3

Is the Next Frontier in Neuroscience a 'Decade of the Mind'?[1]

*Jacqueline A. Sullivan**

3.1 Introduction

In 2007, ten world-renowned neuroscientists, including James Albus, George Bekey, John Holland, Nancy Kanwisher, Jeffrey Krichmar, Mortimer Mishkin, Dharmendra Modha, Marcus Raichle, Gordon Shepherd, and Giulio Tononi advocated in a letter published in *Science* for "a major national research initiative called 'A Decade of the Mind'" (Albus et al. 2007, 1321). Their contention was that, despite the successes of the *Decade of the Brain*, "a fundamental understanding of how the brain gives rise to the mind [was] still lacking" (ibid.). They identified four areas of research to be the focus of this new decade including: (1) healing mental disorders; (2) understanding "aspects of mind believed to be uniquely human" including "the notion of self, rational thought processes, theory of mind, language and higher order consciousness"; (3) "enriching the mind through education"; and (4) "modeling the mind by means of computational models and artificial intelligence" (ibid.). The proposed decade was to be "transdisciplinary and multi-agency in its approach," incorporating insights from neuroscience, medicine, cognitive neuroscience, psychology, computer science, engineering, mathematics, robotics, systems biology, cultural anthropology and social science (ibid.).

Six years have passed since the publication of the letter, and only one direct response to it has appeared in the scientific literature. However, that response, provided by the German psychiatrist Manfred Spitzer, did not take up the most interesting issues raised by the proposal. Spitzer interpreted the authors as merely suggesting that, given that research

C. T. Wolfe (ed.), *Brain Theory*

during the decade of the brain had been directed primarily at understanding the mechanisms of "perception and motor control," it was time for "investigators in systems neuroscience" to "turn their attention and a powerful arsenal of methods towards what traditionally were regarded as 'mind-functions'" (Spitzer 2008). In other words, Spitzer interpreted the authors of the proposal as merely advocating for the broader application of those investigative strategies prevalent in cognitive and systems neuroscience during the 1990s to more complex phenomena such as theory of mind, the self and higher-order consciousness. While it may just be that Albus and colleagues' call for a decade of the mind was supposed to involve the extension of imaging technologies to study higher-order mental functions and dysfunctions, that hardly seems like the "paradigm-shifting progress" and interdisciplinary call to arms that the proposal was intended to instigate. Thus, the only response to the proposal in the review literature downplayed what was truly interesting about it, namely, that a group of world-renowned neuroscientists had acknowledged that a major change in how neuroscientists study the mind-brain relationship was needed.

The aim of this chapter is to remedy this oversight by addressing what I regard as a set of interesting questions that the proposal for the decade of the mind prompts – questions that I believe, when answered, enable a more cogent case for a decade of the mind and provide insight into how to implement the proposal in neuroscientific practice.[2] First, what is "the mind" that is supposed to be the target of this new decade? Second, what was missing during the decade of the brain that would prompt the need for a separate decade of the mind? Third, what should this new decade look like?

3.2 What is 'the mind'?

One fundamental aspect of being human is that we learn to adopt what Daniel Dennett (1987) has dubbed "the intentional stance." Specifically, we come to believe that human beings as well as some non-human animals have some special quality – a *mind*, *consciousness*, *awareness* – that other kinds of things – rocks, stars, and trees – lack. We describe ourselves as having beliefs, desires, feelings, and intentions. We ascribe similar internal states to other human beings and some non-human animals. We appeal to these states to explain our own and others' behaviors. This conceptual-explanatory framework, which has been elevated to the status of a theory by cognitive psychologists and philosophers (for example Churchland 1981), plays a fundamental role in introspection and in our development of a concept of self and a personal identity. It also has and continues to serve as a principal basis for philosophical thinking about the nature of consciousness, subjective experience, mental causation, cognition and knowledge, ethics, the nature of the self, personal identity, and mental disorders (e.g., delusions as false beliefs). Insofar as components of this framework may vary across cultures, races, and religions, and such differences may be put forward to explain social, economic, political, historical, and cultural phenomena, it plays a fundamental role in the humanities and in the social sciences,[3] including those areas of science the proponents of the Decade of the Mind (DoM) would like to be involved in this new initiative.

The DoM proposal may be misunderstood as a call for a revival of folk psychology as a conceptual-explanatory framework for at least two reasons. First, the authors define what they mean by "mind" ostensively rather than discursively, and the terms put forward by them including "mind," "higher-order consciousness," "the notion of self," "rational thought processes," and "mental disorders" overlap with the mixed ontology of entities and activities that function as causes of behavioral phenomena in the folk-psychological worldview. A second and related reason why the proposal appears sympathetic to folk-psychological theory is that the initiative is supposed to be "transdisciplinary" – it is intended to instigate interdisciplinary interactions among neuroscientists and practitioners coming from more disparate research areas in which "talk about belief is ubiquitous" (Dennett 1987, 13). The success of the decade as conceived by the DoM authors purportedly requires crosstalk among investigators coming from areas of science that take folk psychology seriously and areas of science that do not. It would be disingenuous for investigators to attempt to include practitioners who regularly appeal to folk-psychological explanations in the DoM initiative while at the same time devaluing this approach. Third, one aim of the initiative is "educating the general public on legal and ethical issues involving the brain and the mind" (Albus et al. 2007, 1321). Effectively communicating scientific results to the general public whose primary understanding of the mind, mental disorders, consciousness, and the self is rooted in folk psychology seems to require some attempt to take that conceptual-explanatory framework seriously.[4]

Other facts, however, may be taken to indicate that the proponents of the decade of the mind are not interested in the conceptual explanatory framework of ordinary folk.[5] First, what we find when we look across the contemporary neurosciences are multiple distinct areas of science ranging from molecular genetics to behavioral neuroscience and "a plurality of incompletely articulated and partially contradictory, partially supplementary theories and models" (Wimsatt 2007, 180) that are all purportedly directed at understanding the neural mechanisms of consciousness and cognition as well as the brain dysfunctions that give rise to mental disorders. Folk psychology as a conceptual-explanatory framework is notably absent from the diverse array of explanations that the contemporary neurosciences yield for these phenomena. This makes sense because a primary aim of the biological and physical sciences is to move beyond folk understandings and explanations of phenomena to discover their *physical* mechanisms.[6] The mental states of the folk do not fit into the world in any interesting way; they are not parts of the ontological hierarchy that runs the gamut from molecules to behavior. Secondly, given that the investigators who authored the DoM proposal work (or did work) in a variety of areas of the cognitive and neurosciences including: artificial intelligence (e.g., Modha), engineering (e.g., Albus, Holland), computer science and robotics (e.g., Bekey, Holland), cognitive science (e.g., Krichmar) cognitive neuroscience (e.g., Kanwisher, Mishkin, Tononi) neuropsychology (e.g., Mishkin), psychiatry (e.g., Tononi), neurobiology (e.g., Shepherd), and neurology (e.g., Raichle), it is reasonable to conclude that they are endorsing the conceptual-explanatory framework of cognitive science and with it the *cognitive notion of mind*.

According to the cognitive notion, mind is the total set of an organism's cognitive states and processes that are causally responsible for, but not identical to, its overt behavior. The conceptual-explanatory framework of cognitive science contains a basic set of assumptions that practitioners coming from those diverse fields represented in the DoM proposal share in common. These assumptions include most basically that (1) human beings have specific kinds of cognitive capacities, (2) these capacities involve "representational structures and processes," and insofar as (3) these representational structures and processes carry information about what they represent (4) thinking of the mind as an information-processing device is a fruitful analogy (von Eckhardt 1993, 1). This set of assumptions is clearly not without its problems, because the "representational structures" of cognitive science that "carry information" are purportedly *mental* states. What relationship mental states bear to neural states and the question of how neurons can assume the representational capacities normally ascribed to them remain subjects of scientific and philosophical debate (for example Bechtel 2008; DeCharms and Zador 2000). So, proposing to revive the cognitive notion of mind in neuroscience has the potential to revitalize a host of thorny problems that the decade of the brain did not resolve, but left behind. However, a persuasive case can be made for reviving the cognitive notion of mind in contemporary neuroscience. In order to make this case, I regard it as first relevant to consider how philosophers have conceived of the relationship between folk psychological and cognitive psychological explanations of cognitive phenomena, so I can be very clear about precisely what aspects of the conceptual-explanatory framework of cognitive science that go hand-in-hand with the cognitive notion of mind need to be revived and where.

Some philosophers, most notably Jerry Fodor (1975), equate the mental states of the folk with the so-called "propositional attitudes," statements of the form "*S* ______ *that P*." where *S* is the subject (e.g., Wayne, Mandy, Tim), *P* is a proposition (e.g., "All mental states are brain states.") and "_______" may be filled in with any one of a number of verbs expressing a disposition towards that proposition (e.g., believing, hoping). These same philosophers have implied that insofar as cognitive psychological explanations make reference to mental states, and mental states are nothing over and above propositional attitudes, cognitive psychological explanations and folk psychological explanations are similar.

Robert Cummins (1983) most notably has responded that such claims are based on a misunderstanding of the nature of psychological explanation, which has as its aim to explain the cognitive capacities of organisms by *functional analysis*. Explanation by functional analysis, according to Cummins, "does not traffic in explicit propositional attitudes" and "does not involve intentional characterization" (Cummins 1983, 82). Rather, the aim of functional analysis is to ascribe cognitive capacities to organisms and identify the sub-capacities that enable the realization of those capacities. To take a simple example, organisms can remember things; they have the cognitive capacity of memory. However, organisms remember different kinds of things – not only how to perform certain procedures but also that certain declarable facts obtain. In other words, memory may be broken down into at least two sub-capacities, which are not propositional attitudes and are not necessarily characterized in terms of them. Ideally, the concepts of "procedural memory" and "declarative memory" designate sub-capacities that can be ascribed to divisible systems in the brains of organisms that exhibit the cognitive capacity. Neuroscience has in fact shown this to be the case, with declarative memory being subserved by structures in the medial temporal lobe (e.g., the perirhinal and parahippocampal cortices, the entorhinal cortex, and the hippocampus) and procedural memory being subserved by the striatum, basal ganglia, and cerebellum. Cognitive psychological explanations thus differ from folk psychological explanations because they ultimately require investigators to take the *physical stance* towards their objects of inquiry.

Although Cummins is correct about the basic aims of explanation in cognitive psychology, I think the claim that cognitive psychological explanation "does not involve intentional characterizations" is false. For it is based on the false assumption that taking the intentional stance towards an organism and ascribing propositional attitudes to it is the same thing. But why equate them? Even the common person deploys a mixed ontology of mental states (anger, happiness, attention, memories) and mental processes (believing, desiring, wanting, remembering, attending) in order to explain and predict behavior. Thus, it is incorrect to claim that when ordinary folk adopt the intentional stance they are "trafficking" exclusively "in propositional attitude ascriptions." Propositional attitude theory offers just one conceptual-explanatory framework for explaining the behavior of the folk and understanding what they do when they explain their own and others' behaviors. However, when cognitive scientists "talk ... about human cognitive activities" they may be interpreted simply as "speak[ing] about mental representations and ... posit[ing] a level of analysis wholly separate from the biological or neurological" level (Gardner 1985, 6). Propositional attitude theory does not capture what they do, but claiming that they take the intentional stance towards their objects of inquiry, namely human cognizers, does. If we grant this, then, contrary to what Cummins claims, it *is* possible that cognitive psychologists in particular and cognitive scientists more generally take *the intentional stance* towards those organisms whose capacities they are interested in explaining. Let's call this a modified version of the intentional stance insofar as the achievement of their explanatory aims requires the positing of abstract mental states and mental processes. Perhaps sometimes this includes propositional attitude ascriptions, but that is not a requirement. This brings us to the question of *when* in the process of doing science cognitive psychologists take the intentional stance, given that, as Cummins rightly claims, explanations by functional analysis, in the final analysis, do not involve intentional characterizations – a point I, and recent proponents of mechanistic explanation (e.g., Bechtel 2008; Craver 2007; Piccinini and Craver 2011) think he is correct about.

One answer to this question that is implicit in Cummins (1983) and explicit in Fodor (1968) is the idea that the intentional stance is involved *only in the earliest stages of* psychological explanation. Cummins (1983) refers to two kinds of functional analysis: *dispositional analysis*, which involves breaking down a capacity into abstract sub-capacities that are ascribed at the level of the whole intact organism and *componential analysis*, which requires that these sub-capacities be ascribed to the internal physical components (e.g., brain structures, synapses, neurons) of the organism that exhibits them. Whereas dispositional analysis requires that cognitive psychologists take the intentional stance, componential analysis does not.[7] These two types of analysis correspond directly to what Fodor (1968) identifies as two phases of psychological explanation. In the first phase, "the psychologist is seeking functional characterizations of psychological constructs" and "the criteria employed for individuating such constructs are based primarily on hypotheses about the role they play in the etiology of behavior" (Fodor 1968, 107–108).[8] The constructs employed by psychologists at this stage of explanation are hypothetical and abstract. Determining whether these constructs correspond to actual divisions in the brain is the second phase of psychological explanation. During this phase the psychologist aims to identify "those biochemical systems that do, in fact, exhibit the functional characteristics" of interest and she "'looks inside' to see whether or not the nervous system does in fact contain parts capable of performing the alleged functions" (ibid., 109). Given Cummins' and Fodor's claims, then, it is reasonable to conclude that whereas the first stage of psychological explanation requires that investigators take the intentional stance towards their experimental subjects, the second phase requires them to take the physical stance and assume that an organism's behavior has internal physical (e.g., neural, biochemical) causes.[9]

I want to suggest, however, an even earlier and more important role for the intentional stance in cognitive psychology in particular and in the neurosciences of cognition more generally. As Fodor claims, in the early stages of explanation psychologists posit "hypothetical constructs." In cognitive psychology in particular and cognitive science more generally, a construct is a postulated capacity or attribute of an organism that may also be a target of psychological explanation. Intelligence, working memory, innateness, and spatial memory are all examples of constructs. A construct typically originates, "with a vague concept which we associate with certain observations" (Cronbach and Meehl 1955, 286), and this vague concept serves as basis not only for theory building in psychology (as well as other social sciences) but also for designing experiments to study cognitive capacities of interest, and in the case of the neurosciences of cognition, localizing them in the brain and determining the neural, synaptic, cellular, and molecular mechanisms that give rise to them. While the origin of such vague concepts may be the conceptual-explanatory framework of folk psychology, investigators also rely on definitions already available in the scientific literature – particularly definitions that are commonly deployed in that same field of research.

When a cognitive psychologist goes into the laboratory, she will aim to design an experimental paradigm or cognitive task in order to investigate the cognitive capacity of interest. An experimental paradigm is roughly "a standard set of procedures for producing, measuring and detecting" a cognitive capacity "in the laboratory" that "specifies how to produce" that capacity, "identifies the response variables to be measured during pre-training, training, and post-training/testing" and includes instructions on "how to measure [those response variables] using equipment that is designed for this purpose" (Sullivan 2010b, 266). It also specifies how to detect the cognitive capacity "when it occurs, by identifying what the comparative measurements of the selected response variables have to equal in order to ascribe" that capacity to a subject (ibid.). Ideally, the investigator will aim to design an experimental paradigm capable of reliably individuating the cognitive capacity of interest that the construct purportedly designates. Individuating cognitive capacities requires the development of experimental paradigms or cognitive tasks that are *reliable* for this purpose. Ensuring the reliability of an experimental paradigm requires, I am claiming, that an investigator adopt the intentional stance towards her subjects. This is supported by the fact that when cognitive neuroscientists design experimental paradigms they regard it as important to consider what may ultimately go on inside the head of a hypothetical subject when that subject is trained and tested in that paradigm. The vast majority of cognitive neuroscientists, particularly those with training in cognitive psychology, assume that there is such a thing as mental function, that it "is composed of distinguishable fundamental processes" and that "these processes can be selectively engaged by properly designed experimental task manipulations" (Carter et al. 2009, 169). In light of this assumption, some cognitive neuroscientists engage in rigorous "theoretically guided task analys[es]" intended to provide "a clear specification of the processes thought to be engaged by an experimental task and how these processes will be influenced by the variables to be manipulated in the experiment" (ibid., 169).[10] Designing tasks that successfully individuate a cognitive function and allow for its localization in the brain requires a significant amount of ingenuity. If an investigator fails to consider the possible mental states of her hypothetical subject or she neglects to itemize other hypothetical or actual processes that may be involved in the execution of the task (e.g., attention, working memory), then the task that she designs will likely be unreliable for individuating the cognitive capacity of interest to her. The task will lack *construct validity* if there is a discrepancy between the hypothetical construct that the task was intended to measure and those cognitive processes actually measured.

Investigators across the contemporary cognitive and neurosciences have substantial freedom to produce, detect, and measure cognitive functions using the experimental paradigm or task that they take to be most reliable for achieving their investigative aims. Not all investigators will agree that a particular experimental task or paradigm is subject to one exclusive task analysis or that it measures a discrete cognitive function. In fact, disagreements about the potential functions that play a role in the execution of a given task may prompt revisions to that task and/or the development of new tasks. For example, the Stroop task,[11] which was for a long time widely thought to individuate the cognitive capacity of *selective attention*, has also been described as measuring *response inhibition* and *context processing*. Precisely what cognitive function it individuates remains a subject of debate (Perlstein et al. 1998; Cohen et al. 1999; Barch et al. 2004). Many tasks have prompted similar debates. However, it is well recognized that developing tasks that successfully individuate cognitive capacities is an iterative, trial-and-error process. What I am claiming is that in cognitive psychology and cognitive neuroscience it is an iterative process that intimately involves the intentional stance.

Insofar as cognitive neuroscientists adopt the intentional stance in the experimental context, we can be rest assured that the cognitive notion of mind is alive and well in contemporary cognitive neuroscience. If this is correct, then the proposal for a decade of the mind in neuroscience must be directed at some other target. For a variety of reasons, I think the primary target is *cognitive neurobiology*. First, the authors of the proposal stress that the aim of this new decade is to move past the achievements of the decade of the brain, which "focused on neuroscience and clinical applications" (Albus et al., 2007, 1321). It is reasonable to suspect that they have low-level neuroscience in mind here, which focused primarily on the development of treatments for mental disorders and the identification of the cellular and molecular mechanisms of learning and synaptic plasticity during the decade of the brain. As John Bickle claims, the real "'revolution'" in neuroscience during that decade occurred in "cellular physiology and molecular biology" (2003, 2) rather than in cognitive neuroscience. Bickle (2003, 2006) also has correctly pointed out that investigators in cognitive neurobiology do not make reference to mental functions or to the mind. Rather, they take the physical stance towards their experimental subjects, they operationally define cognitive capacities in terms of observable changes in behavior and then look directly into the brain to determine the cellular and molecular activity implicated in the production of those changes in behavior (see also Sullivan 2009, 2010a, 2010b; Sweatt 2009). Third, the authors of the proposal stress that this new decade should be "transdisciplinary and multi-agency" in its approach. In other words, it requires scientists in all the areas that are to be involved in the initiative to take seriously the cognitive notion of mind. This makes sense because determining how the brain gives rise to the mind will be a non-starter if investigators in the business of localizing cognitive functions in the brain (i.e., cognitive neuroscientists) cannot effectively communicate with investigators who are discovering their cellular and molecular mechanisms (i.e., cognitive neurobiologists).

The call for "paradigm-shifting" progress in neuroscience conjures up Kuhn's (1962) notion of a paradigm and the problem of incommensurability: investigators working in radically opposed scientific traditions that are associated with different foundational assumptions and investigative approaches have trouble communicating to the extent that it sometimes seems as if they are living in different worlds. In fact, cognitive neuroscience and cognitive neurobiology do indeed emanate from separate and radically opposed historical traditions, and appreciating this history enables us not only to recognize potential obstacles to this new decade but also to appreciate that overcoming such obstacles requires reviving the mind in the experimental context in cognitive neurobiology.

3.3 Why revive the cognitive notion of mind?

Cognitive science as we know it today began to emerge in the second half of the 20th century when a group of scientists representing a diverse array of disciplines including mathematics, computer science, neurophysiology, and psychology began to search for methods to understand the mind, the brain, and behavior that were alternatives to the then-dominant experimental psychology of behavior (Gardner 1985). The "behaviorists," including J.B. Watson, E.L. Thorndike, Edwin Guthrie, Clark Hull, and later B.F. Skinner, aimed to develop a scientific psychology that was on a par with the physical sciences. The behaviorists took the achievement of this aim to require the rejection of 19th-century introspective psychology,[12] its method of introspection and "its subject matter *consciousness*" (Watson 1924[1970], 2). In its place they put forward a science that was supposed to rely exclusively on publically verifiable methods, explaining behavior solely by appeal to stimuli and responses. They supposed that the experimental learning paradigms of classical and operant conditioning – two standard sets of procedures for producing, measuring, and detecting forms of associative learning in the laboratory – could be used to investigate the causes of an organism's behavior without investigators needing to concern themselves with what was going on inside those organisms' heads. Another way of putting it is that they rejected "the intentional stance" in psychology; it played no role in the process of designing and implementing experimental learning paradigms or in explaining the data.

Behaviorism and the experimental paradigms of classical and operant conditioning came under attack for a number of reasons. First, it was difficult for experimental psychologists to reliably isolate learning conditions that were favorable exclusively to stimulus-stimulus and stimulus-response type explanations because it was often difficult to identify precisely what the independent variables were. For example, as J.J.C. Gibson (1960) and Charles Taylor (1964) claim, it was difficult to assess to which stimuli or which aspects of those stimuli organisms trained in classical and operant conditioning paradigms were actually responding. Another limitation was that the experimental paradigms were neither ecologically nor externally valid. The highly artificial conditions of the laboratory, which included raising organisms in impoverished environments and depriving them of an experiential history, was regarded by critics as an obstacle to generalizing from learning in the laboratory to learning in the world (see e.g. Hinde 1973, Lorenz 1965). A more severe problem was that learning in classical and operant conditioning paradigms required the repetition and contiguity of stimuli or stimuli and responses. However, as the psychologist Karl Lashley argued in his lecture at the Hixon Symposium, more complex forms of learning (e.g., Kohler 1947) did not appear to conform to this associative learning model. As the historian Howard Gardner claims, "so long as behaviorism held sway [...] questions about the nature of human language, planning, problem solving, imagination, and the like could only be approached stealthily and with difficulty, if they were tolerated at all" (Gardner 1985, 12).

Physiologists trained in the behaviorist tradition, however, were not moved by Lashley's worries about the limitations of associative learning theory for explaining complex cognitive capacities (e.g., learning a language or learning how to play a sport or an instrument). Some also doubted the reliability of Lashley's ablation experiments, which he used to refute the widely accepted idea that memory traces were stored in the brain in those neurons activated during training in associative learning paradigms. In a seminal review paper in 1968, Eric Kandel and W. Alden Spencer sought to discredit Lashley in claiming that his "experimental techniques and his conclusions had been seriously questioned" (1968, 67). Feeling that they had successfully turned back the only challenge to the idea that "the synapse" plays "a crucial role in information storage" (1968, 66), they emphasized the importance of the repetition and contiguous presentation of stimuli to forge new synaptic connections in the brain.

Kandel and Spencer essentially echoed the ideas of the neobehaviorist Donald Hebb. In 1949, Hebb, borrowing insights from behaviorism, gestalt psychology and physiology, suggested that learning and memory are achieved by physiological changes in brain synapses. Hebb claimed that when two cells, A and B, which communicate under normal conditions, undergo a period of repeated and concurrent activation, as may happen during classical or operant conditioning, the result will be a strengthening of the connection between the two cells. According to Hebb, this strengthening is reflected in a subsequent change in the way the one neuron excites the other (Hebb 1949). The crux of this postulate, often referred to as "Hebb's rule," is that each associative learning event is accompanied by the brief associated activation of two neurons that comprise a synapse, which together, effectively store information in the form of a physiological change at that synapse.

Although Hebb's postulate was attractive, at the time of its introduction, no plausible candidates for a neural mechanism of associative learning that satisfied his described conditions had been located in the mammalian brain. However, in 1966, in the context of investigating the physiology of the dentate gyrus in the hippocampus of the adult anesthetized rabbit, Terge Lømo observed an artificially induced physiological equivalent of a strengthening in synaptic efficacy (Lømo 2003). This discovery of a "long-lasting potentiation" in area CA1 of the rabbit hippocampus *in vivo* led to Lømo's famous publication with Tim Bliss in 1973 in which they described the phenomenon of long-term potentiation (LTP), which instantiated all of the features of the mechanism of associative learning that Hebb (1949) had described (see Craver 2003).

Kandel's seminal research on the cellular and molecular mechanisms of learning and memory in the sea mollusk *Aplysia Californica* in combination with Tim Bliss and Terge Lømo's (1973) discovery of long-term potentiation (LTP) in area CA1 of the rabbit hippocampus *in vivo* form the cornerstones of modern cognitive neurobiology. Although the paradigms of classical and operant conditioning and associative learning theory were rejected by those attendees of the Hixon symposium who founded modern cognitive science, they found a good home in cognitive neurobiology where versions of the two learning paradigms are widely used to this day in conjunction with electrophysiology experiments that are used to induce LTP.

Paradigms intended to investigate more complex cognitive functions also have been introduced into the cognitive neurobiological literature. Social recognition memory paradigms (Bickle 2006; Sullivan 2009) and the Morris water maze (e.g., Craver and Darden 2001; Craver 2007; Sullivan 2010a, 2010b) are some widely celebrated examples. However, their introduction has led to certain kinds of problems in part because of the failure on the part of cognitive neurobiologists to take the intentional stance. Specifically, investigators in cognitive neurobiology are not interested in "what" cognitive capacity is operative when an organism is trained in an experimental learning paradigm. They are satisfied just so long as they can use those paradigms to produce robust behavioral effects in which they can pharmacologically or genetically intervene. When they train organisms in experimental learning paradigms, they do not assume that these organisms have minds; they do not take the intentional stance, but the physical stance.

One representative example is the hidden condition of the Morris water maze. The water maze is an open field maze consisting of a large circular pool filled with opaque water. The pool is placed in a room containing a discrete set of fixed distal visual cues. In the hidden condition of the water maze, a silvery-white platform is placed just beneath the water's surface so as to be undetectable to a rodent placed in the pool. During training in the hidden condition, the location of the platform remains fixed across trials and the placement of the rat in the pool varies randomly with respect to the four cardinal positions (i.e., N, S, E, W). When a rat is placed into the pool, it will attempt escape, and thus swim about the pool. On each training trial, the swim path of the animal in the maze, the length and direction of the angle of that path, and the time it takes it to find the platform ("escape latency") are measured. A significant decrease in the amount of time it takes the animal to find the hidden platform across training trials is taken to indicate that the rat has learned the location of hidden platform solely on the basis of the distal room cues. Morris originally referred to this set of behavioral effects observed in the water maze as "place learning," which was intended to capture the idea that the rats learned the place of the hidden platform solely on the basis of the distal room cues, rather than by stimulus-stimulus or stimulus-response associations.

An historical analysis of the Morris water maze, however, reveals that over a 30-year time span, across the experimental and review literature in cognitive neurobiology, the term used to designate the phenomenon under study in the hidden condition of the maze oscillated. Candidate terms included place learning, place navigation, spatial learning, spatial memory, spatial navigation, water maze navigation, and water maze performance (Sullivan 2010b). Such oscillations suggest not only that investigators were unclear what cognitive function was under study in the water maze, but also that over a 30-year time span, only slight efforts were directed at achieving clarity. This makes sense given that cognitive neurobiologists are not concerned with "what" rodents trained in the water maze learn or what cognitive functions are involved in the production of the behavioral effects. After all, they do not assume organisms have minds, nor do they take the intentional stance towards their experimental subjects. However, this lack of concern is an impediment to explanatory progress because to discover the mechanisms of a cognitive function, it is necessary to know what the function is (see for example D'Hooge and De Deyn 2001).

One of the reasons the water maze is such an interesting experimental paradigm is that it teaches us lessons about the challenges of the scientific study of cognitive functions. The first lesson is that when an experimental paradigm is designed, establishing its reliability for individuating a discrete cognitive function requires a consideration of "what" an organism trained in the paradigm is learning. This suggests that investigators must appeal to a cognitive understanding of the mind, take seriously the potential mental states and information processes of the whole intact organism and engage in a thorough task analysis, much like cognitive neuroscientists do. Secondly, when an experimental paradigm is used to reliably produce a discrete set of behavioral effects, investigators often assume that they have individuated the cognitive function of interest, and the search for the systems, synaptic, cellular, and molecular mechanisms productive of those behavioral effects then begins. However, cognitive functions are not identical to behavioral effects that result from training an organism in an experimental paradigm. The causes of the changes in behavior likely include many more changes in internal states and processes than are captured by the term designating the cognitive function under study in the paradigm (see for example Taylor 1964). This is why taking the intentional stance is not only important in the context of designing experimental learning paradigms, but it is also fundamental when such paradigms are being implemented in the laboratory and the results of these experiments are being interpreted.

Currently, the possible limitations of all the experimental paradigms used in cognitive neurobiology are likely being missed precisely because cognitive neurobiologists fail to take an organism's mental states seriously once they have identified an experimental paradigm that seems to produce robust behavioral effects. However, the problems that arise from failing to raise questions about "what" an organism is learning, remembering, or doing, and what representational processes are involved, are an impediment to individuating discrete explanatory targets and identifying their mechanisms. To eliminate such problems and thus answer the question of where we need to revive the mind in contemporary neuroscience, I think the answer is that the cognitive notion of mind and the intentional stance ought to play a fundamental role in the experimental context in cognitive neurobiology when investigative strategies are being designed and/or implemented.

Ironically, the proposal for a decade of the mind in contemporary neuroscience is roughly similar to the original call to revive the mind made by the diverse array of scientists who attended the Hixon Symposium in 1948 where the seeds of contemporary cognitive science were sown. The scientific backgrounds of those members of that original group are also similar to the backgrounds of those who gathered at George Mason University in 2007 to discuss the need for a new decade of the mind. Thus, in answer to the question of why a world-renowned group of cognitive scientists would call for a new decade in contemporary neuroscience, I think we need only look to the reasons put forward by the attendees of the Hixon symposium who recognized that the only way to overcome the limitations and problems with the then-current conceptual-explanatory framework for understanding learning and memory was to revive the mind in psychology and bring back the intentional stance.[13]

3.4 What should a decade of the mind look like?

Proponents of decade of the mind argue that this new decade should be "broad in scope" and "transdisciplinary in nature." However, they offer no positive proposals with respect to what form such interdisciplinary interactions should take. Given the insights revealed from my analysis of the Morris water maze, I have indicated that one appropriate venue for collaboration is the experimental context. More specifically, I am claiming that to increase the reliability of experimental paradigms for individuating cognitive functions, practitioners from a variety of different areas of the mind-brain sciences, including but not limited to: cognitive psychologists, cognitive neuroscientists, experts in animal behavior, computational scientists, and molecular and cellular cognitive neurobiologists, should combine forces to develop and implement experimental paradigms in the laboratory. Furthermore, extensive dialogue across research teams and disciplines using the same experimental paradigm should be on-going. Such interdisciplinarity makes good sense for several reasons. First, in any area of science that uses whole intact organisms, one must be privy to the fact that a variety of different processes – molecular, cellular, synaptic, network, systems, representational, informational, and behavioral – co-occur simultaneously. Second, different investigators, given different areas of expertise, have different explanatory interests, and they face different obstacles in developing and implementing experiments that work for their distinct explanatory purposes. These explanatory interests and obstacles require a forum within which solutions may be located and the impact of such solutions on the phenomena under study (i.e., cognitive capacities) may be considered.

It makes sense for such "perspectival pluralism" (Giere 2010; Wimsatt 2007) to be implemented in the context of experimentation rather than the context of explanation in neuroscience for several reasons.[14] To date, the various areas of science that comprise the contemporary neurosciences have yielded a plurality of piecemeal explanations of cognitive phenomena that do not fit together in any interesting way – in part because there is no standardization of the use of concepts designating constructs in contemporary neuroscience[15] and investigators working at all levels of analysis are free to develop cognitive tasks working with whatever assumptions about their experimental subjects they regard as germane to their research.[16] Yet, what investigators seem to want are coherent multi-level mechanistic explanations of cognitive phenomena (Bechtel 2008; Craver 2007). Although philosophers of neuroscience have suggested that such explanations are on the horizon and that "explanatory unification will be achieved through the integration of findings from different areas of neuroscience and psychology into description[s] of multilevel mechanisms" (Piccinini and Craver 2011, 285), they never specify how one area of science that strives to ensure that its investigative strategies individuate discrete cognitive functions will be able to readily integrate its findings with an area of science that does not.

However, we can point to examples in the neuroscientific literature that support the claim that maintaining the current status quo in neuroscience is insufficient and that "paradigm-shifting progress" is necessary if we want to solve important problems like eradicating mental illness.

One representative example is an interdisciplinary research initiative that has evolved during the past 6 years with the aim of developing effective "pro-cognitive" agents to eliminate cognitive deficits in schizophrenia. The teams of investigators involved in the *Cognitive Neuroscience Treatment Research to Improve Cognition in Schizophrenia (CNTRICS)* Initiative regard schizophrenia as involving a core set of cognitive dysfunctions. A primary aim of the initiative is to develop experimental paradigms/cognitive tasks capable of individuating those cognitive functions that are disrupted in schizophrenia. Given that the ultimate aim of these research treatments, which first have to be tested in animal models, the initiative has required cognitive neuroscientists to interact with cognitive neurobiologists, experts in animal behavior, clinical pharmacologists, and members of industry. Their first aim was to develop cognitive tasks that could be used to identify which cognitive functions are disrupted in schizophrenia. However, in developing these paradigms they also had to concern themselves with what kinds of cognitive functions it is possible to study in animal models. This required matching tasks and engaging in task analysis across species. The initiative is still in progress, with different task forces directed at solving different kinds of practical problems. However, what is important for my purposes is that the investigators who are involved in the initiative are interested in bringing about "paradigm-shifting" progress in neuroscience, they believe that this requires interdisciplinary dialogue in the contexts of designing and implementing experimental learning paradigms, and they also believe that on-going interdisciplinary dialogues are fundamental to the success of the initiative (see Sullivan forthcoming).

On a final note, some of the aspects of mind that the DoM authors identify, including "the notion of self" and "higher-order consciousness," are not as readily construed as cognitive capacities, and little work has been undertaken to design investigative strategies for individuating them. This is where insights and perspectives from cognitive psychology, clinical medicine, cultural anthropology, and even philosophy may be helpful. To take one interesting example, Sadhvi Bahtra and colleagues (2013) have developed a semi-structured interview to operationalize the self, so that it may be elicited in terms of behavioral responses, much like cognitive capacities are. One reason for developing this questionnaire is to improve the treatment of patients who suffer from memory disorders like Alzheimer's disease (AD). Even despite the fact that "Alzheimer's disease is often characterized as leading to 'a loss of self' [...] people with dementia continue to refer to themselves with 'I' and often do not recognize the cognitive deficits ascribed to them"

(Bahtra, Geldmacher, and Sullivan 2013). Developing an investigative tool that allows for the qualitative assessment of the self may provide clinicians with clues on how to improve the conditions of life of persons who have AD. Of course, the benefits of developing such investigative tools may not stop there. If successful, such self-assessment questionnaires may one day evolve into investigative tools that may be effectively combined with functional imaging technologies, with the ultimate aim being to determine what brain structures subserve the self. However, the self is a construct, and developing a tool/procedure/task that effectively individuates it will be a challenging iterative process. Fundamental to the success of this process will be a plurality of different perspectives involved in designing and implementing investigative strategies for studying it as well as on-going interdisciplinary discussions about how such strategies succeed and/or fail.

3.5 Conclusion

The proposal for a decade of the mind in 2007 generated little attention among practicing neuroscientists and practitioners working in areas of science targeted as fundamental players in this new decade. Implicit in the questions I have sought to answer in this chapter is a criticism of the proposal for failing to answer a set of questions that are relevant for generating interest in such a decade and making a case as to why such a decade is essential to the future success of neuroscience with respect to understanding how the brain gives rise to the mind.

The first problem with the proposal is that the proponents of the initiative did not explain what they meant by mind. I have here sought to identify what notion of mind they regard as relevant. My claim is that they want to revive the cognitive notion of mind in contemporary neuroscience and with it the intentional stance.

The proposal also lacked a clear motivation. I have provided one such motivation in demonstrating that the cognitive notion of mind and the intentional stance offer two valuable perspectives that should be operative when investigators are designing and implementing tasks that are intended to successfully delineate cognitive functions. Furthermore, when they fail to be operative in the experimental context, the explanatory aims of neuroscience cannot be realized.

A third problem with the proposal is that the authors never specified what this new decade was supposed to look like. They failed to provide a set of guidelines for how to get such an interdisciplinary initiative off the ground. I have suggested that one way in which we can increase the probability of investigators from multiple disciplines coming together to solve common explanatory problems of interest is to sanction perspectival pluralism in experimental contexts across the neurosciences. I have indicated in broad strokes what such perspectival pluralism might look like. The specific details will have to wait for another occasion.

Notes

*The author would like to thank Charles Wolfe for his helpful and insightful comments on an earlier version of this paper.

1. This chapter is based in part on a poster that I presented with Edda Thiels (Department of Neurobiology, University of Pittsburgh) at the 2011 Annual Society for Neuroscience meeting in Washington, DC (Sullivan and Thiels 2011).
2. The members of the group have held several additional meetings, but these meetings have not, as far as a literature search revealed, yielded published proceedings.
3. I think at best what can be said here is that *many* philosophers and social scientists take the intentional stance towards human organisms. However, the extent to which this stance is informed by cognitive scientific thinking about the nature of mind varies across practitioners.
4. The precise aim of the *Decade of the Brain* was, according to G.W. Bush, "to enhance public awareness of the benefits to be derived from brain research" (Presidential Proclamation 3168, 1990). In recent years, as governments and granting agencies have cut funding to scientific research, being able to communicate effectively with members of the general public has become key to generating funds for scientific research.
5. The lack of interest in folk psychology as a viable conceptual-explanatory framework may be regarded as an impediment to the DoM. At a bare minimum, trying to understand the relationship between the *intentional stance*, which common folk and some scientists take towards human organisms, and the *physical stance* – the assumption that an organism's behavior has internal physical (e.g., neural, biochemical) causes – seems prerequisite for effective interdisciplinary communication. No reasons exist to think practitioners in areas of science outside of neuroscience will completely abandon their appeals to folk psychological explanations of behavior, nor is the eliminative materialism for which Paul Churchland (1981) advocates obviously in the offing. Furthermore, given that misunderstandings between neuroscientists and ordinary folk who are looking towards neuroscience for answers may also arise, it seems legitimate for the sake of clarity for neuroscientists to be clear about how they understand the mind and how and in what ways that differs from how non-scientists think about it.
6. As I explain later in the chapter, I think neuroscientists are interested primarily in explaining the capacities of organisms. Recent work in philosophy of neuroscience supports this conclusion (e.g., Bechtel 2008, Bickle 2006, Piccinini and Craver 2011). So, while neuroscientists may be concerned with determining where in the brain *believing*, *wanting*, and *intending* occur, they

are not interested in localizing beliefs, desires, and feelings in the physical world.

7. Clearly, if dispositional analysis is a prerequisite for componential analysis when it comes to cognitive capacities, then proceeding to the explanatory stage that involves componential analysis does require the intentional stance. That this is true is clear in Fodor 1968.
8. To use Fodor's example, "a psychologist might seek to explain failures of memory by reference to the decay of a hypothetical memory 'trace'" (1968, 108).
9. However, as I aim to show in Section 3, the story is a bit more complicated. For psychological explanation will only proceed *successfully* to the second phase in those cases in which an investigator has arrived at the correct functional decomposition, which requires that the intentional stance be adopted in the first phase. Sometimes, however, explanation will proceed to the next stage before one has arrived at the correct functional decomposition, and in the process of investigating the mechanisms of the purported function, investigators will realize they have to "reconstitute the phenomenon" (Bechtel 2008; Craver 2007; Sullivan 2010b). Implicit in my claims in this paper is that reconstituting the phenomenon requires an investigator to consider the intentional stance in addition to considering the physical constitution of the object of inquiry.
10. Piccinini and Craver 2011 argue that a task analysis is nothing more than an incomplete mechanistic explanation of a cognitive capacity. However, to introduce this relationship has the consequence of not keeping investigative strategies, of which task analysis is one, separate from explanatory strategies or models of explanation like mechanistic explanation. Task analysis serves an important function in the context of experimental design – a function it is not obvious that mechanistic explanation can replace.
11. The Stroop task consists of three different types of stimulus conditions that vary across trials. In the *congruent condition*, subjects are presented visually with a word and the color of the text of the word matches the color word (e.g., "red" is presented in red-faced type). In the *incongruent condition* the color of the text differs from the color word (e.g., "red" is presented in green-faced type), and in the *neutral condition* a color-neutral word is presented in either red or green type. The subject's reaction time from the point of presentation of the stimulus on a given trial to the point of responding with the correct word for the color seen (but not read) is measured, and errors in identifying the correct color word are recorded.
12. Watson took this to include introspectionist psychology's "illegitimate children," namely, "functionalist psychology" and "gestalt psychology" (Watson [1924]1970, p. 1).
13. I do not mean here to deny other important critiques of contemporary neuroscience (e.g. Bennett and Hacker 2003), only to point to the irony that the criticisms of behaviorism raised at the Hixon symposium are still applicable to modern cognitive neurobiology (see also Machamer 2009).
14. I want to thank Muhammad Ali Khalidi who, in response to a talk I gave at York University in January 2012, encouraged me to read Wimsatt 2007 as a means to get clear on the kind of pluralism I was advocating.

15. As evidence for this lack of standardization, neurobiologist Yadin Dudai published a dictionary in 2002 in an effort to standardize concepts across those areas of neuroscience that study learning and memory.
16. We see this problem perhaps most clearly in the neuroscientific study of mental disorders in which there has been to date a lack of coordination, particularly between cognitive neurobiological and cognitive neuroscientific approaches (Sullivan forthcoming).

References

Albus, J.S., G.A. Bekey, J.H. Holland, N.G. Kanwisher, J.L. Krichmar, M. Mishkin, S.M. Dharmendra, M.E. Raichle, G.M. Shepard, and G. Tononi. (2007). A Proposal for a Decade of the Mind Initiative. *Science*, 317, 1321.

Barch, D., C. Carter, and J. Cohen. (2004). Factors Influencing Stroop Performance in Schizophrenia. *Neuropsychology*, 18(3), 477–484.

Batra S., J. Sullivan, and D.S. Geldmacher. (2013). Qualitative Assessment of Self-Identity in Advanced Dementia. Poster presented at *Dementia Care at the 13th Alzheimer's Association International Conference: Translating Research to Practice,* Boston, MA.

Bechtel, W. (2008). *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. New York: Lawrence Erlbaum.

Bennett, M.R. and P.M.S. Hacker. (2003). *Philosophical Foundations of Neuroscience*. Malden, Mass.: Blackwell.

Bickle, J. (2003). *Philosophy and Neuroscience: a Ruthlessly Reductive Account*. Dordrecht: Kluwer.

Bickle, J. (2006). Reducing Mind to Molecular Pathways: Explicating the Reductionism Implicit in Current Cellular and Molecular Neuroscience. *Synthese*, 151, 411–434.

Bliss T. and Lømo, T. (1973). Long-lasting Potentiation of Synaptic Transmission in the Dentate Area of the Anaesthetized Rabbit Following Stimulation of the Perforant Path. *Journal of Physiology*, 232 (2), 331–56.

Bush, G. (1990). *Presidential Proclamation 6158*, Library of Congress website. (http://www.loc.gov/loc/brain/proclaim.html).

Carter, C., J. Kerns, and J. Cohen. (2009). Cognitive Neuroscience: Bridging Thinking and Feeling to the Brain and Its Implications for Psychiatry. In *Neurobiology of Mental Illness* (pp. 168–178), 3rd edn, ed. D. Charney and E. Nestler. Oxford: Oxford University Press.

Churchland, P.M. (1981). Eliminative Materialism and the Propositional Attitudes. *The Journal of Philosophy*, 78(2), 67–90.

Cohen, J., D. Barch, C. Carter, and D. Servan-Schreiber. (1999). Context-Processing Deficits in Schizophrenia: Converging Evidence from Three Theoretically Motivated Cognitive Tasks. *Journal of Abnormal Psychology*, 108 (1), 120–133.

Craver, C. (2003). The Making of a Memory Mechanism. *Journal of the History of Biology*, 36(1), 153–195.

Craver, C. (2007). *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Oxford University Press.

Craver, C. and L. Darden. (2001). Discovering Mechanisms in Neurobiology: The case of Spatial Memory, in P.K. Machamer, R. Grush, and P. McLaughlin (eds), *Theory and Method in the Neurosciences. Pittsburgh: University of Pittsburgh Press.* Cronbach, L. and P. Meehl. (1955). Construct Validity in Psychological Tests. *Psychological Bulletin*, 52, 281–303.

Cronbach, L. and P. Meehl. (1955). Construct Validity in Psychological Tests. *Psychological Bulletin* 52: 281–302.

Cummins, R. (1983). *The Nature of Psychological Explanation*. Cambridge, Mass.: MIT Press.

DeCharms, C.R., and A. Zador. (2000). Neural Representation and the Cortical Code. *Annual Review of Neuroscience*, 23, 613–647.

Dennett, D. (1987). *The Intentional Stance*. Cambridge, Mass.: MIT Press.

D'Hooge, R., and P. De Deyn. (2001). Applications of the Morris Water Maze in the Study of Learning and Memory. *Brain Research Reviews*, 36, 60–90.

Dudai, Y. (2002). *Memory from a to z: Keywords, Concepts and beyond*. Oxford: Oxford University Press.

Fodor, J. (1968). *Psychological Explanation*. New York: Random House.

Fodor, J. (1975). *The Language of Thought*. New York: Cromwell.

Gardner, H. (1985). *The Mind's New Science: A History of the Cognitive Revolution*. New York: Basic Books, Inc.

Gibson, J.J. (1960). The Concept of a Stimulus in Psychology. *The American Psychologist*: 694–703.

Giere, R. (2010). *Scientific Perspectivism*. Chicago: University of Chicago Press.

Hebb, D.O. (1949). *The Organization of Behavior: A Neuropsychological Theory*. New York: Wiley.

Hinde, R.A. (1973). Constraints on Learning: An Introduction to the Problems. In R.A. Hinde and J. Stevenson Hinde, eds *Constraints on Learning* (pp. 1–19). London: Academic Press.

Kandel, E.R., and W.A. Spencer. (1968). Cellular Neurophysiological Approaches in the Study of Learning. *Journal of Physiology*, 48, 65–134.

Kohler, W. (1947). *Gestalt Psychology: An Introduction to New Concepts in Modern Psychology*. New York: The New American Library.

Kuhn, T. (1962). The Structure of Scientific Revolutions. Chicago: University of Chicago Press.

Lømo, T. (2003). The Discovery of Long-Term Potentiation. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 358, 707–714.

Lorenz, K. (1965). *Evolution and Modification of Behavior*. Chicago: University of Chicago Press.

Machamer, P. (2009). Neuroscience, Learning and the Return to Behaviorism. In *The Oxford Handbook of Philosophy and Neuroscience*, ed. J. Bickle. Oxford: Oxford University Press.

Perlstein, W., C. Carter, D. Barch, and J. Baird. (1998). The Stroop Task and Attention Deficits in Schizophrenia: A Critical Evaluation of Card and Single-Trial Stroop Methodologies. *Neuropsychology*, 12(3), 414–425.

Piccinini, G. and Craver, C.F. (2011). Integrating Psychology and Neuroscience: Functional Analyses as Mechanism Sketches. *Synthese*, 183, 283–311.

Spitzer, M. (2008). Decade of the Mind. *Philosophy, Ethics, and Humanities in Medicine*, 3, 7–11.

Sullivan, J.A. (2009). The Multiplicity of Experimental Protocols: A Challenge to Reductionist and Non-Reductionist Models of the Unity of Neuroscience. *Synthese*, 167, 511–539.

Sullivan, J.A. (2010a). a Role for Representation in Cognitive Neurobiology. *Philosophy of Science*, 77(5), 875–887.
Sullivan, J.A. (2010b). Reconsidering Spatial Memory and the Morris Water Maze. *Synthese*, 177(2), 261–283.
Sullivan, J.A. (forthcoming). Stabilizing Mental Disorders: Prospects and Problems. In *Classifying psychopathology: Mental kinds and natural kinds*, eds H. Kincaid and J.A. Sullivan. Cambridge, Mass.: MIT Press.
Sullivan, J.A. and E. Thiels. (2011). The Place of the Mind in Contemporary Neuroscience. (Poster presented at the *Society for Neuroscience Annual Meeting* Washington, DC).
Sweatt, J.D. (2009). *Mechanisms of Memory* (2nd edn). London: Academic Press.
Taylor, C. (1964). *The Explanation of Behavior*. London: Routledge and Kegan.
Von Eckardt, B. (1993). *What Is Cognitive Science?* Cambridge, Mass.: MIT Press.
Watson, J.B. (1924 [1970]). *Behaviorism*. New York: W.W. Norton & Company.
Wimsatt, W. (2007). *Reengineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge, Mass.: Harvard University Press.