Chapter 7: The Most Important Thing Neuropragmatism Can Do: Providing an Alternative to 'Cognitive' Neuroscience

Eric P. Charles, Sabrina Golonka, and Andrew D. Wilson

I. The Long-Known Problem

> The frantic search for sources of motivation and of emotion in visceral activity, though initiated by introspective analysis, has been supported by the faith that the nervous system is only a conductor having no sources of energy within itself…. [Theories derived from this approach] may or may not be true, but their truth must be demonstrated by experiment and cannot be assumed on a background of questionable neurology…. The facts of both psychology and neurology show a degree of plasticity, of organization, and of adaptation in behavior which is far beyond any present possibility of explanation. <u>For immediate progress it is not very important that we should have a correct theory of brain activity, but it is essential that we shall not be handicapped by a false one</u>.
>
> – Karl Lashley[1]

Karl Lashley was, by several accounts, the greatest psychologist-neurologists of the early 20th century. He is most well-known for his search for the 'engram,' the biological 'memory trace' in the brain. He eventually abandoned this search, concluding correctly that the engram did not exist. Despite that, his search has been reinvented in modern times, with a slightly different vocabulary, in light of modern advances in neuroimaging. The researchers on this new quest seem to think that Lashley's failure was due to technological limitations – if only that brilliant man had a 10 tesla fMRI machine, then he would have found it for sure! But that was not the problem. The problem was that the central metaphor guiding his search did not match reality in important ways; there simply is not a place in the brain where memories are stored. The biological events involved in the adaptive change of behavior over time, and even the biological events involved in re-experiencing the past do not resemble the process of retrieving information from a reliable, representational storage system. What was once a perfectly viable hypothesis about how the mind works is simply not true.

Even though cognitive scientists now increasingly recognize that this hypothesis was false, many neuroscientists continue to use language that is derived from it. While some will protest that we are exaggerating, an examination of the literature reveals that neuroscientists continue to rely on the implications of the (false) mind-is-computer hypothesis. Why do they do this? We suggest that they continue because they have no alternative, and that the greatest service neuropragmatism can offer is to provide a better way of talking about the role of the brain.

I.A. Cognitive Neuroscience – The Science of a Bad Hypothesis

Cognitive neuroscience was known foolish before it began. 'Cognitive psychology' in its modern form, is an aggressive use of a computer metaphor to explain how the mind operates. This approach had some important virtues. It led to many discoveries about mental phenomenon, and it freed researchers from the oddly restrictive behaviorism of a prior era. However, while we might one day have machines that work like brains, our brains do not work much like current computers, and certainly not much like the computers of the late 1950's and early 1960's. While the new metaphor has been very productive, when it is pushed to the extreme it is as silly and restrictive as the reductive behaviorism it was supposed to replace. It consists, for the most part of a proliferation of 'hypothetical constructs' – such as reliable 'memory stores' or 'intelligence' – which are conceived as inherently non-observable things-in-the-head that can be offered to explain behavior.

MacCorquodale and Meehl wrote the first paper to meaningfully distinguish hypothetical constructs from related notions.[2] They argued persuasively that there is nothing inherently wrong with hypothetical entities, and that that the hard sciences use them quite effectively. However, they caution psychologists against two mistakes:

> First… there is the failure explicitly to announce the postulates concerning existential properties, so that these are introduced more or less surreptitiously and ad hoc as occasion demands. Secondly, by this device there is subtly achieved a transition from admissible intervening variables to inadmissible hypothetical constructs. These hypothetical constructs, unlike intervening variables, are inadmissible because they require the existence of entities and the occurrence of processes which cannot be seriously believed because of other knowledge.[3]

They later rephrase the second point more forcefully, reminding psychologists that:

> It is perhaps legitimate, even now, to require of a hypothetical construct that it should not be manifestly unreal in the sense that it assumes inner events that cannot conceivably occur.[4]

Without passing judgment on the early days of cognitive psychology, we can report that the current field makes those mistakes flagrantly. It is quite likely that individual research programs are relatively consistent in their use of hypothetical constructs, but across programs the divergence is tremendous. As has been pointed out elsewhere, the terms of cognitive psychology give a false sense of concreteness, while really being hopelessly vague:

> For example, if you were interested in 'how memory affects perceived meaning,' how would you go about investigating it? 'Memory' could be operationalized in a wide variety of ways; memory could refer to what people remember in the present, any change in behavior following an event you (the investigator) are sure happened, the effects of bodily (neuronal?) alterations on future behavior, the retention of a conditioned response, etc. 'Perception' and 'meaning' are similarly, if not more, ambiguous. Perhaps you will ask people to recall certain events in their lives, then have them interpret ink-blots; perhaps you will flash lists of words for 50 milliseconds at a time, then see if they feel positively towards words you repeated several times; perhaps you will ask people the perspective from which they view specific memories, then test them for their accuracy in recalling important aspects of the event; the variety is almost infinite. Despite this incredible lack of specificity of the question, it is easily transformed into a concrete empirical endeavor. Whatever the form of the study, the results could be reported in the

> local newspaper and the average high school student reading it will nod their head as if they are learning something important about human nature. Further, if you also give the task to chimps or dolphins it can make national headlines – 'Chimps perception of meaning less affected by memory manipulations than teenagers' – despite the headline saying nothing concrete, it is perceived as understandable and straightforward.[5]

Cognitive neuroscience is, alas, the process of looking inside the brain to try to find the hypothetical constructs of cognitive psychology. As the quote from Lashley at the beginning of this paper should make clear, this process did not start with the relatively recent advent of neuroimaging, it was alive and well at the start of the 20th century. To the extent that the terms of cognitive psychology can be made concrete, they are arbitrary; and, to the extent that the original meanings of the terms are held to tightly, they don't refer to anything that actually happens in the brain. Thus, the terms either don't mean anything obvious, or they are obviously wrong.

Many cognitive neuroscientists, to their credit, clearly recognize this problem. The amazing models of brain activity currently being offered by cognitive neuroscientists retain the terminology of cognitive psychology, while jettisoning most of the terms' obvious meanings. The importance of neural network modeling, for example, is that it provides a model for behavioral change without anywhere having stored 'memories,' a central 'thinking' system, or a clearly identifiable 'representation' of the difference-in-the-world that leads to the different responses of the system. And yet, those words are commonly used in discussions of neural networks; the terms remain, though they are stripped of any semblance of their original, intuitive, natural meaning. The terms remain and cause endless confusion amongst those trying to present and understand the implication of the research being conducted. The terms remain, because there is no other language to use.

Is the language really so corrupting? Yes. We did a cursory review of Journal of Cognitive Neuroscience, Volume 20 (2008), a journal with an impact factor to make all but the very top psychology-specific journals envious. Focusing first on memory-related papers:

- We encountered much talk of 'false memory.' This seemed perfectly fine at first, until we realize that in most other articles 'memory' clearly indicates accuracy with respect to past events. In those articles, when one is mistaken about the truth of the situation, one has failed to remember. Thus, depending on the article, a mistake about past events is either defined as 'false memory' – a variety of memory – or the lack of any variety of memory.
- We also learned that the hippocampus and perirhinal cortex contribute to 'episodic item and associative encoding.' While this suggests that items, and the associations between items, are somehow stored in the identified areas of the brain, it actually just means that those areas of the brain fire more frequently shortly after you see something new. It is quite clear that, despite Cognitive Neuroscientists' renunciation of dualism, they cannot communicate without the notion that something is somehow stored in the brain. But the brain is not a bank vault into which things are placed for later retrieval, the brain is an organ that adapts to the environment over time.
- Distracted, we citation-jumped a bit into the future, learning that the ventral regions of the posterior parietal cortex 'binds episodic features stored in disparate neocortical regions.' All fine and good on the surface, until we remembered that features are not stored anywhere, and that the metaphor of 'binding' suggests that we have reified the

memories as things that can be attached to other things while some third-thing wraps around them.

Returning to JOCN 2008, we examined work on the 'central executive,' which was supposed to be a place-holder, a term for stuff we didn't understand.

- We found out that the thalamus is particularly crucial to activities of this non-entity.

'Processing,' another common term, might seem neutral, but, in fact, use of the term assumes a very odd notion of what the brain does. For example:

- One article that assumed the brain was an information-processing system starts by telling us that there is 'disagreement regarding the nature of the information that is maintained in linguistic short-term memory.' Of course there is! That is because 'information,' 'maintain,' 'linguistic,' and 'memory' are all highly ambiguous terms, which likely, even in their less ambiguous moments, do not reflect what the brain was doing during the intervals of interest to the study (i.e., waiting patiently in a brain scanner to see if a second stimulus matches a first).

Not all is bleak though. There were somewhat more acceptable suggestions that a given part of the brain could be 'associated with,' 'correlated with,' 'coded to,' 'involved in,' 'activated during,' 'participating in,' 'underlying,' 'playing a critical role in,' or 'supporting' different psychological happenings. There was also discussion of how one area of the brain ‘inhibits’ or ‘activates’ another area of the brain, which is at least better than the claim that a brain area ‘inhibits’ or ‘activates’ a psychological happening.

The latter language is more neutral, and hence better than the cognitive language. Also on the positive side, a brief scan of more recent volumes of this journal suggest that they now favor use of the neutral references to what the brain does, whereas earlier volumes seemed to favor the egregious descriptors. However, while this is a step in the right direction, the neutral descriptors are not very satisfactory. They are about as useful as telling us that sparkplugs ‘participate in’ the movement of a car.

I.B. What to do?

As many of the other articles in this book argue, your mind is more than the activity of your brain; your mind extends into your body, and into your environment, including the social elements thereof. However, neuroscientists interested in the mind continue to use the inaccurate, brain-focused language given to them by cognitive psychology. They can be forgiven for one obvious reason: there is no alternative. Neuroscientists want to do their work, publish it, get funding, get coverage by the press, etc. They want to study the brain. Hence, they can be forgiven for not fixing the language of psychology; they can be forgiven for stretching the terms they are given beyond any semblance of the terms’ original meanings because, frankly, neuroscientists have better things to do. The psychologists and philosophers of psychology, on the other hand, have been purely negligent in their duty. Philosophers have not only been complacent in the abuse of language, they have facilitated it willfully. For example, philosophers have worked hard to warp the definitions of ‘representation’ until it bears no obvious resemblance to the notion of presenting-something-again, as in the way a picture of your childhood home re-presents certain aspects of the home. Psychologists as well, or at least those who have recognized the problem, have shied from their clear duty.

Which psychologists are most to blame? Presumably those most guilty are those who complain the most about the existing language. Rather than sit idly by, complaining about silly things their colleagues say, these psychologists should be providing more intelligent ways to discuss what the brain actually does. There are many groups of psychologists who fit into this category, but we can only speak in more depth about those with which we are most familiar, Ecological Psychologists and Radical Behaviorists. While the opinions represented by these fields are considered 'minority' opinions in psychology, they are quite large and empirically successful minorities, generating experimental and observational research at the cutting edge of psychology, kinesiology, education, and neuroscience. Moreover, it should be noted in this context that these traditions are intellectually descended from Pragmatism.[6] Thus, you might expect them to provide a good starting point for neuroscientists looking for alternatives to cognitive psychology.

Why, then, have these approaches not provided viable alternatives? Because, while it is clear that the fields of behaviorism and ecological psychology can generate alternative means for speaking about the brain, there are no individuals in the fields clearly obligated to do so. Further, due to the history of these two fields, individual researchers tend to shy away from brain work, focusing more on a behavioral level of analysis. This is not to say that brain work is not done in these fields, but to point out that existing work gives only limited hints as to better ways to talk about the brain. What is needed is a fully formed alternative language. While we cannot possibly provide that here, we can anticipate certain aspects of what will come.

II. Ways of Speaking about the Brain

> I don't think it's meaningful to say I currently (sitting at my desk) have the ability to hit a softball; that ability isn't stored somewhere for me to access on demand. I do think it's closer to say I have the ability to become something that can hit a softball.
>
> – Andrew D. Wilson[7]

We believe there are two ready sources of insight into how we can better speak about the brain. First, there is a rich foundation of conceptual insights set up in the early 1900's by those attempting to reimagine psychology in light of, among other things, insights from pragmatic philosophy. Second, there is a rich variety of recent experimental and observational results that reveal flaws in the now-accepted vision of the psychology-neuroscience relationship, suggesting that those early intellectual risk takers were really on to something. Our presentations of both of these foundational elements must be brief, and selective, but we hope even in this limited space to offer a taste of the potential for neuropragmatism to lead us to better ways of speaking about the brain and its function.

II.A. Speculative Neuroscience, circa 1936 – Levels upon levels of loops and dynamics

Perceptual research under the name 'Ecological Psychology' emerged out of the work of James J. Gibson in the early 1960's, and the movement has tremendously advanced our understand of perception-action systems. It is related to some of the more molar approaches to behaviorism, and it overlaps with the perceptual-control theory that grew out of Power's work around the same time period.[8] A key lesson out of the ecological approach is that organism and environment are not connected by simple input or output relations, but rather by dynamic, circular relations. This is not a trivial observation, as circular relations allow for easy solutions to problems that would be very difficult for simple input-output mechanisms to solve. A body that is part of such a circular relation can rely on stable aspects of the world to accomplish many tasks without needing much 'computation' on the part of the brain.[9]

Perhaps the most aggressive expression of this idea was made by E. B. Holt, a proto-behaviorist, proto-ecological psychologist, and proto-epigeneticist.[10] Almost eighty years ago, in a chapter dedicated to the Russian physiologist Beritoff, Holt tried to describe the multiple, embedded levels of circularity available to control behavior.[11] In 'Eight Steps of Neuro-muscular Integration,' Holt attempts to show how the basic principles of a sensory-motor loop can account for numerous aspects of behavior, and behavioral development. His proposed system is ambitious, and ultimately a bit unwieldy, but it shows how much a system using the organism-in-the-world approach of pragmatism can simplify what we think the brain needs to do.

II.A.1. Lowest level loops: Circular reflexes

In the lowest level loops, a motor neuron causes the stimulation of sensory neuron, which in turn affects the firing of the motor neuron. For example, when a neuron's firing tenses a muscle it squeezes proprioceptive cells, and activation of that proprioceptive cell will typically affect the firing of the original motor neuron, leading to muscle relaxation. Two neurons, properly connected, can thus maintain static positions or smooth forces, with no additional input required. Such loops greatly simplify, for example, the challenge of holding your posture. This suggests that, while a brain might adjust the strength of connections between neuronal loops at this level, depending on the posture-to-be-obtained, it might need to do little else to maintain such postures.

II.A.2. Level 2: Reciprocal innervation

As circular reflexes of antagonistic muscle groups are integrated together, neurons that flex one muscle are linked to the neurons that relax other muscles. Such loops greatly simplify, for example, the challenge of lifting and lowering limbs. This suggests that the brain need not tightly control all muscle groups involved in a movement, as the muscles are innervated to handle much of that coordination peripherally.

II.A.3. Level 3: Adient reflexes

To be ‘adient’ means ‘to approach.’ Further integration of sensory-motor loops can easily explain how organisms orient-towards and approach objects and events. Without guidance from a central system, motor neurons strengthen their connection with sensory neurons until any sources of a given stimulation can produce the same movements by which an organism self-produces that stimulation. For example, a human fetus makes a lot of random movements, including finger flexes that result in palm touching. ‘Cannalization’ of that loop results in a ‘reflex’ whereby fingers flex in response to palm stimulation, i.e., the newborn baby’s grasping reflex. This same principle can be applied to other perceptual systems to help explain phenomena such as auditory imitation.

While we are here dealing with processes that, at least in higher organisms, will likely involve neural activity in the brain, it should be clear that no ‘top-down,’ homunculus-like process has been invoked. Even without that, these loops provide quite a bit of sophistication:

Because they will naturally develop so as to function effectively given any 'special posture from which a movement starts,' they can accommodate in-process perturbations equivalent to changes in initial posture, all without intervention from a 'higher-function.'

II.A.4. Level 4: Concatenation of reflexes

As Level 3 processes become chained together, we find it ever-more accurate to describe what an animal is doing in terms of behavior directed towards some external thing. That is, while Level 1, 2, and 3 loops can be described in terms of the immediate stimulation at a sensory neuron, the integration of such systems leads the organism to be responding simultaneously to several properties of particular object – and to respond to the conjunction of several different properties in a particular physical location is to be responding to the object in which those properties conjoin. In this context, Holt specifically invokes William James's notion of Habits, and connects habit formation with broader discussion of the 'association' of 'images' and 'ideas.' We believe that Holt is pointing out how the strengthening of neuro-muscular integration at this level can explain behaviors ranging from a dog drooling at the sight of a particular laboratory assistant, to the tipping of a hat at the sight of an old acquaintance. The behavior in each case is controlled by a loop from the presence of a complex, multi-propertied object or situation to a particular behavior or set of behaviors, with no higher control needed.

II.A.5. Level 5: Movements of exploration

Here we combine loops from the previous levels and add in a bit of random movement to create the exploratory behavior one sees, for example, when a ball is placed in a toddler's hands. If these experiences repeat, then new patterns of action are set up, such that the same 'exploratory' movements can happen even in the absences of the ball. (If you can hold your hands out in front of you and pretend to be touching the outside of a ball, then you know what Holt is talking about.) We now start to see the implications of perception-action loops for more clearly 'psychological' processes. As Holt points out, to perform the initial exploratory movements is, literally, to feel the contours of the ball. To be able to perform the same motions without the ball present is then, in some sense, to re-member the ball (to feel-it-again) despite its absence.[12] Holt is thus providing initial hints as to how organism-environment loops can produce the rudiments of 'memory.' This gets our foot in the door to higher functions, without needing to invoke a store of re-presentations, or other similar myths, and certainly without some central system that indexes, retrieves, and re-interprets what is 'stored' in other parts of the brain. The loop itself, strengthened sufficiently through development, can account for the re-behavior.

II.A.6. Levels 6-8

Alas, at this point it becomes clear that while Holt will continue to ambitiously try to explain ever 'higher' psychological functioning through higher level loops – all without needing a central control mechanism to direct such functioning – he will not move us further towards developing a vocabulary to describe the role of the brain in such processes. Level 6, interconnecting of multiple exploratory levels, explains, for example, how hearing a song invokes a visual 'image' of past events.[13] Level 7, the strengthening and interconnecting of several Level 6 mechanisms, allows for what we might call 'imagination,' as such loops allow people to act in elaborate ways towards objects that are not currently present. This level also begins the appearance of 'self-determination,' as the triggering stimuli have little obvious connection with the behaviors themselves, as when a certain smell leads someone to quietly sing a song to themselves and perhaps even begin to recreate movements from a high-school dance.

In such a situation, there are specific causal relations to be discovered, created by the canalization of certain neural connections due to past events, but the causes of the behaviors are 'non-obvious.'[14] Finally, Level 8 deals with the integration of responses to multiple objects or events, which will occur in any sufficiently complex organism. Holt admits that this might not be impressive as a final level of integration, at least at first blush. However, he believes it is key for understanding phenomenon like successful navigation of the environment, and intelligent behavior in general (i.e., behaviors that must occur with respect to several simultaneous aspects about the world).

Though Holt could lay all this out in 1936, he was working largely in a vacuum, based on his own personal intellectual connections with William James. Most of his colleagues were creating the foundation for cognitive psychology, and his students had not yet developed the behavioral and ecological systems that would lead to their future empirical successes. There is now much evidence to support Holt's notion of loops from bottom-to-top; loops that provide flexibility and redundancy, with continuous environmental support playing an essential role all the way from bottom to top. Top researchers are currently struggling valiantly in an effort to come to grips with the role of the brain in such a system, and in so doing, a new vocabulary is slowly, and awkwardly, starting to emerge.

II.B. A Modern Update: What is the Brain Doing?

It is clearly the case that the brain is up to something important. In humans, it consumes approximately 20% of our body's energy[15] and our bodies are thoroughly innervated by projections to and from the nervous system. So what, precisely, is it up to?

Real brains are always busy[16] and are utterly integrated with a wide variety of other internal systems (e.g. musculo-skeletal, circulatory, and cardiovascular) as well as to the external world around it, and also to some systems that sit ambiguously between internal and external (e.g., the digestive and respiratory systems). In each of these cases, the integration of the systems is mediated via perception. Successful activity, therefore, entails perceiving and controlling both these internal and external resources to solve specific tasks; controlling both the inherent dynamics of the system and the incidental dynamics due to the situation.[17]

Brains presumably play a critical role in this process, but they aren't the only player in the game, and there are a lot of things brains simply don't need to do because some other system takes care of it. For example, in the 'equilibrium point' model of limb control, the brain does not need to explicitly set a series of joint angles to move into a given joint configuration; instead, the nervous system merely needs to set a force-length relation for the muscles at the joint, and the limb will smoothly move to its new equilibrium point, with the specific trajectory determined completely by the local anatomical constraints.[18] One is tempted to say that the local constraints 'select' the trajectory, but of course almost everything implied by the term 'select' does not apply, and the more useful description is that the trajectory 'emerges' from this constellation of constraints on a system with a specific dynamic. In effect, the anatomical and physiological organization of the arm greatly simplifies the situation, allowing the brain's role in smooth movement to be efficiently minimal.

This principle implies the following: In order to work out what the brain does need to do, we need a clear theory about what everything else is up to, so that we don't waste time looking for the neural implementation of something taken care of elsewhere in the body. In other words, we need a job description for the brain. To develop this job description, we will first lay out what the non-neural components of extended, embodied cognitive systems do to support behavior. This will then leave us with work to be done by the brain, and the resulting gap suggests a very different role than the traditional cognitive model.

II.B.1. Bodies as task-specific devices

Cognitive systems are not general purpose systems – we are, instead, at any given moment, one kind of task-specific device (TSD).[19] A TSD is a smart solution to a specific problem. Smart devices[20] take advantage of external, task-specific resources, as well as the internal resources described below, to find and use locally optimal solutions to problems. The kind of device we currently are is a function of what we've been up to recently; the specifics of the device reflect the nature of multiple subsystems and how these respond to the task space we are currently embedded within. These devices must therefore be assembled and controlled with respect to task-specific constraints (this requires perception of those constraints). Specifically, these task-specific devices must be softly assembled. To be 'softly assembled' is to be built out of whatever resources are available to solve the task at hand, resources which can be removed and used again later as part of a different TSD. This is in contrast to a device which is built to do one particular thing, and can only ever do that thing. Take a crude example: I can use my hand to pat someone on the head or to pull their hair. The resource (the hand) is the same, but adjustments to the larger system (the body) have placed it in a different relation to the other element of the task (the head). The most likely primary role for the nervous system is going to be in supporting this soft assembly of task resources; our first task, therefore, is to determine what kind of resources we are talking about.

II.B.2. 'External' Resources (Gibsonian Information)

An essential and vastly underappreciated aspect of the world is the flux of energy in the 'empty space' around us, the light and pressure waves crisscrossing, chemical gradients, magnetic fields, etc. Gibson labeled the structural elements of this ambient energy that can support behaviour 'information.'

This information flow is surprisingly stable, as is our access to it (following the requisite learning).[21] We are not adrift in a 'blooming, buzzing confusion' (to use the oft misunderstood William James quote). As with other stable aspects of our ecological niche, we can (and evidence shows that we do) rely on this information to do a lot of work for us. We exist in and move through a flow of information and our behavior emerges as we interact with that flow. There is no need to construct a model of our environment; as Rodney Brooks famously claims, we can let the world be its own model.[22]

We are not passive blank slates, however. At any given moment in time we are very specific measurement devices, sensitive to some information variables but not others, and capable of responding in some ways but not others. We spend our days being one device, then another, then another, in response to changes in the information flow, and in accordance with our capacities. The information flow alters as our location in space and time changes, and is specific to the current environment.

II.B.3. Internal (non-neuronal) Resources

The human body has various subsystems with different capacities from which to assemble the various task-specific devices that we can become. Each has their own dynamical properties and typically operate over different time scales. These subsystems are non-linearly coupled to each other in a variety of ways, but can be separated out for analytical purposes. What follows develops the taxonomy described by Bingham.[23]

The link-segment skeletal system provides a very stable base for the other systems; the bones are all different sizes and shapes and are attached to each other with very particular joint configurations. The overall configuration does not radically alter over our lifetimes without serious injury, although it does change in response to prolonged use (e.g. thickening with weight training, thinning in extended weightlessness). The role of this system is to provide a physical substrate for the transmission of the forces involved in moving and interacting with the world. The musculo-tendon system is also globally quite stable in its configuration, but its form is more responsive to use on time scales ranging from days to weeks (e.g. weight-bearing exercise increases the size of muscles and stretching alters the give of tendons). The role of this system is to generate the forces required to move the link-segment system; it's organized in ways that solve many of the 'degrees of freedom' problems inherent in controlling a complex system.[24] The circulatory system is highly responsive to current events: it only takes seconds for your heart rate to accommodate current energy demands, and veins and arteries change size and shape in response. The role of this system is to deliver energy on demand to the musculo-tendon system so it can generate the forces required to move the link-segment system. The respiratory system is another highly responsive system which adapts to current requirements on very short time-scales. Its role is to provide oxygen and remove carbon dioxide, to enable the continuing metabolic processes powering the muscles. Nutritional systems operate over longer time scales, and their role is to provide the nutrients and energy required for the above systems. Hormonal systems operate over a wide variety of time scales, and are typically involved in cyclical regulation (e.g.

circadian and seasonal rhythms, alternation between rapid-response-readiness and relaxation, etc.). Hormonal systems have broad affects across the other systems.

II.B.4. Internal (neuronal) Resources.

Where does the nervous system fit in this taxonomy? The nervous system responds to changes in its environment on millisecond time scales, which makes it suitable to serve as a fast-response system, and, we suggest, the primary medium supporting the informational (as opposed to direct, mechanical) couplings between the various subsystems. These need to be both stable and flexible, so that an organism can maintain informational couplings as long as needed, but de-couple and reform into a different device when the context changes. Structurally, then, the nervous system is dynamically stable, probably edge-of-chaos stable.[25] This stability does not reflect specific wiring connections being preserved, but rather the preservation of global function despite change in the underlying wiring, and it is actively maintained by the form of the informational flow. If you change that flow, the dynamic structure smoothly alters in response; as, for example, when you have someone pick up a tool,[26] or provide information suggesting a person actually has three hands.[27]

This structural flux still supports stable function; in biology, this capacity is referred to as functional homeostasis or, a bit oddly, as degeneracy.[28] This is defined as the capacity of a system to perform similar functions despite variation in how the functions are implemented.[29] The most remarkable demonstration of functional homeostasis in a nervous system comes from Prinz, Bucher, and Marder,[30] who simulated the functional output of the 20,500,000 different ways there are to combine the neurons involved in the pyloric (swimming) rhythm of the lobster's stomatogastric ganglion. Astonishingly, 20% of these combinations preserved pyloric-like activity (i.e., rhythmic swimming motions). Even when the simulations were constrained by data from actual lobsters, there were still 452,516 different ways to achieve the same functional outcome. Pyloric-like activity occurred for all 150 possible combinations of neurons, and for synapse weights across the entire functional range (with the exception of one particular synapse weight which had to be very weak: this matches the biology of the actual system very nicely). In effect, the behavior of the network exhibited extraordinary functional homeostasis (450,000 ways to produce the same rhythm) and the parameters were stable over several orders of magnitude (so long as the appropriate compensation occurred elsewhere in the network). Prinz et al. suggest this flexibility shows that the nervous system is <u>not</u> trying to maintain a specific set of wiring connections and weights; <u>the nervous system is trying to maintain the capacity to perform a specific function</u>. Sporns notes that 'such homeostatic mechanisms are essential for the long-term stability of the brain given the continual remodeling and structural turnover of its cellular and molecular components.'[31] Function is primary, and how that function is achieved is of secondary importance.

With some historic context and modern research insights in hand, we can now turn our attention to better explaining what something like the brain would do in this system.

III. A Job Description for the Brain

Functional homeostasis seems to be an organizing principle that extends across multiple systems and levels in biology. The specific structural implementation is rarely that informative: It may simply be the solution that evolution and development happened to find first, by chance. <u>The real question of interest is: What is the function performed by the whole system, and how does the configuration of the system relate to the task?</u> This makes it all the more important that we adopt a task-specific approach[32] to understanding brain and behavior. A task-specific research strategy addresses the overwhelming and non-linear 'degrees of freedom' problem in the study of the human perception-action system. The proposed solution is to study the organization and composition of particular, carefully described examples of the human perception-action system at work, and then to describe the kinds of resources the system has access to for building such devices. In other words, figure out how a given task is achieved, and by doing so, come up with a list of things you know the system as a whole is capable of, which can then inform your study of related tasks. Typical model systems in perception-action research include coordinated rhythmic movement,[33] locomotion,[34] throwing and catching .[35]

A given perception-action task uses specific informational and motor components, and it explicitly does not use others. Formation of a stable perception-action system that will succeed in a particular task often requires that the task-relevant components become functionally 'walled off' from irrelevant components (chewing gum doesn't interfere with walking, unless you are new at both). This enables the system to accomplish its goal stably and reliably for the duration of the task (i.e., by resisting perturbations).[36] Assembly is this process of coupling resources together to achieve some function. These resources are not mechanically connected, however (the way 'the hip bone is connected to the thigh bone,' for example), and they must therefore be coupled <u>informationally</u>.[37] This type of coupling affords the second critical aspect of assembly – <u>softness</u>, or the capacity to re-purpose your resources as required.

The nervous system is perfectly placed to implement these informational, temporary couplings between disparate resources. Information about these resources (both 'internal' and 'external') flows through the nervous system, which shapes itself in response. When this flow alters (i.e. when different task resources are available, or the same resources are in different states), the nervous system re-shapes itself in response, on millisecond time scales. This dynamically stable system is also, critically, everywhere in the body. So, in principle, if there is information about a resource flowing through the nervous system, it can be coupled to any other resource for which there is information.

This, then, changes the kinds of questions we ask about what the brain is doing. We are no longer interested in where and how the brain stores knowledge and learned skills, but in how the brain engages in the real-time coordination and control of behavior in response to stable but evolving patterns of information flow. To understand what the brain is doing, we need to first understand the forms of information to which it resonates. The ecological approach[38] is an example of a research program that is identifying one form of information (perceptual, typically visual information about the dynamics of events in the world). The answers to these questions should then inform the questions we ask about the brain. These new kinds of neuro-questions need a language with which to be asked, and that work remains to be done, hopefully with one eye on what pragmatism can contribute.

IV. Connecting with Pragmatism

The approaches discussed above have been developing in bits and spurts since the early 1900's as part of pragmatism's intellectual lineage within psychology. Let us translate some of the above to make this connection with pragmatism more obvious, acknowledging that we can do so in only a cursory fashion.

There is a lot of research within the ecological tradition looking at the catching of balls. Most people, most of the time, are not prepared to respond rapidly to incoming projectiles. We don't recommend you try this, but we assure you that if you randomly toss baseballs at strangers' heads, you will hit them in the head a lot of the time. In contrast, well-trained people, detecting a quick movement in their periphery and hearing a loud 'Head's up!' will rapidly become a task-specific device specialized in catching things. Under such circumstances, people's accuracy in catching balls is remarkable given the precision required and the degrees of freedom involved. Imagine the ball is coming from slightly behind the person's side, right at the limits of their peripheral vision. Through a diffuse and highly flexible array of neural outputs, the orientation of the body shifts slightly. Included in this, the shoulder and back tense, pulling back the upper arm; the forearm naturally rotates upward without any central control needed, due to the configuration of bones and tendons. Similarly, the hand naturally opens. As the visual system engages more fully, fine motor control becomes coupled with the optical acceleration and optic expansion of the approaching objects so that the hand ends up at the correct location (with no 'computation of the trajectory' required). When the flying object hits the center of the palm, the natural impact begins to close the fingers around it, and quick connections through the spinal cord finish the grasp. As you can see, the nervous system has a pervasive role in temporarily turning you into a ball-catching device, and it will similarly have a role in rapidly turning you into the next type of device as well, but it does so within the context of several nested systems.

But aren't there people who always seem ready to perform this task? Have we not heard of people making highly improbable catches of incoming projectiles before they even knew what they were doing? Yes, of course. Some people are prepared to perform certain types of tasks at more or less any moment. This is what Peirce called 'habit,' and his discussion of habit formed a central part of his influence on psychology. For Peirce, your habits, which also in some sense constitute your 'character' and your 'knowledge,' are those things you do without hesitation. The person who is always at the ready to catch incoming objects is 'that type of person' and clearly knows how to catch a ball; that vigilance is an important part of his or her character. Neuropragmatists are thus obligated to answer the question: What is the role of the brain in the things you do without hesitation?

The modern research presented above, with some addition of Holt's vision towards explaining higher processes, lays a foundation for answering such questions, and more. Thus neuropragmatists are in a good position to help us find better ways of talking about 1) how we transform from one task-specific device to another, 2) how we develop the ability to become new types of devices, and 3) how we sometimes become rigid in our ability to switch between device-types. Or, to use Peirce's phrasing, we need ways to talk more accurately about the role of the nervous system in our developing and solidifying habits. This would be a bare beginning of connecting neuroscience and pragmatism.

An adequate language for this discussion will make explicit 1) the muti-leveled connections between the nervous system, the other bodily systems, and the wider world, 2) the circular nature of those connections, with some equivalent of a perception-action loop supporting behavior at each level, and 3) the tremendous redundancy in the nervous system and the wider system as a whole, which allows organisms to behave adaptively despite wide variations in the activity of the nervous system.

V. Summary

The nervous system develops, in both the short and long term, so as to preserve function, not structure. In order to understand the brain, or at least to talk about it correctly, we have to understand what those functions are, and we have to understand the many ways in which the nervous system can combine with other resources to fulfill those functions. The ecological approach proposes quite specific challenges that the organism has to solve: The organism must detect higher-order invariant patterns of ambient energy that specify action-relevant properties of the world, then use this action-scaled information to coordinate and control softly assembled task-specific devices. Behaviorist approaches to psychology, in contrast, tend to focus on the developmental support of behavior. While behaviorist approaches typically neglect physiology, they emphasize the dynamic nature of developmental process, and point out the incredible flexibility of the nervous system across environments. Neuroscience, in contrast, currently speaks a language derived from the cognitive approach, with its physiologized homunculi and rigid linear mechanisms. Reorienting cognitive science along more pragmatic, ecological lines alters the job description for the brain entirely, and redefines what we should be looking for in the brain. The above is a preliminary attempt to lay out a new way of conceptualizing what the brain is up to. It is time that the non-cognitive approaches to psychology stepped up and helped create a more honest way of engaging neuroscience.

Where to go from here? If you have read this far, the major impression we hope you have gotten is that there is a vocabulary vacuum that needs to be filled. Neuroscientists do not lack knowledge about the brain, and the psychologists do not lack knowledge about behavior. What we lack is a language to bridge the gap. We have so many things we want to say about what the brain does, and if we could articulate these things better, it would provide a more honest way to discuss the results of current neuroscience, and provide a framework to guide future neuroscience. However, we have not struggled with the problems long enough and deeply enough to have a solid, widely agreed upon vocabulary that can serve these purposes. Neuropragmatism, if pursued, will necessarily lead researchers to conclude that the conceptual apparatus of cognitive psychology is inadequate and misleading, especially when applied to the brain. Maybe, however, the pragmatists can stand and deliver where others have not. Maybe neuropragmatism can finally provide us with a sensible way of talking about the role of the brain in psychological processes.

Acknowledgements

This collaboration developed from ideas explored on the author's blogs, 'Fixing Psychology' (fixingpsychology.blogspot.com) and 'Notes from Two Scientific Psychologists' (http://psychsciencenotes.blogspot.com). We wish to thank our readers, and especially those who have left comments and engaged us in discussion. If our page hits are any indication, there is a wide, international web of people interested in these ideas, and we are grateful that readers would find our contributions worthy of their valuable time and attention.

NOTES

1. Lashley 1930, pp. 23–24, emphasis added.

2. MacCorquodale and Meehl 1948.

3. Ibid., p. 106, emphasis added.

4. Ibid., p. 105.

5. Charles and Dege 2008, p. 195.

6. Heft 2001, Smith 1986.

7. Wilson 2011.

8. E.g., Powers 1973.

9. Wilson and Golonka 2013.

10. Charles 2011a.

11. Holt 1936.

12. See also Charles 2011a.

13. For a more modern argument of similar points, see Tonneau 2011.

14. For a solid introduction to modern work on non-obvious causes of behavior, see Miller 1997.

15. Clark and Sokoloff 1999.

16. Raichle et al. 2001.

17. Bingham 1988.

18. Feldman 1966, 2011.

19. Bingham 1988.

20. Runeson 1977.

21. Gibson and Pick 2000.

22. Brooks 1997.

23. Bingham 1988.

24. Feldman 1966, 2011.

25. Alexander and Globus 1996.

26. Carlson et al. 2010.

27. Gutersham, Petkova, and Ehrsson 2011.

28. Mason 2010.

29. Sporns 2011.

30. Prinz, Bucher, and Marder 2004.

31. Sporns 2011, p. 68.

32. Bingham 1988.

33. Bingham 2004b, a, Haken, Kelso, and Bunz 1985.

34. Fajen and Warren 2003, Wilkie and Wann 2005

35. Zhu and Bingham 2008, Fink, Foo, and Warren 2009

36. Wilson and Bingham 2008.

37. Kugler 1986.

38. Gibson 1979, Golonka and Wilson 2012, Turvey et al. 1981.

REFERENCES

Alexander, David M., and Gordon G. Globus. 1996. 'Edge-of-Chaos Dynamics in Recursively Organized Neural Systems,' in Fractals of Brain, Fractals of Mind: In Search of a Symmetry Bond, ed. E. MacCormac and Maxim I. Stamenov (Amsterdam, Netherlands: John Benjamins Publishing Company), pp. 31–73.

Bingham, Geoffrey P. 1988. 'Task-Specific Devices and the Perceptual Bottleneck,' Human Movement Science 7(2–4): 225–64.

Bingham, Geoffrey P. 2004. 'Another Timing Variable Composed of State Variables: Phase Perception and Phase Driven Oscillators,' in Advances in Psychology, Vol 135: Time-to-Contact, ed. Heiko Hecht and Geert Savelsbergh (Amsterdam, Netherlands: Elsevier Science Publishers), p. 421–442.

Bingham, Geoffrey P. 2004. 'A Perceptually Driven Dynamical Model of Bimanual Rhythmic Movement (and Phase Perception),' Ecological Psychology 16(1): 421–42.

Brooks, Rodney A. 1997. 'Intelligence without Representation,' in Mind Design 2: Philosophy, Psychology, Artificial Intelligence (Cambridge, MA: The MIT Press), pp. 395–420.

Carlson, Thomas A., George Alvarez, Daw-an Wu, and Frans A. J. Verstraten. 2010. 'Rapid Assimilation of External Objects into the Body Schema,' Psychological Science 21(7): 1000–1005.

Charles, Eric P. 2011a. A New Look at New Realism: The Psychology and Philosophy of E. B. Holt Piscataway, NJ: Transaction Publishers.

Charles, Eric P. 2011b. 'Seeing Minds in Behavior: Descriptive Mentalism.' Review of General Psychology 15(3): 267–276.

Charles, Eric P., and Martin Dege. 2008. 'The Beyond Must Also Be Useful: The Burden of Alternative Approaches,' Integrative Psychological and Behavioral Science 42(2): 194–199.

Clark, D. D., and L. Sokoloff. 1999. 'Circulation and Energy Metabolism of the Brain,' in Basic Neurochemistry: Molecular, Cellular and Medical Aspects, ed. G. J. Siegel, B. W. Agranoff, R. W. Albers, S. K. Fisher and M. D. Uhler (Philadelphia: Lippincott), pp. 637–70.

Fajen, Brett R., and William H. Warren. 2003. 'Behavioral Dynamics of Steering, Obstable Avoidance, and Route Selection,' Journal of Experimental Psychology: Human Perception and Performance 29(2): 343–62.

Feldman, Anatol G. 1966. 'Functional Tuning of the Nervous System with Control of Movement or Maintenace of a Steady Posture, Ii: Controllable Parameters of the Muscles.' Biophysics 11: 565–78.

Feldman, Anatol G. 2011. 'Space and Time in the Context of Equilibrium-Point Theory,' Wiley Interdisciplinary Reviews: Cognitive Science 2(3): 287–304.

Fink, Philip W., Patrick S. Foo, and William H. Warren. 2009. 'Catching Fly Balls in Virtual Reality: A Critical Test of the Outfielder Problem,' Journal of Vision 9(13): 14.

Gibson, Eleanor Jack, and Anne D. Pick. 2000. An Ecological Approach to Perceptual Learning and Development. New York: Oxford University Press.

Gibson, James Jerome. 1979. The Ecological Approach to Visual Perception Boston, MA: Houghton, Mifflin and Company.

Golonka, Sabrina, and Andrew D Wilson. 2012. 'Gibson's Ecological Approach–a Model for the Benefits of a Theory Driven Psychology.' Avant 3: 40–53.

Guterstam, Arvid, Valeria I. Petkova, and H. Henrik Ehrsson. 2011. 'The Illusion of Owning a Third Arm,' PLoS ONE 6(2): e17208.

Haken, H., J. A. Kelso, and H. Bunz. 1985. 'A Theoretical Model of Phase Transitions in Human Hand Movements,' Biological Cybernetics 51(5): 347–56.

Heft, Harry. 2001. Ecological Psychology in Context : James Gibson, Roger Barker, and the Legacy of William James's Radical Empiricism. Mahwah, N.J.: L. Erlbaum.

Holt, Edwin B. 1936. 'Eight Steps in Neuro-Muscular Integration,' in Problems of Nervous Physiology and of Behavior (Tiflis: Georgian Branch, Academy of Sciences, USSR), pp. 25–36.

Kugler, P.N. 1986. 'A Morphological Perspective on the Origin and Evolution of Movement Patterns,' Motor development in children: Aspects of coordination and control, pp. 459–525.

Lashley, K. S. 'Basic Neural Mechanisms in Behavior,' Psychological Review 37(1): 1–24.

MacCorquodale, Kenneth, and Paul E. Meehl. 1948. 'On a Distinction between Hypothetical Constructs and Intervening Variables,' Psychological Review 55(2): 95–107.

Mason, Paul H. 2010. 'Degeneracy at Multiple Levels of Complexity,' Biological Theory 5(3): 277–288.

Miller, David B. 1997. 'The Effects of Nonobvious Forms of Experience on the Development of Instinctive Behavior,' in Evolving Explanations of Development: Ecological Approaches to

Organism-Environment Systems (Washington, DC: American Psychological Association), pp. 457–507.

Powers, William T. 1973. Behavior: The Control of Perception. Oxford: Aldine.

Prinz, A. A., D. Bucher, and E. Marder. 2004. 'Similar Network Activity from Disparate Circuit Parameters.' Nat Neurosci 7(12): 1345–1352.

Raichle, M. E., A. M. MacLeod, A. Z. Snyder, W. J. Powers, D. A. Gusnard, and G. L. Shulman. 2001. 'A Default Mode of Brain Function,' Procedings of the National Academy of Science, U.S.A. 98(2): 676–682.

Runeson, Sverker. 1977. 'On the Possibility of 'Smart' Perceptual Mechanisms,' Scandinavian Journal of Psychology 18(1): 172–179.

Smith, Laurence D. 1986. Behaviorism and Logical Positivism : A Reassessment of the Alliance. Stanford, CA: Stanford University Press.

Sporns, Olaf. 2011. Networks of the Brain. Cambridge, MA: MIT Press.

Tonneau, Francois. 2011. 'Holt's Realism: New Reasons for Behavior Analysis,' in A New Look at New Realism: The Psychology and Philosophy of E. B. Holt, ed. Eric Charles (Piscataway, NJ: Transaction Publishers), 33–55.

Turvey, M. T., R. E. Shaw, E. S. Reed, and W. M. Mace. 1981. 'Ecological Laws of Perceiving and Acting: In Reply to Fodor and Pylyshyn,' Cognition 9(3): 237–304.

Wilkie, Richard M., and John P. Wann. 2005. 'The Role of Visual and Nonvisual Information in the Control of Locomotion,' Journal of Experimental Psychology: Human Perception and Performance 31(5): 901–911.

Wilson, Andrew D. 2011. 'What Does the Brain Do, Pt 2: The Fast Response System,' in Notes from Two Scientific Psychologists. Available at: http://psychsciencenotes.blogspot.com/2011/08/what-does-brain-do-pt-2-fast-response.html.

Wilson, Andrew D, and Sabrina Golonka. 2013. 'Embodied Cognition Is Not What You Think It Is,' Frontiers in Psychology 4(58): 1–13.

Wilson, Andrew D., and Geoffrey P. Bingham. 2008. 'Identifying the Information for the Visual Perception of Relative Phase,' Perception and Psychophysics 70(3): 465–476.

Zhu, Qin, and Geoffrey P. Bingham. 2008. 'Is Hefting to Perceive the Affordance for Throwing a Smart Perceptual Mechanism?' Journal of Experimental Psychology: Human Perception and Performance 34(4): 924–943.