot in the same way conjure up a *real* dragon, or a real bag full of gold, or a real new car. The world emits patterns of physical energy. Our cognitive system can detect and organize some of that energy. The correct approach navigates between a pure objectivism that puts all of the work on a strictly objective real world and a pure idealism that puts all of the work entirely on our cognitive system. Perceptual objects are born from a stable, external world plus a properly tuned and trained cognitive system. Any talk of ‘manifesting’ an object can only be made under the caveat that such manifestation requires a world and mind and that such a world is independent from us in the sense of being outside of our direct control.

We live in a world of objects by the graces of our cognitive systems which, constrained by our biology, conspire to enable us to discern objects, under proper conditions of stimulation, for the benefit of our successful navigation in the environment. Evolution has predisposed us to learn to detect some objects more readily than others—apples over neutrinos, potential mates over imaginary numbers. Objects are joint constructions between the world and our cognitive system, important constructions. There are those, certainly, who will even dispute the claim that experienced objects are subjective in the minimal sense that the subject's cognitive system partly determines the object. But I take it to be self-evident, an updated variant, in broad outline, of Kant's basic position that our experience of the world is in part determined by our contribution.

How does our cognitive system contribute? The exploration into the experience of objects began with the work of the gestalt psychologists who, reacting to the stimulus-response model of the behaviorists and what was seen as the atomistic approach of Wundt, were impressed by the sum-greater-than-the-parts effects, standardly recounted in examples like that of the implicit triangle below:

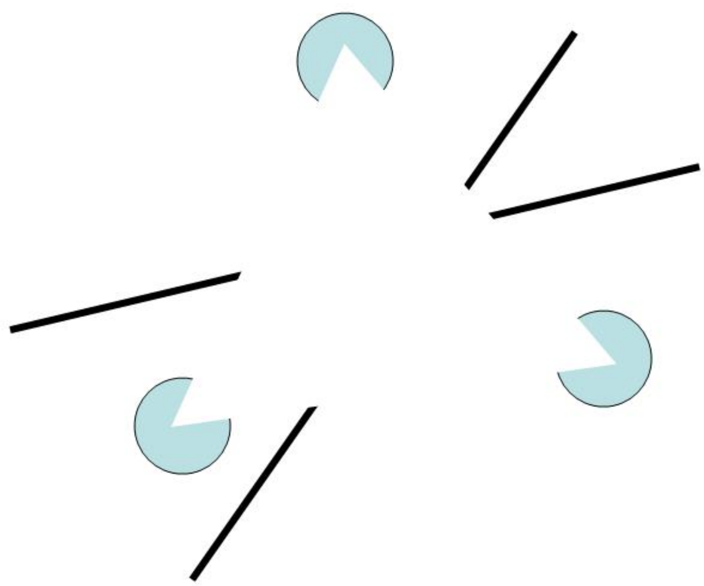

We perceive an equilateral triangle in the illustration despite the lack of lines to demarcate one.

Although the early 20th century gestalt psychologists like Cal Stumpf, Kurt Koffka, Max Wertheimer, and Wolfgang Köhler organized their work around the tenet of the unified character of experience and the imposition of our cognitive system on relatively impoverished sensory input to give structure to that input, their guidelines are often reinterpreted in contemporary cognitive science as sketches for algorithms that our central cognitive processor executes, like software applications executed on the computer of our central nervous system. Our 'cognitive computer', which just happens to be biologically instantiated, takes sensory input, runs calculations over that input, and then provides output.

Input → Processing → Output

The input → process → output model is a variant and elaboration of the stimulus-response model pursued by the behaviorists. They wanted objective, observable and rigorously measurable experimental variables, which could only be acquired in the input and output stages, and so they treated the middle cognitive processing stage as a black box. The computational approach, which took off just as behaviorism started to wane, reinforced the stimulus-response model by reinterpreting it as computational processing. Stimulus was interpreted as input. The black box was elevated to the level of an important computational processing stage (which provided a boon of research freedom for the computer scientists compared to the black box prohibition years of behaviorism). And response was reinterpreted as the output stage.

It's no surprise, then, that the role of our cognitive system in the development of our ability to discern objects has been downplayed. The stimulus-response model, understood as cause and effect, sets us up to be on par with the rest of the natural world. Like a billiard ball that doesn't move from rest unless it is acted upon, or a patch of earth that only makes a 'thud' if it's hit with a rock, we don't manifest a response unless some perceptual input sets into motion within us a causal chain that leads, in the typical case, eventually to motor output.

This is all wrong, I claim. We don't passively receive pre-formed objects, hand-delivered to us from the external world. We don't make simple copies of those objects in the form of representations. Cognition is not a stimulus-response machine. We do not merely react to external inputs. When we do react like an input-output machine, during a knee-jerk test, for example, it's because our cognition is the *least* involved it can be, the impact of the doctor's mallet triggering a nerve activation loop that runs from knee to spine and back down to the motor nerves in the leg, leaving out the brain until after the fact.

II. An Analysis of Objects

If an object cannot be understood via the stimulus-response model how are we to understand an experienced object? The question is not, 'what is an object?', but rather, 'what is an object *for an organism*?'.

Let's take a paradigmatic object, like the tennis ball. It's perceivable, lying smack in the middle of our human-specific sensitivity range for both vision and touch. We can easily manipulate it. We can grasp it, squeeze it, move it back and forth, hit it across a net with a racquet, we can even bite it. It's part of what we understand a tennis ball to be that it affords a gamut of such interactions. It doesn't exhaust our perceptual experience of the tennis ball but it provides structure to what would otherwise be a raw flood of phenomenology. Though if we did not receive 'raw stimuli' from the tennis ball we would never come to have a tennis ball as an object in our world.

But there is more than simple raw stimuli that is necessary for the manifestation of objects. Imagine that there was, at all times since birth, the stimuli of a spot in your visual field that never varied. It was always in the same position in your visual field. It never changed, and you could in no way control your experience of it. The question is: would you notice it? Well, there is such a spot in each eye where the optic nerve passes out the back of the eye, the blind spot. And we don't notice it. It's the constancy of it, the lack of control we have over it, the ever-present permanence of it that brings about its invisibility as an object of our experience. In psychology, sensory habituation is a well-known phenomenon that can dull even the most glaring and obnoxious dynamic stimuli.

When I was a child my grandmother would travel from a noisy neighborhood in Brooklyn and visit us in the countryside. Whenever she visited I couldn't sleep because she would take a radio, set it between stations at high volume so that it generated white noise not entirely unlike the blended cacophony of the street traffic heard from her high apartment window, and place it next to her head on her bed before she went to sleep. She had developed sensory habituation to the Brooklyn traffic and now needed the noise as part of her expected baseline auditory experience, even to fall asleep. I, of course, was used to nothing but the occasional cricket, and couldn't sleep at all in such noise from down the hall.

But my grandmother, if her attention was drawn to it, could pick out the noise in the field of her experience. By contrast, without very carefully designed techniques, we can't pick out the blind spot in our visual field. Why not? What's the difference? The difference is that the spot can't be controlled. But the perceived noise can be controlled, altered, modified. It can't be completely controlled (though we wish it could be), but it exhibits certain patterns of expected variability in the face of other sensory modifications. If I have the sensation of placing a pillow on either side of my head, then I expect, and indeed I'm rewarded with, the reduction in the volume of the noise. Tilting my head, moving toward the noise, covering my ears, placing a large obstacle between me and the source all affect the noise and what's more, do so in ways that I can anticipate, control, and thereby understand.

The first step to getting a handle on the manifestation of objects in our experience is to acknowledge that for there to be an object the raw stimuli from the object must make itself available within the detectable range of one of our sensory modalities and it must not be invariant like the blind spot on our eyes or the purported sound of the Earth spinning on its axis.

The second step is to realize that our engagement with the world is first and foremost one of interactions, and that those interactions eventually lead to sensorimotor regularities and patterns of engagement with the world.

What is it that leads one from raw, random uncoordinated bumblings in the world with crude parameters of control to the stage when one possesses refined object-utilizing skills that can be dispatched at will? Different perceptual inputs avail themselves to different ways of controlling them. Take the following scenario: Imagine you are snorkeling when, unexpectedly and suddenly, the light is blocked by some kind of diffuse mass. Is it a play of the light in the water or a genuine object, or multiple objects? You explore it. You swim around it, you reach out to it. You move, it moves. As you engage it in different ways and from different angles you gather an understanding of it. If you had no interaction with it, if you did not, e.g., see it, it would not exist in your cognitive landscape. But your slight movements cause slight alterations in your perceptions of the object. Your large movements cause large alterations in your perception of the object. Because the interactions co-vary with your perceptions in ways that you come to predict, and hence control, there is an object there for you, one whose perceptual alterations you can utilize.

This simple example illustrates the nub solution to the problem of how objects manifest in our world. It's illustrative because in it we can begin to see what must take place for an organism to move from a world barren of objects to its mature state, a condition where patterns of interaction are routinely executed under its control. This is the third step. With enough engagement in the world, the organism manifests habituated patterns of interaction. These patterns of engagement are the fundamental cognitive structures. These set up *expectations* of sensations such that a given pattern is expected to follow a prior pattern. If an expectation is violated in an incomprehensible way there is merely confusion. But if it is violated in a way that can be modified by stages, if it can be replicated and controlled, you have a new sensory regularity that can, to some degree or other, be initiated by organism. If the sensory regularity is unlike our sensation of the blind spot and more like the displacement of light from an underwater mass or the multi-modal input from a tennis ball, if it is in any way a *controllable* sensation, then you have a reliable set of patterns of interaction. If you have a patterns of interaction that operate over sensory input where there is a well-defined sensory boundary, you have an object. If it falls in the 'sweet spot' range of highest perceptual sensitivity across both vision and touch (the dogmatic duo of perception) with a well-defined boundary, you have a typical medium-sized physical object. The answer then, at last, in a mouthful, is this: *an object is a pattern within a field of stimuli and expected stimuli that is triangulated upon through a series of graded control-response interactions.* Of course, if the interactions are controlled by the organism then they are controlled *toward some end or purpose*. So, stated more completely:

> *An object is a pattern within a field of stimuli and expected stimuli that is triangulated upon through a series of graded control-response interactions initiated by the needs and desires of the organism.*

This will be spelled out in chapter 3. The 'expected stimuli' plays a greater role than one might think. The basic idea is that each pattern of interaction can be seen as an ingrained homeostatic pattern with the expected sensory input at any point in the duration of that pattern as the target stasis point which the system strives to achieve. When a smooth, continuous, pattern of interaction, like moving your hand through the air, is interrupted, especially when it's interrupted across more than one sensory modality, as when you suddenly feel resistance on your hand and at the same time see something new where your hand is, you eventually encode a discontinuity into your pattern of interaction. If enough overlapping discontinuities are encoded across enough patterns of interaction, an object, whose boundaries take the "shape" of the boundaries of those sensorimotor control discontinuities, manifests itself.

Think of each interaction you've had 'across' an object as a tether and the more interactions you've had the more lines are tied to the object until it is tied down in the net of your cognition. Or think of a charcoal trace you make over an unknown object. With each swipe of the charcoal you gain another impression of the object behind the paper until with enough strokes you have a complete picture.

The pivotal notion of this interactive view of objects is 'control'. In fact, what some may find shocking is that this account implies the following (rather strong) claim.

> The Control Thesis: The degree to which an object exists for us is the degree to which we can control our perceptions of it.

The Control Thesis in turn has strong implications. First, there is no object for you if you have no interactions that tie you to it. What's more, there is no object for you if you cannot *control* your interactions with the object. Your interactions must be active, not passive. This will strike many as patently absurd—"there is a 'moon' but I can't *control* the moon, can I?". No, but you can control your perception of the moon but simply moving your eyes or your head. A growing collection of experiments show that our cognitive system begins to collapse rapidly when we lose our ability to actively engage the world. It's my hope to dispel the air of absurdity that initially surrounds the interactive theory of object development. Indeed, I hope to persuade you that the strength of the Control Thesis claim is ultimately a sign of its robustness rather than its vulnerability. It (or some variant of it) is, I think, the only viable explanation for the formation of objects in the experienced world of a developing organism.

One of the benefits of this interactive view of objects is that it pays quick dividends in the form of a robust theory of object semantics. An organism acquires an understanding of an object derivative from the interactions in which that object plays a role given its needs, desires, and values. Interactions in general, in the world, are infused with meaning by virtue of the fact that an organism finds some interactions desirable (eating when hungry, e.g.) and other interactions undesirable (getting eaten by a predator, e.g.). As Cisek (2005) says:

> Animals have physiological demands which inherently distinguish some input (in the sense of 'what the animal perceives as its current situation') as 'desirable', and other input as 'undesirable'. A full stomach is preferred over an empty one; a state of safety is preferred over the presence of an attacking predator. This distinction gives *motivation* to animal behaviour — actions are performed in order to approach desirable input and avoid undesirable input. It also gives *meaning* to their perceptions — some perceptions are cues describing favourable situations, others are warnings describing unfavourable ones which must be avoided. The search for desirable input imposes functional design requirements on nervous systems that are quite different from the functional design requirements for input-output devices such as computers. In this sense, computers make poor metaphors for brains. For computers *there is no notion of desirable input within the computing system*, and hence there is the riddle of meaning, a.k.a. the symbol grounding problem (134).
>
> Meaning comes long before symbols in both phylogeny (evolutionary history of a species) and ontogeny (developmental history of an individual). Animals interacted with their environment in meaningful ways millions of years before they started using symbols. Children learn to interact with their world well before they begin to label their perceptions. The invention of symbols, in both phylogeny and ontogeny, is merely an elaboration of existing mechanisms for behavioural control (135).
>
> These relationships between animals and their habitats may be considered precursors to meaning. They are properties of the environment which make adaptive control possible and which guide that control. They make it possible for an organism to establish a behavioural control loop which can be used to approach favourable situations and avoid unfavourable ones. Because these properties tend to come packaged along with semi-permanent physical objects, we can speak of the 'meaning' that these objects have to the organism in question. However, the crucial point is that the '*meaning of an object' is secondary to the meaning of the interactions which the object makes possible.* (136).

With one important clarification, this is right. Interactions are not infused with meaning *solely* because of a creature's desires or aversions to the interaction but also because the interaction itself is a stable, structured pattern with which the animal is familiar (directly or by similarity). Any point in the unfolding of the interaction it is has meaning in the sense that the creature forms expectations about how the interaction will play out. The interaction is a skill, a tool deployed to achieve certain ends. This approach bears two interesting similarities, one to later Wittgenstein's relentless exploration of meaning of linguistic expressions as *use* of those expressions. The other similarity is to Gibson's (1979) theory of affordances. An affordance is an objective feature of the environment that affords opportunity for action by a creature. A creature perceives these affordances and thereby perceives and understands the world as ripe with opportunities for action. One benefit I have that Gibson didn't is a quarter of a century's worth of intervening experiments and a burgeoning theory of 'simulation semantics' whose evidence (at higher levels of cognition) dovetails nicely with my lower-level account.

Concomitant with a readily available story about object semantics, another benefit of the interactive account of objects is its strength in the face of, and arguably its dissolution of, the frame problem (McCarthy & Hayes, 1969). The problem has been taken up by Fodor and Dreyfus as a serious challenge to artificial intelligence because it points to a deeper flaw: the problem of relevance. If a computer operates by manipulating meaningless syntactic elements how will it determine which changes are relevant and which are irrelevant? Relevance depends on understanding and, as Searle (1980) has argued, syntactic transformation rules are no substitute for genuine understanding. The gulf between syntax and the distant continent of semantics, where relevance is evaluated, is the Bermuda Triangle of artificial intelligence that has silently swallowed otherwise powerful, brilliantly-crewed AI research vessels.

Per the interactive account, objects are phenomenological targets caught in the overlapping crosshairs of interactions. Every object is inherently meaningful to an agent by virtue of being captured in a net of meaningful interactions controlled by the agent. This is 'meaning' in the limited sense of 'viable possibilities of control'. Any controlled interaction is driven by some collection of needs, values, goals, and habits of the agent, and these offer an additional dimension of meaning above and beyond mere possibilities of physical engagement. A change in an object affects a change in the possibilities of control. And a change in the possibilities of control that are driven by needs and values of the organism will influence the success of satisfying those needs and values. If the likelihood of the success of satisfying those needs or values is strongly affected then the change is relevant. If not, then the change is irrelevant. Relevance is determined ultimately by needs and values which motivate an organism's behavior.

The interactive theory of object development is related to various strands of 'embodied' approaches to cognition and control theory, like the embodied theory of perception developed by Varela, Thompson, and Rosch (1991), and continued by Noë (2001a,b)--

though its roots go back to Heidegger, Dewey (1896), Merleau-Ponty (1945), Dreyfus (1979) and (1992), Lakoff and Johnson (1999), Johnson (1987), Feldman (2006), and Searle (1992) (who does not explicitly ally himself with the embodied movement). One common line of similarity among these approaches is that they all reject the classic computational and functionalist assumption that the body is an incidental machine on which the software cognition programs happen to run (This holds true for Feldman's computational account as well, though he may not realize it). On the embodied approaches the body is not incidental but a key feature that structures our cognition.

To take stock, I make the following claims in this project:

- Our engagement with the world consists first and foremost of interactions with it (not computations over representations of it).

- Objects (basic-level physical objects, that is, as perceived by an organism) develop as higher-order patterns over interactions in the world.

- Interactions are guided, that is, controlled, and control is essentially the manipulation of our own perceptions (not the manipulation of our objective externally-viewed behavior).

- An object is a pattern within a field of stimuli and expected stimuli that is triangulated upon through a series of graded control-response interactions initiated by the needs and desires of the organism.

- The degree to which an object exists for us is the degree to which we can control our perceptions of it. This is the Control Thesis.

III. An Analysis of Interactive Behavior

On this account, objects are made possible by a creature's behavioral interactions in the world. A complete account, then, must be grounded in a general account of interactive behavior. Living creatures are unique in the universe in that they proactively make use of the resources in their environment. Inanimate objects do not. They do not *use* anything; they only manifest the effects of causal impingement. Geologists expect the island of La Palma, one of the westernmost Canary Islands, to lose one of its flanks due to deformational instability, probably during the next volcanic eruption, and for that huge flank to slide into the Atlantic causing a 1 to 2 mile high mega-tsunami which will travel across the ocean at ferocious speeds until hitting the east coast somewhere on North or South America with horrific devastation. That mega-tsunami will be the simple causal result of dropping a few billion tons of rock and earth into the sea from a good height all at once. It is a straightforward natural *reaction*.

Living creatures are different. Not because we have special exemption from the natural laws of the universe. We can't escape the causal web any more than a rock can, but it's also true that *simple* cause and effect models in the form of input-output style reaction psychology models are very limited in application to us. Why? *Because we, like all living things, actively control our environment to our own benefit*. A tree brings nutrients up through its roots. A beaver builds a dam for a better home. But a rock just 'sits' there . . . until acted upon. Without much thought, and caught in the turbulence of constantly changing conditions, we keep ourselves upright during the day. We acquire enough food to maintain our blood glucose in a range of about 4 and 8 millimoles per liter. We blink enough to keep our eyes moist but not so much as to blind ourselves. We keep our cars between the lane lines. My dog drinks enough water to keep her cells functioning optimally but not so much as to kill herself from hyponatremia (the dilution of sodium to dangerously low levels). Our body keeps small variations in blood pH in check by means of releasing bicarbonate into our blood, and larger variations in check with the help of the kidneys, since a drop in our blood pH (hydrogen ion concentrations) down below 6.8 would kill us. Wolves grow dense coats in the cold winter months and lose fur in the summers to help stabilize their year-round temperature.

These are all examples of *homeostatic* systems. French physiologist, Claude Bernard discovered that although blood sugar could be raised or lowered it would inevitably gravitate back to a narrow range. Bernard referred to this as 'le milieu intérieur' and extended it to other components of the blood that he hypothesized were similarly regulated. It wasn't until 50 years later that Harvard physiologist W. B. Cannon revived the notion to explain the constancy of blood pressure and coined the term, 'homeostasis'. Homeostasis is the active regulation (control) of a biological parameter such that if that parameter is moved outside its target range the regulatory system will resist that change and recover back to target.

The advantages of controlling one's immediate environment by means of these homeostatic systems is enormous. It frees the creature from the oscillations of the external world and allows it to hold fixed, in the face of environmental variability, conditions it can rely on to more easily thrive in the world. As Bernard said (1966), "La fixité du milieu intérieur est la condition de la vie libre, independante." "Free, independent life" was achieved 3-4 billion years ago by the first life forms—single-celled microorganisms that enclosed themselves in a phospholipid (fatty) membrane which provided a clearly-demarcated boundary between that portion of the environment that could be controlled via their homeostatic systems and the rest of the world. That stable collection of conditions inside the membrane made for a dependable platform upon which the single-celled organism could more successfully thrive. "All the vital mechanisms, varied as they are, have only one object: that of preserving constant the conditions of life".

In general, the greater the number of homeostatic regulatory systems a creature has the more evolutionarily advanced the creature and the greater its ability to succeed in diverse conditions. Cold-blooded creatures (ectotherms) like snakes, don't have a homeostatic system for regulating temperature so they are slowed to the point of immobility in cold weather, while an endotherm like you and I can function pretty much the same as we do when we're in warmer climes. That's the advantage of a temperature regulating system that keeps our body temperature within a tight range of variability.

Bioscience has made a bit of progress since the time of Claude Bernard and W.B. Cannon. Now, *every* physiological system is commonly understood as a homeostasis system—the formation of our bones, the beating of our heart, the blinking of our eyes, the release of adrenaline, the acidity of our stomachs, the firing of our neurons, the rate of cell turnover, the population of digestive bacteria in our guts, the secretion of the amylase enzyme, the levels of our hormones, the concentration of CO2 in our blood, an so on (for homeostasis in neural structures c.f. E. Marder and A.A. Prinz, Modeling stability in neuron and network function: the role of activity in homeostasis, *Bioessays* 24, 2002, pp. 1145–1154).

Every physiological system is now understood as a system that 'strives' to maintain control over select biochemical parameters to keep them within a 'desired' range . . . with one rather notable exception. Can you guess what that exception is to the rule of life as a collection of interwoven homeostatic systems? It's the mind. But why should the mind be given the right of exceptional exclusion from the otherwise exceptionless rule of physiology? Clearly the brain, the extended nervous system, and its enormous variety of interrelated homeostatic systems can't be excluded from the natural order. So, unless one takes up the unenviable position of arguing for a division between the brain and mind so sharp as to leave the former in the domain of natural law and the latter outside it, then the mind, too, is a homeostatic system, or more accurately, a collection of homeostatic regulatory control systems.

> All physiological systems are homeostatic systems.
> The mind is a physiological system.
> Therefore, the mind is a homeostatic system, too.

Let's call this the Homeostasis Thesis: The mind, like all interesting physiological systems is a homeostatic system.

If the Homeostatic Thesis is right (and I think it is) then the two models that still hold sway to varying degrees in psychology and cognitive science are wrong—the stimulus-response model of behavior and the input→processing →output computation model. Put more charitably, these models are, at the very least, radically under-equipped in that they leave out the effect of the output on the input to close the feedback loop that comprises any regulatory control system. And they miss the most important structure, the reason for the very existence of any homeostatic system—the target parameter that the system exists to control. Dewey tried to warn of us against relying on a stimulus-response model of the mind more than 100 years ago in his paper, "The Reflex Arc Concept in Psychology" (1896):

> The discussion up to this point may be summarized by saying that the reflex arc idea, as commonly employed, is defective in that it assumes sensory stimulus and motor response as distinct psychical existences, while in reality they are always inside a coördination and have their significance purely from the part played in maintaining or reconstituting the coördination; and (secondly) in assuming that the quale of experience which precedes the 'motor' phase and that which succeeds it are two different states, instead of the last being always the first reconstituted, the motor phase coming in only for the sake of such mediation. The result is that the reflex arc idea leaves us with a disjointed psychology, whether viewed from the standpoint of development in the individual or in the race, or from that of the analysis of the mature consciousness. As to the former, in its failure to see that the arc of which it talks is virtually a circuit, a continual reconstitution, it breaks continuity and leaves us nothing but a series of jerks, the origin of each jerk to be sought outside the process of experience itself, in either an external pressure of 'environment,' or else in an unaccountable spontaneous variation from within the 'soul' or the 'organism.' As to the latter, failing to see unity of activity, no matter how much it may prate of unity, it still leaves us with sensation or peripheral stimulus; idea, or central process (the equivalent of attention); and motor response, or act, as three disconnected existences, having to be somehow adjusted to each other, whether through the intervention of an extraexperimental soul, or by mechanical push and pull.

Dewey points out that the stimulus-response model (the "reflex arc" model) of human psychology fails to understand perceptual motor skill for what it is. Rather than an arc, it's a closed *loop* where motor responses affect sensory stimulation which in turn affect motor responses in an ongoing perceptual-motor pattern or skill, a *coordination*. A coordination, pace Dewey, involves ongoing closed loop feedback control over a desired stasis target behavior. A behavior, if it unfolds in an unintended direction, is corrected and brought back into the desired pattern progression. A coordination, in other words, is what we, with greater familiarity, would refer to as a homeostatic system.

But is Dewey's point relevant 100 years later? The intervening period has seen not only the tremendous growth of psychology as a discipline apart from philosophy and seen the birth of cognitive science, it has seen the refinement of experimental methodologies. The bulk of these are reaction time response tests (though survey-question-based experiments and ability tests are used as well). Neuroscientists make heavy use of fMRI scans which rely on fact that oxygenated blood hemoglobin possesses a slightly negative charge and deoxygenated hemoglobin possesses a slightly positive charge, allowing them to see, by means of the magnetic field in an fMRI machine, which areas of the brain are active, relative to baseline activity. The typical fMRI study proceeds by measuring the subject's brain activity *as a reaction* to a motor, perceptual, or a cognitive task. EEG experiments measure brain frequencies through the skull, and single cell recordings of animals measure the firing of small clusters of neurons. Lesion analysis looks at the effects of natural lesions in humans, manufactured lesions in lower animals, and virtual lesions in humans and other species by disrupting brain activity temporarily with trans-cranial magnetic stimulation (TMS).

The bulk of all of these experiments in cognitive science and neuroscience *are still reaction-based*. And those that are not, like the ability tests, still don't have any way of teasing apart the principal features of homeostasis. These experiments have provided invaluable help and key empirical evidence, but if Dewey was pointing us in the right direction toward the Homeostasis Thesis, the experimental methods we've relied on haven't been methods that could elucidate much at all about the rampant homeostatic systems that constitute our mind. Our reaction-based experimental methods were simply not built to be able to detect features of homeostasis. And the classic computational model at the basis of cognitive science, the input→processing→output model, is the wrong model. Dewey would have told us as much.

Dewey, Heidegger, and Meleau-Ponty all emphasized the importance of a skill, or in Dewey's words, a 'coordination', i.e., a structured pattern of routinized perceptual-motor activity. Take the sensorimotor habit of putting on a sock. You grasp the sock. You insert a thumb from each hand into the sock. You slide your foot into the opening made by your thumbs, etc. As you proceed you have expectations of what the appropriate sensory feedback is that indicates that your task is on track. If you suddenly feel a loss of resistance against one of your hands or your leg kicking away from you then something is amiss (maybe the dog ran off with your sock, e.g.). The expected sensation at each stage of the process constitutes the variable that is regulated as part of this small homeostatic behavioral pattern, the 'put-the-sock-on' homeostatic system. Deviations from the expected sensations are regulated and brought back in line toward the expected sensations. They are controlled. All our interactions in the world are the discovery and execution of patterns of control as manifest in the control of our own sensations.

Standard computational and robotics approaches assume that the object of control is motor behavior. The simpler view advanced by Powers (1973) and (1998), Cisek (1999), and Bourbon (1995), under the rubrik of Powers' 'Perceptual Control Theory' is that the target of control is not behavior, but perception itself. When one is driving a car the goal is not to move the steering wheel two inches to the left then one and half inches to the right, etc. (although these movements happen as by-products). The goal is simply to keep the car between the lane lines, and this is what a driver typically thinks. That means that one must receive sensory input that is in line with keeping the car between the lane lines. If your car drifts outside the lane you act in such a way that you soon receive perceptual input that your car is between the lanes again. You act so as to control your perception, not a myriad of leg and arm movements.

Unfortunately, the overly-grand claims of paradigm-shift by those advocating Perceptual Control Theory have distanced many who are currently applying modern variants of Powers' model—people like Van Gelder, Kelso, Freeman, Thelen & Smith, and others that are using dynamical systems feedback control approaches to understand the mind. The important insight of Powers is the simple phenomenological point that the only means by which one can keep a behavior on track is through perception. The sensations received are evaluated relative to the desired sensations and behavior is adjusted accordingly. So, from the first-person perspective, it's perception that is controlled (not behavior as viewed from a third-person perspective). Actions are always enacted by agents and so the appropriate language for describing an action is that from the perspective of the agent, the description of their sensations.

Imagine you're an engineer on a team whose task is to design a humanoid robot that can ice skate. Your job is to build the location and balance systems. How do you get the robot to determine its position and its balance? Should you design it to detect absolute location on Earth at any moment and determine its balance by running complex calculations on how, when standing on only one skate, its total mass is distributed over, say, the 360 degrees of an imaginary circle that is projected onto the ice with the contact point of its skate as the center of the circle? In other words, should you build the robot's skill on top of objective third-person descriptions of its movements and positions? Or should you instead teach it where it is by having it look at how far away the walls of the ice rink appear—that they grow larger when closer and smaller when farther away—and teach it balance by building delicate sensors into the sole of its foot that detect subtle changes in the pressure distribution inside its skate? In other words, should you use 'first-person sensation input' that is prevalent and readily available? Well, input that the robot receives is the only input available. There's no other input to choose from. The real question is whether, based on that input, it will perform slow, complex, calculations that derive difficult-to-determine external third-person descriptions of its behavior or whether it will do simpler calculations that are easily derived and rapidly updatable from changing 'sensations'. Clearly, building third-person descriptions is an unnecessary computational burden. The more transformations you need to make to the sensory data usable to the organism the more complex the system will need to be.

It's easy to see its importance of first-person-based control systems. The subtler distinction often missed in Powers' suggestion is about how action should be initiated. What is the object of control? What should your robot control in order to initiate a jump on the skating rink? You have two choices. You can build a control system that activates its servo-motors directly in very complex ways at very low levels of detail. Calculating the detailed activation patterns of its motors is incredibly difficult. Or you can have the robot simply 'bring about' that stream of sensory input that is already bound up in the behaviors it's learned that you want it to perform. Having the robot increase the pressure on the sole of its left foot forces it to shift more weight onto that foot. Asking it for a strong-enough spike of pressure on that foot is also thereby asking it for a jump. Instead of forcing huge computational burdens on your robot's processor, asking it to calculate enormously complex sets of low-level detailed activation sequences you can rely on the existence of simple intertwined 'sensation' plus motor patterns—*coordinations*. If these are in place, a creature can't help but act in certain ways in order to bring about the sensation sequence. Try engendering *the sensation sequence* of hitting a tennis serve without really hitting an actual tennis serve. Or try engendering *the sensation sequence* of riding your bicycle without actually riding your bicycle. It's very hard to get the right *feel*, the right sensorimotor feedback without actually engaging the activity that typically provides that feedback. That's why flight simulators at NASA are so expensive.

The distinction is between initiating action via direct computationally-taxing third-person descriptions versus initiating action only indirectly by bringing about the relevant sensation sequence (which carries with it associated actions). The approach here is reductive. Motor initiation is best approached as sensory initiation. This simplifies the system and shifts the emphasis from a discrete serial system of:

sensation → processing → motor output

to one built from the ground up on entangled sensorimotor patterns. This will be discussed in more detail in chapter 4.

Each sensorimotor pattern is a small homeostasis system whose target pattern is the pattern of expected sensory input. A developing infant hasn't yet accumulated a repertoire of patterns of interaction with the world like these that we take for granted, and hasn't yet carved the world into objects. All it can do is erupt in undirected bursts and take the merciless salvos of sensory onslaughts. Even a newborn, though, will display generalized behavior indicative of perceptual input that is either outside or inside of a desired target equilibrium zone. E-coli bacteria initiate spastic, chaotic, random "tumbling" behavior when there is no food. This has the advantage of increasing the probability that they will find food simply because they are moving away from a place that definitively doesn't have food. This increases the chances of settling back into stasis (a food-rich location). Not dissimilarly, infants will flail and scream until their general evaluation system detects that they are back near equilibrium. This is one of the only means they have for moving back to stasis (which relies of the evolutionary fact that their parents will keep trying things until the screaming stops). With enough practice, enough trial and error, those bursts of will mature into genuine interactions, i.e., skilled engagements with the world that are part of a hierarchy of other skills refined to take control over regularities in the world. Instead of screaming and flailing the infant will eventually exhibit control and move itself away from disequilibrium and ask for food when it's hungry or grab a blanket when the source of disequilibrium is the fact that it's cold. Living organisms are homeostatic systems that are constantly struggling to regain homeostasis against the destabilizing forces of the external world and their own changing needs and drives.

What distinguishes psychological homeostasis from standard non-psychological homeostasis, like the blood sugar regulation system, is that standard homeostasis systems are permanent fixtures of our physiology whereas psychological homeostasis patterns are brought out temporarily for the service of some end, so their existence is intermittent, and they are constantly disrupted and replaced. I am not continuously putting on a sock throughout my life; I put on a pair each morning and, if I'm vacationing in the tropics, my sandals instead. Perhaps it's the dynamic, intermittent nature of skills like these, embedded in the nervous system, that leaves physiologists ambivalent about their status as genuine homeostasis systems. One standard physiology text says, "[The nervous system] is especially important in detecting and initiating reactions to changes in the external environment. Furthermore, it is responsible for higher functions that are not entirely directed toward maintaining homeostasis, such as consciousness, memory, and creativity" (Sherwood, 9).

Notice the implicit reliance on a stimulus-reaction psychology model in the phrase, "initiating reactions". My claim is *precisely* the opposite: the mind, like every other biological system, *is* a homeostatic system through and through, in at least four senses:

(a) The neural structures are homeostatic systems.
(b) The perceptual invariants are homeostatic.
(c) The sensorimotor contingency patterns are homeostatic.
(d) The behavior of the organism itself is homeostatic

First, neural activity consists of dynamic homeostatic sub-systems. 'Dynamic' here might be thought to be in conflict with the claim of homeostatic equilibrium, but if we understand a 'dynamic homeostatic system' as referring to a stable pattern that repeats itself, and so 'dynamic' in the sense of intermittent (albeit reliable), then there's no real conflict. This first point is supported by an increasing number of studies show that neuronal activity 'strives' toward a target level of activation and that if that activity is experimentally hampered it returns to baseline within a few days (Rich and Wenner, 2007). It does so in surprisingly malleable ways—increasing the number of postsynaptic receptors and increasing the payload of neurotransmitters to compensate for the surgical removal of neurons, for example.

Second, the mind is a homeostasis system in the sense that perceptual experience is constrained by simple perceptual invariance relationships (also known as 'perceptual constancies'). The best known is size constancy which states that the size of an object remains constant despite appearing larger as it moves closer and smaller as it moves farther away.

stimulus size = actual size / distance

We interpret our experience such that our understanding of it conforms to this constancy. We don't think that the train magically grows in size as it approaches us on the platform, even though the size of the image on our retina is increasing, because we implicitly understand the world around us according to such relationships—the train is just getting closer. The stasis target of this perceptual invariant as a homeostatic system is that set of ways of interpreting our experience that coheres with it. My interpretations are limited. I can see the train as growing, but then it would not be getting closer. But I can't see the train as moving away and staying the same size with it's image on my retina is enlarging.

The invariants are quite powerful and out of one's conscious control. Interpretations of experience out of accord with the perceptual constancies are rejected and quickly re-interpreted in line with them, even if, by artificial intervention as when for example one puts on a pair of inversion goggles, they no longer help us function in the world, and, instead, impair our functioning. The sense of disorientation when perceptual regularities fail us is overwhelming and nauseating.

Third, the mind is a homeostatic system in the sense that it is a collection of hierarchically organized temporally extended patterns of interaction with the world (like the put-on-the-sock motor skill), and almost all of which are, on the one hand, part of larger habits of interactions (like the put-on-my-clothes activity), and on the other hand, decomposable into smaller patterns (like reaching and grasping skills). Each activity is homeostatic because in each case perceptual input is maintained within a narrow range of the target perceptual pattern.

But what patterns of interaction, from amid the myriad, are triggered at any given point? How does the organism decide which 'coordinations' to activate? This brings up the fourth sense in which the mind is a homeostatic system. It is a homeostatic system *in its entirety*. If so, the natural question is: What, then, is the overarching stasis target of this homeostasis system, of the mind itself? Over what is it striving to maintain balance? This question asks no less than what the ultimate on-going purpose of the mind is. The answer to this grand question that I shall try to make persuasive, is that *an organism always strives to maintain 'goodness-of-fit' between its needs, values, and desires on the one hand, the bulk of which are unconscious, and its ongoing conscious experience of the world on the other hand*. The stasis target then is the pulsing, shifting structured aggregate of needs/values. Behaviors are chosen that satisfy the most prominent needs/values at any given time. This is, incidentally, the reason why those who are hampered by emotional instability often have less success achieving long-term goals than those whose values and desires vary little. Deep restructuring of values (often a trait of those who are emotionally unstable) leads to a more widely oscillating target for goodness-of-fit, a less focused engagement with the world that makes it difficult to make consistent progress on any one value-specific goal. Emotional variability engenders oscillating values which makes for shifting goodness-of-fit targets.

What purpose does it serve to analyze the mind as a homeostatic system? What work does it do? Well, it's a clean way of providing a unified account of the mind. It nicely bridges psychology with physiology in that the mind, like the rest of our physiological systems, can be understood according to the same basic principle of homeostasis. It inclines us toward a naturalistic approach. The stable patterns of neural firing are reflected in the stable sensorimotor invariants by means of which we 'interpret' our sensory stimulation, stable patterns of interaction by means of which we understand and engage the world, and stable needs and values that an organism strives to satisfy. A homeostatic approach is more accurate than an atemporal computational approach in that a homeostatic system is, by definition, a temporally extended system. It is, arguably, a much better way to understand human behavior—not as a reactive input-output system, but as an activity that strives to bring into existence the manifestation of needs and desires of the organism. In short, the Homeostasis Thesis is a powerful, comprehensive approach to the mind that offers the possibility of radically streamlining diverse areas of research and explaining why living organisms are successful at the kinds of things that computer programs and artificial agents have proven so unsuccessful at doing.

In summary, I make the following claims in this project:

- Our engagement with the world consists first and foremost of interactions with it (not computations over representations of it).

- Objects (basic-level physical objects, that is, as perceived by an organism) develop as higher-order patterns over interactions in the world.

- Interactions are guided, that is, controlled, and control is essentially the manipulation of our own perceptions (not the manipulation of our objective externally-viewed behavior).

- An object is a pattern within a field of stimuli and expected stimuli that is triangulated upon through a series of graded control-response interactions initiated by the needs and desires of the organism.

- The degree to which an object exists for us is the degree to which we can control our perceptions of it. This is the Control Thesis.

The Homeostasis Thesis says the following:

> The mind, like every other biological system, is a homeostasis system. It's a homeostatic system in four ways.
> (a) The neural structures are homeostatic systems.
> (b) The perceptual invariants are homeostatic.
> (c) The sensorimotor contingency patterns are homeostatic.
> (d) The behavior of the organism itself is homeostatic

Chapter 2 explores potential answers to the question of how objects come to exist in our perceptual world and ultimately rejects the solutions explored there. That work, however, provides clues to the correct account, an interactive account of objects, which is investigated in chapter 3.

Of the four points above, point (a) is already well-established in neuroscience and shall not be discussed. Chapter 3 investigates points (b) and (c) in the course of advancing the interactive account of objects. And in chapter 4, we examine the last point, (d), the claim that the fundamental purpose of the mind is to serve as a homeostatic regulator for our surging, shifting collection of drives. This is the essence of the Homeostasis Thesis, the key to understanding human behavior, and the structuring principle behind all of our activity while we are alive. There I advance the following:

1) The behavior of a living organism is not random or merely a collection of stimulus-driven responses, but by and large directed under its own control.

2) An organism exerts control through its behavior in order to satisfy its needs/desires/values. Controlled interactions in the world are guided by the relentless imposition of an organism's changing needs/desires/values onto its environment.

3) This control system is best approached the way any physiological control system is—as a homeostatic system, albeit a dynamic one.

4) That homeostatic system of organism behavior has the following features:
   - A regulated parameter
   - A target goal state
   - A correction system
   - A large collection of underlying regularities that are relied upon
   - An evaluation system

As a homeostatic system, an organism strives to achieve equilibrium. The ever-present equilibrium target of the organism is the dynamic, shifting subset of perceptual patterns that are activated by the organism's needs/desires/values. That which the organism strives to regulate, to bring in line with its values, are its conscious focal experiences.

A Note about Style

I see the broad, foundational strokes painted by a naturalistic philosopher as an activity that has a characteristic freedom—the freedom to attempt to provide a unifying theoretical context for disparate areas of empirical research without the burden of doing that research. But this armchair license has a price. I see this freedom as demanding concomitant responsibilities: if you are going to do broad, foundational, work a la a philosopher, then you incur the responsibility of constraining your reasoning by the principles of induction and empirical evidence rather than, say, speculative reasoning guided only by first-order logic. You have a responsibility to present an in-principle falsifiable theory. And, most importantly, you bear the burden of pursing unifying theories that expand the number of theoretical options for everyone, theories that are currently too brazen for the careful detail-oriented experimental researcher to in good conscience defend, theories that are too empirical for the analytic philosopher to defend, theories that must be advanced so the researchers can begin thinking about attempting to falsify them and the philosophers can begin thinking about their logical implications. The price is one of taking on the role of putting creative solutions on the table for the rest of the community to evaluate using the specialty tools of their respective specialized sub-domains. Where the current work succeeds as a piece of naturalistic philosophy it succeeds because it has taken these responsibilities seriously. Where it fails it's because it has failed to take them seriously enough.

References

Alerstam, T. (1990). *Bird Migration*. New York: Cambridge University Press.

Baars, B. (1998). *A Cognitive Theory of Consciousness*. New York: Cambridge University Press.

Baillargeon, R. (1987). "Object Permanence in 3.5- and 4.5-month-old infants", vol. 23:163, 665-664.

Bahrick, L. (1988). Intermodal learning in infancy: learning on the basis of two kinds of amodal invariant relations in audible and visible events. *Child Development*, 59, 197-209.

Bahrick, L. & Pickens J.N. (1994). "Amodal Relations: The Basis for Intermodal Perception and Learning in Infancy", in *The Development of Intersensory Perception: Comparitive Perspectives*, 205-33. Hillsdale, N.J.: Lawrence Erlbaum Associates.

Bakken, G.S. & Krochmal, A.R. (2007). "The Imaging Properties and Sensitivity of the Facial Pit of Pitvipers as Determined by Optical and Heat-Transfer Analysis." *Journal of Experimental Biology*. 210: 2801-2810.

Balch, W.R., Myers, D.M., & Papotto, C. (1999). Dimensions of mood in modd-dependent memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 70-83.

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

Becker, Suzanna, & Hinton, Geoffrey E. (1992). Self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature,* 355, 161-163.

Bergen, B. & Feldman, J. A. (2006). It's the body, stupid: Concept learning according to cognitive science. *International Computer Science Institute,* Technical Report, ICSI TR-06-002.

Bergen, B., Matlock, T., & Narayan, S. (In Prep). Linguistic cues for perceptual simulation.

Berkeley George. (Howard Robinson, ed.) (1710/1999). *Principle of Human Knowledge & Three Dialogues.* Oxford University Press.

Bernard, C. (1966). *Leçons sur les phénomenes de la vie communs aux animaux et aux vegetaux*. T. II. Paris: Bailliére.

Bloom, P. (2000). How children learn the meanings of words. Cambridge, MA: MIT Press.

Bourbon, W.T. (1995). "Perceptual Control Theory", in *Comparitive Approaches to Cognitive Science*, ed. H.L. Roitblat and J.A. Meyer (Cambridge: MIT Press).

Boles, L. & Lohmann, K.J. (2003). "True Navigation and Magnetic Maps in Spiny Lobsters". *Nature*, 421, 60-63.

Bower, G.H. (1981). Mood and memory. *American Psychology*, 36:2, 129-148.

Bourbon, W. T., (1995. Perceptual Control Theory. *Comparative Approaches to Cognitive Science: Natural and Artificial.* (Roitblat H.L. & Meyer, J.-A., eds) Cambridge, MA: MIT Press.

Bremner, J.G. (1978). Egocentric versus allocentric spatial coding in nine-month-old infants: Factors influencing the choice of code. *Developmental Psychology* vol 69:34-355.

Bremner, J.G. & Bryant, P.E. (1977). Place versus response as the basis of spatial errors made by young infants. *Journal of Experimental Child Psychology* vol 23:162-177.

Broaks, J. (1992). The autonomy of colour. In K. Lennon & D charles (Eds.), *Reduction, Explanation, an Realism* (pp 421-465). Oxford: Oxford University Press.

Brooks, R. (1991). Intelligence without reason. *Proceedings of the 12th International Joint Conference on Artificial Intelligence*. Sydney, Australia, Aug, 1991, pp. 569-595.

Brooks, R. (1992). Intelligence without representation. *Foundations of Artificial Intelligence*, vol. 47, 139-159.

Campbell, John (2002). *Reference and Consciousness*. Oxford University Press.

Cherry, E. C. (1953). Some experiments on the recognition of speech, with one and with two ears. *Journal of the Acoustical Society of America,* 25, 975-9.

Cisek, P. (1999). Beyond the computer metaphor: Behaviour as interaction. *Journal of Consciousness Studies.* 6(11-12): 125-142.

Conway, F. & Siegelman (2006). *Dark Hero of the Information Age: In Search of Norbert Wiener, the Father of Cybernetics*. Basic Books.

Cooper, L. A., & Shepard, R. N. (1973). The time required to prepare for a rotated stimulus. *Memory & Cognition*, 1, 246-250.

Cooper, L. A., & Shepard, R. N. (1978). Transformations on representations of objects in space. In E. C. Carterette & M. P. Fiedman (Eds.), *Handbook of Perception (Vol. 8)*. New York: Academic Press.

Copi, I. & Cohen, C. (2008). *Introduction to Logic*. Inglewood, N.J.: Prentice Hall

Crane, T. (1992) (ed.) *The Contents of Experience, Essays on Perception*, Cambridge: Cambridge University Press.

Crane, T. (1998). Intentionality as the mark of the mental, in O'Hear, A. (ed.). (1998) *Contemporary Issues in the Philosophy of Mind.* Cambridge: Cambridge University Press.

Cruse, D.A. (1977). The pragmatics of lexical specificity. *Journal of Linguistics* 13:153-64.

Curtis, G.C. & Zuckermand, M. (1968). A Psychopathological Reaction Precipitated by Sensory Deprivation. *American Journal of Psychiatry*, 125, pp 255-260.

Dawkins, R. (1976). *The Selfish Gene*. Oxford: Oxford University Press.

Dewey, John (1896). The reflex arc concept in psychology, *Psychological Review*, 3, pp. 357-70.

Diamond, A. (1991). Neuropsychological insights into the meaning of object concept development. In S. Carey & R Gelman (Eds.), *The Epigeneis of Mind*. 67-110. Hillsdale, N.J.: Erlbaum Associates.

Dolezal, H. (1982). *Living in a World Transformed*. New York: Academic Press.

Dretske, Fred (1981). *Knowledge and the Flow of Information*, Cambridge, Mass.: MIT Press, a Bradford book.

Dreyfus, H.L. (1979). *What Computers Can't Do*. New York: Harper & Row.

Dreyfus, H.L. (1992). *What Computers Still Can't Do*. New York: Harper & Row.

Dreyfus, H.L. (2007). Why Heideggerian AI failed and how fixing it would require making it more Heideggerian, *Philosophical Psychology*. 20:2, 247-268.

Eich, E., & Metcalfe, J. (1989). Mood dependent memory for internal versus external events. *Journal of Experimental Psychology: Learning, Memory, an Cognition*, 15, 443-455.

Eich, E., Maculey, D., & Ryan, L. (1994). Mood dependent memory for events of the personal past. *Journal of Experimental Psychology: General*, 123, 201-215.

Ehrsson, H., Geyer, S. & Naito, E. (2003). Imagery of voluntary movement of fingers, toes, and tongue activates corresonding body-part specific motor representations. *Journal of Neurophysiology*, 90 pp. 3304-3316.

Evans, Gareth. (1985). *Collected Papers*. Oxford: Clarendon Press.

Ewing, Alfred Cyril. (1969). Kant's Treatment of Causality. Oxford, England: Archon Books.

Feldman, J. (2006). *From Molecule to Metaphor*. Cambridge, MA: MIT Press.

Fodor, J.A. (1981). *RePresentations*. Cambridge, MA: MIT Press

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

Fodor, J. 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford University Press.

Freedman, David A. (1971). Congenital and Perinatal Sensory Deprivation: Some Studies in Early Development. *American Journal of Psychiatry*: 127: 1539-1545.

Gale Encyclopedia of Psychology, 2nd ed. Gale Group, 2001. Entry: Sensory Deprivation.

Gallagher, Shaun. (2005). *How the Body Shapes the Mind*. Oxford: Clarendon Press.

Gallese, V. and Lakoff, G. (2005). The brain's concepts: The role of the sensory-motor system in reason and language. *Cognitive Neuropsychology*, Vol. 22, no. 3-4, pp.455-479.

Gibson, E.J. (1969). *Principles of Perceptual Learning and Development.* New York: Appleton-Century-Crofts.

Gibson, James, J. (1986). *The Ecological Approach to Visual Perception*. Hilllsdale, N.J.: Lawrence Erlbaum Associates.

Godden, D.R. & Baddeley, A.D. (1975). Context-dependent memory in two natural environments: Land and underwater. *British Journal of Psychology*, 66, 325-331.

Hebb, D. O. (1970). "The Motivating Effects of Exteroceptive Stimulation." American Psychologist 25(4): 328-336.

Held, R., and A. Hein. (1963). Movement produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology* 56:873-876.

Herz, R.S. (1998). Emotion experienced during encoding enhances odor retrieval cue effectiveness. *American Journal of Psychology*, 110, 489-505.

Heyers, D., Manns, M., Luksch, H., Güntürkün, O. & Mouritsen, H. (2007). A visual pathway links brain structures active during magnetic compass orientation in migratory birds. *Public Library of Science One*. 2(9): e937.

Hinkle, L.E. & Wolff, H.G. (1956). Communist interrogation and indoctrinatino of 'enemies of the state': analysis of methods used by the communist state police (special report). *A.M.A. Archives of Neurology and Psychiatry*, Vol. 76, pp 115-174.

Hopkin, M. (2004). "Squirrels use Infrared against Snakes". *BioEd Online*. http://www.bioedonline.org/picks/news.cfm?art=1018

Horn, R., Yoshimi, J., Deering, M., & McBride, R. (1998). *Can Computers Think?* Bainbridge Island, WA: MacroVU Press.

Hume, David. (1772/1999). *An Enquiry Concerning Human Understanding.* Oxford: Oxford University Press.

Hunn, E.S. (1975). A measure of the degree of correspondence of folk to scientific biological classification. *American Ethnologist* 2, no. 2, pp. 309-27.

James, William (1890/1950). *The Principles of Psychology*. Dover Publications.

James, William (1904). Does Consciousness Exist? *Journal of Philosophy, Psychology, and Scientific Methods*, 1, 477-491.

James, William (1897/1984). The sentiment of rationality, in *The Essential Writings*. Edited by B. Wilshire. Albany: State University of New York Press.

Johnson, Mark (2007). *The Meaning of the Body: Aesthetics of Human Understanding*. Chicago: Chicago University Press.

Kant, Immanuel (1878/1999). *Critique of Pure Reason*. Paul Guyer and Allen W. Wood (Eds.). New York: Cambridge University Press.

Kauffman, S.A. (1993) *The Origins of Order: Self-Organization and Selection in Evolution* (New York: Oxford University Press).

Kellman, P.J. & Arterberry, M.E. (2000). *The Cradle of Knowledge: Development of Perception in Infancy*. MIT Press.

Kelly, S.D. (2001). The non-conceptual content of perceptual experience: situation dependence and fineness of grain. *Philosophy and Phenomenological Research*. 62:3, 601-608.

Klein, M. (1981). In L.T. Benjamin, Jr. & K.D. Lowman (Eds.), *Activities Handbook for the Teaching of Psychology.* Washington, D.C.: American Psychological Association.

Kohler, I. (1951). Über aufbau und wandlungen der wahrnehmungswelt. *Österreichische Akademie der Wissenshaften. Sitzungsberichte, philosophish-historische Klass*, 227, 1-118.

Kosslyn, S. M. (1973). Scanning visual images: Some structural implications. *Perception & Psychophysics*, 14, 90-94).

Kosslyn, S., Ball, T., & Reiser, B. (1978). Visual images preserve metric spatial information: Evidence from studies of image scanning. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 47-60.

Kosslyn, S., Ganis, G., & Thompson, W. (2001). Neural foundations of imagery. Nature reviews. *Neuroscience*, 2, pp 635-642.

Kottenoff, H. (1961). Effect of retinal image stabilization of the appearance of heterochromatic targets. *Journal of Optical Society of American*, 53, 741-744.

Kraft, J.M. & Brainard, D.H. (1999). "Mechanisms of Color Constancy Under Nearly Natural Viewing", *Proceedings of the National Academy of Sciences*, Vol 96:1.

Lakoff, George. (1987). *Woman, Fire, and Dangerous Things*. Chicago: University of Chicago Press.

Lakoff, George & Johnson, Mark. (1999). *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*. New York: Basic Books, a member of the Perseus Books Group.

Langacker, R.W. (1991). Descriptive Application. In: *Foundations of Cognitive Grammar* Vol. 2, Stanford: Stanford University Press.

Leibniz, Gottfried (1896). *New Essays Concerning Human Understanding*. Trans. A. G. Langley.

Lewkowicz D.J. and Lickliter, R. (Eds.) (1994). *The Development of Intersensory Perception.* Hillsdale, N.J.: Lawrence Erlbaum Associates.

Locke, John (1690/1959). *An Essay Concerning Human Understanding*, ed. A. C. Fraser. New York: Dover.

Lotze, M., Montoya, P., Erb, M., Hülsmann, E., Flor, H., Klose, U., Birbaumer, N, & Grodd, W. (1999). Activation of cortical and cerebellar motor areas during executed and imagined hand movements: An fMRI study. *Journal of Cognitive Neuroscience*, 11(5), p 491-501.

Mangan, Bruce. (1991). What is the "feeling of knowing?". *Consciousness and Cognition*, vol. 9, no. 4, pp. 538-544.

Mangan, Bruce. (2001). Sensation's ghost: the non-sensory "fringe" of consciousness. *Psyche,* 7(18).

Marder, E., and A.A. Prinz. (2002). "Modeling stability in neuron and network function: the role of activity in homeostasis", *Bioessays* 24, pp. 1145–1154

Matlock, T. (2004). Fictive motion as cognitive simulation. *Memory and Cognition*, 32:8, pp1389-1400.

McBride, Russ & Aubrey Gilbert. (in preparation). "An analysis of EEG data in tests of the simulation theory of action initiation".

McBride, Russ (in preparation). *The Homeostatic Mind*.

McBride, Russ. (1999). "Consciousness and the state/transitive/creature distinction". *Philosophical Psychology*, vol. 12, no. 2, pp 181-196.

McBride, Russ. (1994). *Consciousness and the Context Problem.* Master's Thesis. Philosophy Dept., Stanford University.

McCarthy, J. and Hayes (1969). "Some philosophical problems from the standpoint of artificial intelligence".

McFarland, D.J. (1971) *Feedback Mechanisms in Animal Behaviour* (New York: Academic Press).

Meltzoff, A. & Moore, M.K. (1977). Imitation of facial and manual gestures by human neonates. *Science*, 198: pp. 75-8.

Merkel, F.W., Fromme, H.G., & Wiltschko, W. (1964). Nichtisuelles Orientierungsvermögen nächtlich zeihender Rotkehlchen (Erithacus rubecula). *Naturwissenshchaften*, 45, pp 499-500.

Miller, George A. (1956). The magic number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, pp. 81-97.

Millikan, Ruth G. (1984). *Language, Thought, and other Biological Categories*, Cambridge, Mass. MIT Press: A Bradford book

Millikan, Ruth G. (1989). "Biosemantics", *The Journal of Philosophy*, 86, pp. 281-97.

Meltzoff, A. and Moor, M.K. (1977). Imitation of facial and manual gestures. *Child Development*, 54: 702-9.

McCollough, C. (1965). The conditioning of color perception. *American Journal of Psychology*, 78, 362-38.

Medin, D. L., Ross, B. H., & Markman, A. B. (2005). *Cognitive Psychology*. 4th edition. John Wiley & Sons.

Moreno, J. D. (2006). Mind Wars: Brain Research and National Defense. New York Dana Press.

Narayanan, S. (1997). *KARMA: Knowledge-Based Action Representations for Metaphor and Aspect*. UC Berkeley Ph.D. Thesis.

Newell, A. and Simon, H.A. (1976). Computer Science as Empirical Inquiry: Symbols and Search, *Communications of the Association for Computing Machinery* 19: 113-126.

Noë, Alva. (2001). "A Sensorimotor Approach to Vision and Visual Consciousness", *Behavioral and Brain Sciences*, 24:5, 939-73.

Noë, Alva. (2004). *Action in Perception*. A Bradford Book.

Nyberg, L., Petersson, K.M., Nilsson, L.G, Sandblom, J., Åberg, C., Ingvar, M. (2001). Reactivation of motor brain areas during explicit memory for actions. *NeuroImage*, 14, 521-528.

O'Regan, J.K. & Noë, Alva (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences.* 24:5.

Penfield, W. (1975). *The Mystery of Mind: A Critical Study of Consciousness and the Human Brain*. Princeton, NJ: Princeton University Press.

Piaget. (1937/1954). *The Construction of Reality in the Child*. Translated by M. Cook. New York: Ballantine.

Piaget, J. (1967). *Biologie et Connaissance: Essai sur les Relations Entre les Régulation Organiques et les Processus Cognitifs* (Paris: Editions Gallimard).

Porro, C., Francescator, M., Cettolo, V., Diamond, M., Baraldi, P., Zuian, C., Bazzocchi, M, & diPrampero, P. (1996). Primary motor and sensory cortex activation during motor performance and motor imagery: a functional magnetic resonance imaging study. *Journal of Neuroscience*, 16, 7688-7698.

Powers, W. T. (1973a). *Behavior: The Control of Perception*. Chicago: Aldine.

Powers, W. T. (1973b). Feedback: Beyond behaviorism. *Science*, 179, 351-56.

Powers, W. T. (1990). *Living Control Systems*: Selected Papers of William T. Powers. Gravel Switch, KY: Control Systems Group.

Powers, W. T. (1992). *Living Control Systems II*: Selected Papers of William T. Powers. Gravel Switch, KY: Control Systems Group.

Prinz, J. (2006). Putting the brakes on enactive perception. *Psyche*, 12:1.

Rich N.M., & Wenner, P. (2007). Sensing and expressing homeostatic synaptic plasticity. *Trends in Neuroscience* 30: 119-125.

Richardson, D., Spivey M., McRae K., & Barsalou L. (2003). Spatial representations activated during real-time comprehension of verbs. *Cognitive Science*, 27, pp767-80.

Rips, L. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behvior* 14:665-81.

Rock, I. (1983). *The Logic of Perception*. Cambridge, MA: Bradford/MIT Press.

Rosch, E., Mervis, C., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology* 7:573-605.

Russell, James. Explaining Mental Life (1984) BF38.R87 1984

Sayre, K. M. (1986). Intentionality and information processing: An alternative model for cognitive science. *Behavior and Brain Science*, 9, 121-66.

Searle, John. (1983). *Intentionality*. New York: Cambridge University Press.

Shahidullah, S. & Hepper, P. G. (1992). Hearing in the fetus: Prenatal detection of deafness. *International Journal of Prenatal and Preinatal Studies* 4 (3/4): 235-240.

Shatz, C. J. 1990. Impulse activity and the patterning of connections during CNS development. *Neuron*, 5:745-56.

Shepard, R. N. & Cooper, L. A. (1982). *Mental Images and Their Transformaitons*. Cambridge, MA: MIT Press.

Shepard, R. N. & Metzler J. (1971). Mental rotation of three-dimensional objects. *Science*, 171, 701-703.

Sherwood, Lauralee. (2001). *Human Physiology: From Cells to Systems*. Pacific Grove, CA: Brooks/Cole.

Smith, Linda. (1994). Forward, in Lewkowicz D.J. and Lickliter, R. (Eds.) *The Development of Intersensory Perception.* Hillsdale, N.J.: Lawrence Erlbaum Associates.

Smith, S.M. (1979). Remembering in and out of context. *Journal of Experimental Psychology: Human Learning and Memory*, 5, 460-471.

Smith, S.M. (1986). Environmental context-dependent recognition memory using a short-term memory task for input. *Memory & Cognition*, 14, 347-354.

Smith, S.M., Glenberg, A., & Bjork, R.A. (1978). Environmental context and human memory. *Memory & Cognition*, 6, 342-353.

Smith, Scott M., Brown, Hugh O., Toman, james E.P., and Goodman, Louis S. The Lack of Cerebral Effects of d-Tubocurarine. *Anesthesiology* Vol. 8, No. 1, (January, 1947) pp. 1-14.

Solomon, P., Kubzansky, Philip E., Leiderman, P. Herbert, Mendelson, Jack H., Trumbull, Richard, & Wexler, Donald , Eds. (1961). Sensory Deprivation: A Symposium Held at Harvard Medical School. Cambridge, MA, Harvard University Press.

Spelke, E.S. (1991). Physical Knowledge in Infancy: Reflections on Piaget's theory. In S. Carey & R. Gelman (Eds.) *The Epigenesis of Mind: Essays on Biology and Cognition* (pp. 133-169). Hillsdale, NJ: Erlbaum.

Spelke, E.S., Born, W.S., & Chu, F. (1983). Determinants of infant perception. In J. Rosenblatt, C. Beer, R. Hinde, & M. Bushnel (Eds.), *Advances in the Study of Behavior*. New York: Academic Press.

Stanfield, R. & Zwaan, R. (2001). The effet of implied orientation derived from verbal context on picture recognition. *Psychological Science*, 12 pp 153-156.

Talmy, L. (2000). *Toward a Cognitive Semantics*. 2 Volumes. Cambridge, MA: MIT Press.

Taylor, J.G. (1962). *The Behavioral Basis of Perception*. New Haven, CT: Yale University Press.

Thelen, E., Schöner, G., Scheier, C. and Smith, L.B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences*, Vol 24:01, 1-34.

Tomasello, M. (2000). First steps in a usage based theory of first language acquisition. *Cognitive Linguistics*, Vol. 11, issue 1-2, pp 61-82.

Vanhatalo, Sampsa & van Nieuwenhuizen, Onno (2000). Fetal pain? *Brain and Development*, 22, 3, 145-150.

Van Gelder, T. (1995). "What might cognition be, if not computation?", *The Journal of Philosophy, 91, pp. 345-81.*

Varela, F.J., Thompson, E., and Rosch, E. (1991). *The Embodied Mind.* Cambridge, Mass: MIT Press.

Volman & Pearson (1980). What the fetus feels. *British Medical Journal*. 233-234.

Walcott, Charles (1977). Magnetic fields and the orientation of homing pigeons under sun. *Journal of Experimental Biology*, 70, pp. 105-23.

Wheeler, M., Petersen, S., & Buckner, R. (2000). Memory's echo: vivid remembering reactivates sensory specific cortex. *Proceedings of the National Academy of Science of the U.S.A.*, 97 pp 11125-11129.

Wheeler, Michael. (2005). *Reconstructing the Cognitive World: The Next Step*. MIT Press.

Wiener, N., Rosenblueth, A., & Bigelow, J. (1943). Behavior, purpose, and teleology. *Philosophy of Science*, 10: 18-24.

Winograd T., & Flores, F. (1986). *Understanding Computers and Cognition: A New Foundation for Design*. Norwood, NJ: Ablex.

Wittgenstein, L. (1953/2001). *Philosophical Investigations*. Blackwell Publishing.

Zwann, R. Stanfield, R., & Yaxley, R (2002). Do language comprehenders routinely represent the shapes of objects? *Psychological Science*, 13, 168-171.