



Enactivism and the New Teleology: Reconciling the Warring Camps

Ralph D. Ellis
Clark Atlanta University
rellis@cau.edu

Received December 2013; accepted September 2014; published Autumn 2014.

Abstract

Enactivism has the potential to provide a sense of teleology in purpose-directed action, but without violating the principles of efficient causation. Action can be distinguished from mere reaction by virtue of the fact that some systems are self-organizing. Self-organization in the brain is reflected in neural plasticity, and also in the primacy of motivational processes that initiate the release of neurotransmitters necessary for mental and conscious functions, and which guide selective attention processes. But in order to flesh out the enactivist approach in a way that is plausible and not merely an epiphenomenon, it is necessary to confront the problem of causal closure in a serious way. Atoms and molecules in the brain do not violate the normal causal principles that govern them in other contexts. The theory of self-organizing dynamical systems must be developed in a way that is compatible with causal closure rather than contradicting it.

Keywords: Enactive; self-organization; dynamical systems; neural plasticity; causal closure.

Introduction

Clinical psychologists tend to view human behavior in teleological terms, while experimentalists (including many of those same clinicians when adopting a more theoretical stance) view all causation as efficient, with human actions appearing not so much as self-initiated actions, but as reactions to stimuli. The enactive approach to cognition can go a long way toward reconciling these different perspectives. When enactivism is grounded in a sober analysis of causal relations and neural mechanisms (as in Kauffman 2003; Monod 1971), it can make room for organismic purposes that drive both action and the nature of re-actions. Multiple realizability and neural plasticity are major components of this analysis. In such processes, the same organizational activity can make use of different sets of micro-level components, which the organ-

ism actively appropriates and replaces as needed, insofar as possible given environmental conditions and the organism's particular needs.

However, the structure of causal analyses within enactivist and self-organizational theory must be taken more seriously than often is done. Vague talk of "higher level" processes that "constrain" causal relations at a lower level can be misleading or even meaningless. The problem of causal closure (Kim 1992, 1998) must be confronted rather than merely talked around. When the causal analyses are done in a careful way, and integrated with the motivational systems of the brain that orchestrate self-initiated and self-energized actions, explanations can remain consistent with normal efficient causal processes that are by no means violated by the atoms and molecules that make up the nervous system. Yet at the same time, enactivist cognitive theory combined with the neurological phenomena of neural plasticity and the serious working out of self-organizational causal theory can make room for a meaningful distinction between action and a merely complicated sequence of re-actions.

This paper will attempt to provide an overview of the path that can lead from self-organizational causal theory, through mechanisms of neural plasticity in the nervous system that serve organismic purposes, to the primacy of motivated selective anticipation in perception, to the enactivist manifestations of these mechanisms in cognitive processes. I will refer here to examples of the timing and interactions of brain mechanisms in perception, imagination, and motivated attention direction (discussed more fully in Ellis 1995, 2005; Ellis and Newton 2010) that illustrate the primacy of purposeful processes of the anticipatory (Freeman 2001) and searching (Panksepp 1998, 2012) functions of the motivational brain. These self-energizing and anticipatory systems lead to views of cognitive functions that fit nicely with enactivist approaches such as those of Noë (2006) and Thompson (2007).

But equally important, I will also discuss the philosophical analyses that are needed to reconcile self-organizational and enactive processes (the "autopoiesis" emphasized by Thompson) with the problem of causal closure of physical systems. More is needed for this reconciliation than merely multiple realization; we need a concept of a *particular kind* of multiply realizable system in which the overall pattern of organization is of such a nature that it shows a tendency to appropriate, replace, and reproduce the components of its own definitive patterns of activity by virtue of the organized way it exchanges energy and materials with the environment (Newton 2000). Such a theory has been extensively developed by Stuart Kauffman (1993) and Scott Kelso (1995).

Science and subjectivity

“Psychology as a science” is a strange expression. It means eliminating from the study of the psyche—the subject of experience—everything that does not present itself in the form of an object of experience. Above all, this means eliminating teleology, because modern science explains things from the standpoint of mechanistic causal relations. Rolls (1999), for example, explains emotion in these terms, and the account of perceptual processing is standardly framed in terms of sequences of efficient causes in the brain (for example, see any standard textbook such as H.R. Schiffman 2001). John Bickle (2008) expresses the philosophical rationale for rejecting teleology in favor of reductionist causal accounts. When I decide to raise my hand, and then do so, we must explain the causes of motion for a measurable and observable physical mass. Modern science therefore must explain the raising of the hand in terms of objectively observable physical causes. And this means in terms of the electro-chemical reactions in the nervous system, which can be reduced to a series of micro-level causal sequences. These causal sequences occur at a level at which things obey the laws of classical Newtonian mechanics, and thus are causally closed. If physical antecedents are necessary and sufficient as an explanation of them, then the conscious correlates of those physical antecedents can play no causal role, unless we completely reduce consciousness to its physical components. This equation of the conscious with its physical components is now a standard solution to this problem, endorsed for example by Bickle as the only reasonable way to make physical and mental causation compatible with each other. If consciousness is reducible to a physical sequence of micro-level events, then it can have the same causal powers as that sequence of events.

And even if we do opt for a strict identity between consciousness and its physical correlates, the past century of work on psychophysical monism still seems to leave us with only mechanistic causes—but with some of them correlating with mental events—and thus without any real teleological causal relations, but only a subjective impression of one. Instead of raising my hand *for the purpose* of reaching the coffee cup, strictly speaking, I am raising it because of neurophysiological events that correlate with the *image* of the coffee cup, which then physically triggers a hand movement, all of which in principle is explainable by physical mechanisms at the level of molecules and chemical reactions.

If the physical components correlating with the thought “I’d like to raise my hand” completely explains the physical raising of the hand, then the idea that the conscious thought caused the hand to raise has only the *appearance* of a teleological process. The relation between the mere appearance of a teleological process and the underlying reality of a mechanistic process is the same as when birds *appear* to migrate “for the purpose” of finding a better climate

to spend the winter. In reality (in the standard mechanistic scientific account) there ultimately must be micro-level causal explanations for what causes them to follow the patterns of flight that they follow. The real explanation for why they migrate in those patterns has no more to do with the birds' achieving a "purpose" than the causal explanation of a giraffe's long neck is that it stretches "in order to" reach the leaves in high branches. To be sure, there are natural selection processes at work, but these in turn are ultimately reducible to a series of mechanistic cause-and-effect relations. Natural selection, far from embracing teleology, actually explains it away.

In no sense is the giraffe's long neck *mechanically caused* by the giraffe's "having a purpose," nor is the birds' pattern of migration mechanically caused by the purpose of the migration. So if all causation is mechanistic, the "purpose" achieved by the birds and the giraffe play no causal role, because this having-a-purpose does not play the causal role that the underlying mechanistic antecedents play.

In theory, it has seemed since the inception of "psychology as a science" that teleological explanations of human behavior and consciousness must work the same as with the giraffe and the birds. It subjectively *appears* that our having a purpose is what causes us to raise our hand, but the real causal work is done by physical antecedents that preceded any subjective impression of "having a purpose," and which are entirely sufficient to explain the resulting behavior without reference to any feeling of teleology. In reality, in the dominant streams of purely empirical psychology, our human behavior is as mechanistic as any other causal relation. While these mechanisms seem purposeful from the subjective perspective, objectively they result from push-pull dynamics. The past is always what causes the future. The same problem presents itself for other apparently purposeful processes in nature, such as the remarkable shunt mechanisms in the metabolism of nutrients in animal systems and many other non-conscious contexts. Our focus here, however, is the need for something resembling a kind of teleology in conscious processes. We will see later that the self-organizational theories that reconcile mental with physical causation are also the same ones that reconcile non-conscious purposeful processes with their micro-level efficient causal mechanisms.

Although many would like to avoid the problems of classical Newtonian mechanics by explaining consciousness at the quantum level (as in Jibu and Yasue 1994; Globus 1992, 2003), in my view this move still cannot avoid the problem of mental causation and other related contradictions between the subjective experience of teleology and the mechanistic explanations of its neurophysiological correlates. Even if quantum behavior does not fit the causally closed picture of Newtonian physics, the fact remains that conscious processes have been correlated thousands and thousands of times with causal mechanisms that *do* occur at a level where causal closure still obtains. We

cannot avoid these correlations by descending to the quantum level of explanation. That move does not erase the correlations.

Classical Newtonian causal relations are causally closed (as emphasized, for example, by Kim 1992, 1993, 1998). If some combination of conditions are sufficient to produce an outcome, then no other condition can be necessary to produce that same outcome. Conversely, if some combination of conditions are necessary to produce an outcome, then no other condition can be sufficient to produce that same outcome. The result is that teleological explanations of why and how I raise my hand can make no reference to my subjective decision to raise it, because the raising of the hand has been entirely accounted for by micro-level physical conditions—unless, of course, I completely reduce the conscious event to a physical one, in which case we still are left with no teleological causal relations, only mechanistic ones.

Adding further to the reductionism of this schema is the passive receiving model of perception, which leads to stimulus-response explanations of all psychological phenomena, including affective and conscious ones. If I am quietly sitting, without any light impinging on my retina, I do not see any object. When the light hits the retina, I see an object. Many psychologists are content to infer that the difference between seeing and not seeing must be caused by the introduction of the light hitting my retina—i.e., the stimulus. My consciousness of the stimulus is then the subsequent product of this causal dynamic. Consciousness in effect is *caused by* physical actions done *to* the body by external objects—not by actions initiated by the organism which would then use or adapt to the external objects. So consciousness must be an epiphenomenon, a mere side effect of the activity of stimuli as they hit my nervous system and in turn lead to cause and effect mechanisms in it. From this standpoint, consciousness is always construed as the last event in a series of causes and effects.

There is of course one easy way out of the standard mental causation problem. If consciousness is identical to a physical process, then whatever is caused by that physical process is also caused by the corresponding conscious act (as in Smart 1970). In that case, even a conscious choice could still become an intermediate cause in a chain, since it would be identical with some of the intermediate physical events in the chain. To be sure, consciousness would only be an intermediate step in a chain of mechanistic causes in that case, but at least it would play a causal role.

Up to this point, the dominant schools of empirical psychology have taken this intermediate causal role of consciousness to mean that we have a subjective reaction to a stimulus input, and a conditioned (mechanistic) response, of which we can form a subsequent subjective impression—as in behaviorism. Or this view can be qualified by allowing innate configurations of brain matter, hormones, and neurotransmitters to partially determine *how* we respond

to a given stimulus—as in the traditional drive theories such as those originally developed by Hull (1952) and Spence (1956; see also Spence and Frith 1999), and currently handed down to us by Rolls (1999), LeDoux (1996) and others.

But of course such a literal psycho-physical identity theory presents a plethora of other problems. For example, there is the “hard problem” pointed out by Chalmers (1995), which is essentially a special case of Levin’s (1983) “explanatory gap” problem. In Chalmers’ formulation, the problem is that, if we can explain everything about a physical event, and if the physical event is identical with a conscious event, then this would mean that we have completely explained the conscious event. But the problem is that something would still be left unexplained: Namely, after we have explained what caused physical event XYZ to occur, we still would not have explained why an event like XYZ should have the property of consciousness, whereas other physical events subject to the same kinds of explanations do not have the property of consciousness. To explain what caused XYZ to occur does not explain why XYZ has the property of consciousness. Those are two different questions. The second one is what Chalmers means by the “hard problem” of consciousness. He explicitly distinguishes this from the “easy problems” of consciousness. Explaining that consciousness always correlates with XYZ is easy. And explaining what caused XYZ to occur is also easy, compared with the “hard problem.” Yet one often hears presentations of solutions to these easy problems accompanied by the claim that they are solving, or will lead to the solution of the “hard problem.” This way of speaking betrays a misunderstanding of what Chalmers means by the “hard problem.”

I use Chalmers’ “hard problem” as only one of many problems that are created by a literal psycho-physical identity theory. As far as teleology is concerned, psycho-physical identity remains at odds with it even if we completely accept psycho-physical identity. Consciousness is still caused by an interaction of mechanistic physical events. And here again, the standard application of empirical psychology to this model has been in effect to use stimulus response explanations, combined with some assumptions about innate physiological predispositions. The result is that “action” is explained as a series of micro-level *re*-actions. We are left with no distinction between action and reaction.

Clinical psychologists have long wished that there *could* be such a thing as teleology in human processes, because our subjective reporting of what we feel to be the purposes of our actions and motivations is largely experienced as teleological. I do not experience myself as deciding to listen to Tchaikovsky’s Sixth Symphony because some classical or operant conditioning history caused something to trigger me to respond to some stimulus in my environment in some way. My experience is that I first *felt* a certain way, and because I felt that way, I wanted to listen to some music appropriate to that feeling. Rather than being caused by the music, the feeling pre-exists and then uses

the music. I engage in voluntary actions needed for access to the right kind of music—driving to a concert hall or playing a disk on a stereo.

The enactive approaches to psychology offer the kind of teleology that is needed here. As the name implies, enactivism begins with the idea that there is such a thing as action—not merely a sequence of *re*-actions. As is now well known, new scientific work on the theory of self-organization in certain kinds of physical systems promises to allow the whole organism to act from its own self-initiated motivations, and not merely to react or to display a complicated system of reactions (see Kauffman 1993; Monod 1971; Newton 1996, 2000; Ellis and Newton 2010; Ellis 2005).

The distinction between action and reaction, construed scientifically, can be grounded in the idea of complex dynamical systems, which I take as synonymous with self-organizing systems. As Newton (2000) defines it, a dynamical system is a thermodynamic system that maintains the continuity of its functional patterns while exchanging energy and materials with the environment. Such a system acts as a whole to seek out, select, appropriate, rearrange, and reproduce the micro-constituents needed to keep its ongoing pattern going with continuity. As Thompson (2007) suggests, the pattern uses its parts to maintain the continuation of the pattern and, within the limits of its situation, replaces the parts as needed for this purpose.

We must be very careful to avoid facile explanations of the relationship between the mechanistic level and the self-organizational level of causal analysis. Locutions like “the higher level of organization *constrains* what can happen at the lower level” (as in Thelen and Smith 1994; Globus 2003) may be satisfying for some purposes, but they do not answer the crucial questions about causation. In classical mechanical causation, which the neurons and neurotransmitters of the brain are observed to obey quite consistently, the cause (or combination of causes) is necessary and sufficient for its effects under the given background conditions. This is as true for turning on an electric light as for shunt mechanisms in biological organisms. So if saying that the higher level “constrains” what the lower level can do simply means that lower level causal relations are sufficient only after we have assumed that certain background conditions are in place (for example, a good bulb in the light socket and a complete electrical circuit), then we still have not distinguished between linear and complex dynamical causal systems. We need to be careful to give a coherent explanation as to how a higher-level of organization “constraining” the lower level is any different from any other causal situation. How are self-organizing systems different from non-self-organizing ones?

In my view, the difference hinges on the fact that, in self-organizing systems, the system as whole displays a tendency for its various components to be rearranged when needed to provide the background conditions that are assumed by a needed micro-level causal relation (Kauffman 1993; Monod 1971).

We do not need for this overall plasticity to be devoid of previous mechanistic explanations for how it initially achieved the plastic condition—for example, the micro-level reactions that occurred during embryonic development (as both Kauffman and Monod point out). The point is that very early in the etiology of a self-organizing system, it acquires a tendency to rearrange its own components in such a way that the structural integrity of the whole is maintained by constantly rearranging the background conditions needed for the given micro-level linear reaction to occur (see Ellis and Newton 2010). This is the sense in which the overall causal situation is different from simply a combination of linear reactions. The whole system seeks out, appropriates, rearranges and replaces its own parts as needed to maintain the pattern of the whole or to change the pattern in ways called for by the needs of the whole system. We see this in many examples of neural plasticity, such as in stroke recovery (Restak 1984: 256; Springer and Deutch 1981: 173-212), in recovery from memory loss from head injuries (Restak 1984: 360ff; Wrightson 1989), and in drug addiction (in which neurotransmitter receptors change their size and number to avoid imbalances resulting from the substances that are over-supplied by the drug user—for example, see Valenstein 2002, Chapter 5).

How can such a dynamical systems approach avoid the problems of previous mechanistic explanations? Is the dynamical system merely an epiphenomenon of the entire series of mechanistic interactions of its components, and thus in reality only a series of reactions after all? Does the system as a whole actually violate causal closure in ways that would contradict a huge mass of empirical evidence suggesting that there are systematic correlations between physical cause-effect sequences and conscious events? Or is the action of the dynamical system somehow “compatible with” the mechanistic explanation of the behavior of its micro-level constituents? Those are pressing problems that enactivism must face if it is to ground itself in action in a meaningful sense, rather than only an epiphenomenal appearance of action.

The new approach has arisen partly from a rejection of some of the assumptions of the stimulus-response and passive-reaction models of the past century. These assumptions, as we have seen, led to the relegating of consciousness to an epiphenomenal role in philosophy and science. They ultimately stem from the assumption that the reality which ultimately must explain all causation is thoroughly “bottom-up,” with the real causal work done at an atomistic-reductionist level; that representational conscious activities (thoughts and perceptions) are clearly distinguishable from non-representational ones (feelings and emotions); and, perhaps most important, that all reality is fundamentally reactive and passive rather than active—i.e., that nothing does anything unless caused to do it by some external force acting on it, that there is no such thing as a pattern of activity which organizes its own substrata rather than the other way around. In short, for “modernist” metaphysics (as postmodernists like Globus, 1992, 2003 call it), there was no important or non-arbitrary dis-

inction between non-living things and living ones (i.e., those which appropriate, rearrange, and reproduce the needed substrata in order to maintain a higher-order pattern of activity); yet the difference between conscious beings and non-conscious ones (e.g., computers) would need to hinge crucially on this distinction.

What is needed for a teleological yet scientific analysis of intentional actions

In the modernist framework as applied to psychology and cognitive neuroscience, consciousness is a causal result of a stimulus input which then gets processed in the brain. Perceptual consciousness is thus assumed to result from stimulation of the occipital lobe and V4 visual areas, as a result in turn of a perceptual object's stimulating the nervous system. Thomas Natsoulas (1994) calls this approach an "appendage" theory of consciousness—the idea that consciousness is a byproduct of a cause and effect mechanism but does not have the power to act as a cause itself. But there are problems with trying to interpret perceptual consciousness in this way.

For example, it has often been observed that occipital activation in response to incoming visual data often results in no perceptual consciousness. Additionally, it is necessary that the parietal and frontal lobes must also play an active role (Farah 1989; Luria 1980; Posner 1980, 1990; Posner and Petersen 1990); and it is well established that this activation of the parietal and frontal lobes is not simply caused by the earlier activation of the occipital cortex (Aurell 1989). As I discussed more extensively elsewhere (Ellis 1995, 2005), the parietal lobe is almost immediately adjacent to the occipital lobe, yet in the process of generation of a conscious perceptual experience, the parietal lobe is not active until about a third of a second after the occipital lobe is activated when a novel stimulus is presented (Runeson 1974; Srebro 1985; McHugh and Bahill 1985). Given that nerve impulses travel at about 100 miles per hour, why should it take a third of a second to travel only a few millimeters?

If the imaginative activity of the parietal lobe were really caused by the nerve impulse which travels to it from the occipital lobe, the impulse should be delivered virtually instantaneously. Whatever is happening during this third of a second that is also needed in order for consciousness of the object to occur cannot be caused by the passive receiving of the nerve impulse to the parietal lobe from the occipital lobe, which in turn receives it from the incoming stimulus. If so, the spreading of activation would be much quicker than a third of a second; it would be virtually instantaneous. Thus it appears that the response is not caused by the stimulus.

This paradox arises only if we assume that the parietal lobe (which is active when we are conscious of visual images) can only be activated as a result of prior occipital activity, which in turn results from prior optic stimulation originating from the environment. But recent research shows that this is not the case. Instead, what happens is that, prior to occipital processing of the visual stimulus, the parietal lobe has already been activated by the frontal lobe (as shown by Aurell 1989, Damasio et al. 2000, and many others), which in turn is activated by the midbrain, which is the focus of emotional-motivational activity as triggered by thalamic arousal by the stimulus only if the stimulus is generally felt as possibly emotionally important for the organism's purposes (Luria 1980; Posner 1990; Damasio 1994). The needs of the organism as a whole must first motivate the asking of questions about what kinds of environmental stimuli might be important for the organism's purposes; at this point, the frontal lobe becomes active. As these questions are formulated with the help of the frontal lobe, the parietal lobe then begins to entertain vague images and/or concepts of the kinds of emotionally important objects that might be present in the environment. If and when this frontal-limbic-parietal activity, once having been developed, finds itself resonating with patterns of activity in the occipital lobe (which reflects sensory stimulation)—only then does perceptual consciousness occur.

The one-third-second delay does not result from any slowing of the incoming nerve impulse as it “travels” from the occipital lobe to the parietal lobe. The parietal lobe (which is active when we are conscious of visual images) is not activated in response to the occipital lobe's activity at all. Instead, the organism must purposely activate the frontal and parietal lobes to “look for” emotionally important categories of objects which the thalamus has already alerted the organism might be relevant, and this “looking for” activity has already begun the forming of visual or conceptual imagery (including proprioceptive and sensorimotor imagery) prior to any occipital activity's having any effect on our perceptual consciousness (since at this point the impulse has not yet “traveled” from the occipital to the parietal lobe). As Thompson (2007), Noë (2006), and Ellis and Newton (2010) formulate this process, the organism is always already engaged in a series of both conscious and non-conscious actions, and then makes an assessment of the environment based on how it seems to facilitate or thwart the process. Neurophysiologically, the prefrontal cortex's searching function and the hippocampal and amygdala emotional processes are already active long before the occipital perceptual processing occurs (Damasio et al. 1990; Haines et al. 1997).

On the enactivist account of intentionality, the organism must act on its environment in order to be conscious of it; consciousness cannot result from a mere passive reaction to incoming input. On this view, the model of the mind as a passive receiver of causal work done by stimulus inputs and other mechanical computations is backwards in some important respects. The or-

ganism must first purposely act, and only then can consciousness of the environment result. It is this fundamental shift in the direction of causation which is now sometimes referred to as the “enactive” view of the mind—a term coined by Varela et. al. (1991/1993). Rather than a stimulus causing a response, it is the response which must occur first, and then act on the incoming afferent signals to produce a stimulus. We might call this enactive approach the current “Copernican revolution” in cognitive theory and neuroscience.

The sensorimotor account of intentionality

The philosopher Natika Newton has elaborated an entire theory of consciousness around the notion that, in effect, we must subliminally use “sensorimotor action images” to imagine the action affordances of objects in order to be consciously aware of them (this careful analysis is developed through a sequence of detailed studies, including Newton 1982, 1989, 1991, 1992, 1993, 1996, 2000). According to Newton, perceptual consciousness is always preceded by an act of imagination. We first imagine ways in which we could possibly act in relation to the object at hand. The action motivated by the action planning process creates an “anticipation” as to possible environmental feedback, and these anticipations constitute mental images. If the anticipations are fulfilled by correlative perceptual input, the result is a percept; if the anticipations are not fulfilled, then the result often is still a state of consciousness, but the state consists only of a mental image of a non-present object or state of affairs. The anticipation, however, must precede the effect of the incoming data from the senses on our perceptual consciousness. Subjunctive ideas about imagined possibilities for action are thus prior to perceptual input, and action planning guides the process of “looking for” instantiations of the subjunctive category (for example, the image) as actually instantiated in the environment. For example, perceptual studies by Runeson (1974), Srebro (1985) and others show that an object previously anticipated is much more readily perceived.

Newton’s theory is a clearly articulated example of an enactivist account in which consciousness must play an active role in the context of the experience of humans and other higher animals. The role it plays is similar to the “pragmatic” role that Marcelo Dascal (1987) assigns to it in bringing about the kind of information processing that is accompanied by conscious experience. But if the organism must act in this sense, it is not merely reacting to stimuli, and therefore is not an easy fit to the notion of an epiphenomenon or “appendage” to a basically non-conscious computational process. This aspect of the enactivist account builds from the foundation established by J.J. Gibson (1986) in his notion of “affordances.” We understand and identify an object by imagining how it would be possible for our bodies to act in relation to the object. Newton and other recent enactivists use neurological work (for example, Damasio

1994; Posner 1990; Farah 1989; Luria 1980; Cytowic 1993), and research on mental imagery (for example, the “mental models” developed by Johnson-Laird and Byrne 1991) to show how action planning grounds our understanding of objects, and ultimately of language, concepts, and logical relations. An infant identifies objects in terms of whether they “afford sucking,” “afford throwing,” etc.

This imagistic account can ground our understanding of logical relations because subjunctive action affordances allow us to imagine what would happen “if” such and such circumstances were to obtain. For example, to anticipate that “If I throw a ball at something it will knock it over,” is very similar to believing that “If I *were* to throw a ball, it *would* knock something over.” Thus anticipations of the future ground our understanding of subjunctives and therefore allow a grasping of abstract logical concepts. In Newton’s approach, the key to this “foundation of understanding” is the process of action planning. This theory is consistent with a host of neurophysiological findings—for example, that the brain mechanisms underlying abstract thought are very similar to those underlying action-planning in the context of body movement (Ito 1993; Damasio 1994).

The crucial point for our purposes here is that, in the enactivist approach, the modernist biases mentioned above become questionable and a new paradigm suggests itself. Because the organism must anticipate actions toward its environment in order for consciousness to occur, consciousness is not merely passively caused by incoming stimuli or unconscious computations performed on incoming stimuli. The body’s organization of stimuli occurs prior to the reception of the stimuli, and if the body does not actively seek to appropriate and rearrange the physiological substrata for its own desired patterns of conscious activity, this consciousness can never occur.

In a sense, we could think of the relationship between neurons and the pattern of activity that constitutes consciousness as similar to the relationship between the molecules of wood in a door and the sound wave that passes through the door. The wave is composed of a pattern of the movements of the wood particles, but we would not want to “explain” the wave by reference to the door. The door does not cause the wave, but on the contrary, the wave causes the particles in the wooden door to conform to its pattern as it passes through the door.

Correlatively, an enactive approach such as Newton’s or Varela’s requires a rejection of the epiphenomenalist account of consciousness as merely the tip of an iceberg which consists of unconscious computational brain processes. Instead, consciousness directs much of this activity, and much of it would never take place without the direction of consciousness; yet it is important that consciousness itself is embodied—not in computational cerebral processes, but rather in emotional and motivational activities of the whole organism.

It is the emotionally motivated process of action planning that directs the focus of conscious attention, not a computer-like computational process.

For the same reasons, enactivism must reject the modernist assumption that representational states (thoughts and perceptions) are caused by perceptual inputs. In fact, in the enactivist account, representational processes are not even clearly distinguishable from non-representational ones (feelings and emotions). Emotion and motivation direct the action imagery that grounds the experience of the action affordances of objects. The purposes of the whole embodied organism come into play in directing conscious attention. Much of our rational processing, for the reasons that I sketched earlier, results from what Newton calls “proprioceptive imagery.” In proprioceptive imagery, we imagine what it would be like to perform a bodily action, and this in turn grounds the understanding of a subjunctive concept of such a movement and of the environmental objects that afford it. The understanding of “what would happen if...” grounds our learning of which kinds of logical inferences can be relied upon as we go through life. And much of this action imagery can occur on a subliminal or pre-conscious basis. Only in the presence of a strong enough motivating selective attention process and a suitable environmental context (or a sufficiently rich imagining of one), does consciousness accompany the action imagery.

In essence, twentieth century empirical psychology was built on an ontology of the physical world in which everything is fundamentally reactive rather than consisting of patterns of activity which appropriate their atomistic components. When scientists tried to apply this passive ontology to the explanation of consciousness, the result was that no explanation was possible. The easy problems of consciousness could be addressed, but no attempt could even be made at the “hard problem.” The next section will consider the advantages of enactivism in making more sense out of the relationship between consciousness and the natural world than was possible in the passive-receiving model of the mind and the epiphenomenalist conceptualization of the relevant causal relations.

How enactivism can confront the “hard problem” head-on

The approach that I have been outlining is equipped to offer a new perspective on Chalmers' (1995) “hard problem” of consciousness. Computationalists had maintained throughout most of the twentieth century that consciousness can be explained either as an epiphenomenon of, or as identical with, a digital computer-like process which uses the hardware of the brain to process its software. The point of Chalmers' “hard problem” is that we can easily imagine such a computational process as occurring in the absence of consciousness. Therefore, some further explanation is required in order to understand why

consciousness does in fact accompany such computational processes in certain cases (for example, in human organisms).

The enactive approach we have been discussing avoids this problem because it focuses on elements of consciousness as actually phenomenally experienced that were left out of the traditional computationalist model—elements that lend themselves to being correlated with action-oriented as opposed to passive-receiving physiological substrata. It is not so easy to imagine that this teleological or action-initiating combination of physiological substrata, including emotionally motivated selective attention, imagery, and resulting action, could have occurred unaccompanied by its conscious correlates. For example, consider the enactive model of consciousness that Newton and I described in a co-authored article and subsequent further elaboration in book form (Ellis and Newton 1998, 2010). On the formulation developed there, consciousness requires

an interested anticipation of possible sensory and proprioceptive input such that the pattern of the subject's interest determines the modality, patterns, and emotional significance of the anticipated input. Specifically, the anticipation takes the form of a sensorimotor, proprioceptive and affective “image” of a state of affairs “looked for” by the subject.... The content of consciousness is vivid to the extent that the activity constitutive of the interest in the future resonates (in terms of holistic patterns of activity) with the activity of incoming (afferent) imagistic data and with activation of memories of past imagistic and conceptual data. (Ellis and Newton 1998: 432)

Subjects can actually experience this effect in perceptual attention experiments. When the subjects are instructed in such a way that they are anticipating an object before it is presented, they perceive it more readily (Corbetta 1990; Pardo et al 1990; Logan 1980; Hanze and Hesse 1993; Legrenzi et al 1993; Rhodes and Tremewan 1993; Lavy and van den Hout 1994). Posner and Rothbart (1992) report that "During vigilant states the posterior attention network can interact more efficiently with the object recognition system of the ventral occipital lobe (96)." This attentional process "increases the probability of being able to detect the relevant signal (97)." Imagining a given object involves a process that is very similar, both phenomenologically and neurophysiologically, to “looking for” an object of that type, even if we are aware that we will not find it in the actual current environment. When I form a mental image of the pink wall as blue, I am executing many of the same brain processes as when I “look for” or “anticipate” blue in the wall, even though I do not find the blue I am “looking for” in this particular wall. As Merleau-Ponty says, "I give ear, or look, in the expectation of a sensation, and suddenly the sensible takes possession of my ear or my gaze, and I surrender a part of my body, even my whole body, to this particular manner of vibrating and filling space known as blue or red" (1962: 212). Later he sums up in this way: "It is necessary to 'look' in order to see" (1962: 232). And "The warmth which I feel when

I read the word 'warm' is not an actual warmth. It is simply my body which prepares itself for heat and which, so to speak, roughs out its outline" (1962: 236). Helmholtz (1962) makes a similar point which is now widely accepted among neurologists: "We let our eyes traverse all the noteworthy points of the object one after another." I.e., the organism must actively search for information in the environment before that information is consciously seen. Vision is active, not passive.

One of the main differences between conscious and non-conscious information processing (as in nuts and bolts computers) is that consciousness involves an emotionally *interested* anticipation of a possible or actual stimulus input (Cytowic 1993). In order to be conscious of a specific input, we must actively and motivatedly "pay attention," as inattentive blindness experiments clearly show (Mack and Rock 1998; Ellis 2005). The action of directing attention, like any other action, must be motivated by the needs of the organism. To be sure, afferent processing involving the occipital lobe is *part* of what produces conscious awareness of an object, but this awareness also requires corticothalamic loops instigated by subcortical motivational activity and the cooperation of frontal-limbic loops with the anterior cingulate (see Bachmann 2000; Lethin 2002, 2005; Posner and Rothbart 1992; Damasio 1994; Farah 1989; Aurell 1989; Luria 1980.) In the enactive approach, the primary organismic need that motivates consciousness of objects is the need to anticipate future data which are considered important for the organism's purposes (Dennett 1996).

The emotionally motivated anticipation of input leads to "imagery" in an enactivist rather than modernist sense. In enactivist terms, an "image" is not a physical replica or "encoding" of the object in the brain, but rather an enacting of the process that one would undergo if perceiving the object—in other words, on my formulation, a sense that we are looking for (or listening for, tasting for, proprioceptively feeling for, etc.) an object or state of affairs in a widely distributed neurophysiological pattern similar to what *would* be executed if we were to actually see, hear, or taste it. This idea is highly consistent with Walter Freeman's (1991) findings, for example that cats activate their olfactory patterns merely in anticipation of the presentation of the smell of fish, before the smell is actually presented. It is also consistent with the now familiar finding that imagining a given action activates most of the same brain processes as actually executing the action. Jeannerod (1994, 1997), for example, explains that when we imagine a movement, we execute all the same brain processes as in actually doing the movement, but at the same time we inhibit the action command at the point when it would feed forward to the body's extremities in order for the action to be carried out.

When we say that consciousness is distinguished from non-conscious processing by being emotionally motivated, this should not be taken as implying that emotions and motivations are necessarily conscious. There are non-conscious yet purposive phenomena in nature, especially in biological organisms. For example, the human organism purposefully does what is necessary to regulate its heartbeat and blood pressure, yet normally is not conscious of doing so. Merleau-Ponty (1941) defines a “purposeful organism” as one which changes, replaces, or readjusts the functioning of its own parts according to what is needed to maintain or enhance the existence and functioning of the whole organism. The purpose-directing role of emotion is not a sufficient condition for consciousness, but in the enactivist account as I have formulated it, emotion and motivation are necessary conditions.

It is true that we can have interested anticipations of the future in this sense without any experienced consciousness. This can occur because, throughout nature, there is purposeful activity without conscious awareness, because self-organization in complex systems can occur in various contexts throughout nature. It is also true that we can process afferent data without consciousness, as in “blindsight” phenomena. And of course there is also holistic processing without consciousness, as in typical holograms. There are also sometimes non-conscious interests alongside of non-conscious data processing, but without consciousness. We can even have non-conscious anticipations of the future, as in operant conditioning, juxtaposed with non-conscious activations of stored information or of present afferent activity, with no consciousness of the process.

So enactivists should be careful not to try to explain consciousness in terms of any *one* of these elements without the needed contextualization in relation to the others—as for example people sometimes speak as if complex dynamical systems could automatically explain what makes something conscious.

Consciousness requires that the interest in the future is neurophysiologically instantiated so as to give rise to an image of a possible future. In the case of perceptual consciousness, this same activity resonates with afferent activity stimulated by input. The interest in the future, the forming of the image, and the processing of the sensory or sensorimotor data must all resonate with each other as the motivation gives rise to the “looking-for” which then directs our attention. The degree of resonance among these activities corresponds to the vividness of the consciousness. The quality of experience as “stretched out” over the three temporal moments—present, immediately retained just-past, and anticipated future—also leads to a feeling of an inseparable blending of feeling with the intending of an object, and even a fuzziness in distinguishing subject from object (as when we attribute the phenomenal redness of an object to the object itself, as if the red were pasted to the surface of the object, or when we attribute the mood that an object produces in us to the object it-

self). This blending accounts for a good deal of the fact that we usually focus our attention on the object of experience rather than our subjective experiencing of it. Here again, what is subjective and what is objective about our experience of reality is in principle not a clear or sharp distinction.

Consciousness cannot be understood as a passive causal result of the actions of the substratum elements that make up the components forming the higher-order process that allows for the active nature of consciousness. Consciousness is self-organizational in a biological way. The higher-level process seeks out, appropriates, re-arranges, and within certain limits even replaces the micro-level constituents needed to enact its dynamical patterns of organization. Consciousness is not simply caused by its micro-level physical components, but neither is it separable from them. The enactive approach meets this demand by conceiving of consciousness as a self-organizing process that is inseparable from its micro-constituents in the sense that it could not occur without them, yet it is not passively caused by the actions of those substrata. But this theory in turn requires a careful scientific grounding of the causal relations required to maintain the regularly observed “causal closure” at the micro level, as our next section will discuss.

The compatibility of dynamical systems with micro-level mechanistic explanation

There is now a mass of empirical evidence pointing to the thesis that processing occurs in a *conscious* way only when it is very globally distributed in the brain. For example, we know that, when impulses caused by optic stimulation set up patterns of activity in the occipital lobe, but without coordinated limbic and frontal-cortex activity, no perceptual consciousness results from the occipital activity (Posner 1990; Damasio et al. 2000; Eslinger and Damasio 1985; Nauta 1971; Luria 1980). Similarly, the transition from sleep to waking consciousness requires that the activities of the hypothalamus and cortex achieve a pattern of synchronization or coordination which was not present during sleep (Asimov 1965: 193; Ellis 1986: 46-52). When we are conscious of dream images during sleep, both efferent and afferent activity throughout the brain are detected, whereas during non-dreaming sleep both the afferent activity and some of the efferent activity are comparatively much less pronounced (Winson 1986: 46ff; Restak 1984: 315-333; Richardson 1991; Jouvet 1967).

Another example is provided by the 1/3-second time delay from the activation of the occipital lobe (in response to a novel stimulus) to the presence of perceptual consciousness of the object, the latter correlating with coordinated limbic, frontal, parietal, and occipital activity (Aurell 1989; Runeson 1974: 14; Srebro 1985: 233-46). EEG and other electrical measures show that parietal

activation does not occur until 1/3 second after the occipital activity, which by itself does not produce consciousness of the object. As we saw earlier, the activation does not merely “travel” from the occipital to the adjacent parietal area; if it did, the distance involved would be traversed in much less time than 1/3 second. Instead, before perceptual consciousness can occur, the limbic system must be aroused, and it in turn must activate the frontal lobe to begin formulating questions about what the nature of the interesting or important environmental stimulus might be, which then activates ideas and/or images in the form of anticipations of possible perceptions with the help of the parietal lobe (Ellis 1995, Ch. 1; Luria 1980; Posner and Rothbart 1992). Only as a result of this symphonic orchestration of global activity can the activity of the parietal lobe be matched against what is happening in the occipital lobe to see whether the image or idea hypothesized is actually instantiated in the environment. If so, perceptual consciousness of the corresponding object occurs. If not, a mere mental image of the object experienced as non-present occurs in consciousness. In either event, consciousness occurs only when brain activities are globally coordinated. What these examples and many others suggest is that consciousness requires globally distributed processes in the brain, combining local mechanisms which under different circumstances would be active in various non-conscious processes. Luria (1980), Posner (1990), Posner and Rothbart (1992), and many other neurologists are now convinced that the prefrontal lobe plays a crucial role in the process of directing attention to what is important. What makes it do so is that it receives rich input from the limbic system (importantly involved in motivational feelings and other “valuations”), and then sends signals that coordinate the remainder of the cortex to be consciously aware of the arousing situation and to devise ways to deal with it.

Neither dualism, nor psychophysical identity theory, nor epiphenomenalism works as an explanation of the relation between consciousness and its physiological correlates, because the modernist concept of atomistic-reductionism does not allow a process to affect the behavior of its own substratum elements, but requires that a process must be caused by the interaction of the discrete movements of its substratum elements, each of which has a sufficient causal explanation of its own, so that the pattern of consciousness, paradoxically, can have no causal power.

But the enactive approach avoids this aspect of the problem of causal closure, because it does allow that a *process* can have causal power. The necessary and sufficient causal antecedent of an outcome in a self-organizing dynamical system is a previous *pattern* of functioning, which can be realized by alternative sets of micro-constituents. None of the particular micro-constituents are strictly necessary, because others could have been used to maintain the same process. But not only is the process multiply realizable. In self-organizing systems, the entity defined by the overall process *makes use* of the fact of multi-

ple realizability by seeking out, appropriating, and replacing the micro-constituents needed to keep the process going in its definitive pattern.

Causal closure is nonetheless still maintained in this case because a process is inseparable, *under the given circumstances* (but only under those given circumstances), from the behavior of its actual micro-constituents or substratum elements. So the process will also be necessary and sufficient for whatever its substratum elements are necessary and sufficient for, under those specific circumstances. Yet this does not necessarily imply that the process is caused by its substratum, or that it is identical with it. Many things are true of a process which are not true of its substratum elements, even taken collectively. For example, a wave on the ocean may travel many miles in a horizontal direction, while its substratum elements, the movements of particles of water, are very small vertical oscillations.

The reason the process can have this kind of power to rearrange its own substratum elements in such a “plastic” way is that causal antecedents are necessary and sufficient for their consequents *only when certain background conditions are in place*. The power of the process is that it can alter the background conditions themselves, because the process is arranged in such a way that shunt mechanisms will kick in when a preferred causal relation (such as a chemical reaction) fails to occur. There are multiple cause-and-effect sequences that can lead to the same pattern being maintained, depending on which sets of background conditions are in place. What is really necessary and sufficient for a given outcome is that a previous *process* have been in place—not simply a series of micro-reactions, as if those micro-reactions were not dependent on a larger set of background conditions being in place.

The process-substratum relation in the case of consciousness is different from the relationship between a wave and the physical medium through which the wave passes, such as a sound wave, in one crucial respect. Consciousness, unlike a sound wave or a wave in the ocean, is a purpose-directed process. Merleau-Ponty (1942/1967: 47ff) defines a purposeful activity as one in which the organism's overall pattern of activity acts in such a way as to rearrange and readjust its various parts in order to maintain or enhance the overall pattern. Purely mechanical processes do not seem to behave in this way. A thermostat, while it will adjust its overall pattern to feedback from the environment, does not seem to be a purpose-directed system because, when one of its parts ceases to function or is removed, the thermostat does not act in such a way as to replace the missing part or try to compensate for its absence; it simply quits functioning. The thermostat does not “care,” in this non-conscious sense of “care,” whether it achieves its ultimate objective or not. It functions or not purely as an additive juxtaposition of the functioning of its parts.

It becomes increasingly clear, as we study the brain, the ecosystem, and the concept of “living organisms” in biology, that at least many patterns of activity maintain their organizational structure across replacements of their own substrata. As Merleau-Ponty suggests, an organism will often rearrange the overall configuration of its parts if an imbalance is created in one part which disrupts the functioning of the whole. “Forms’ . . . are defined as total processes whose properties are not the sum of those which the isolated parts would possess. . . . We will say that there is form whenever the properties of a system are modified by every change brought about in a single one of its parts and, on the contrary, are conserved when they all change while maintaining the same relationship among themselves” (Merleau-Ponty 1942/1967: 47).

Examples of self-directed neurophysiological reorganization following localized brain injury or trauma can be found in Restak (1984: 360ff). Kandel and Schwartz (1981) also emphasize the “plasticity” of the brain in reorganizing itself to accomplish its objectives by getting around disruptions in one way or another. For example, if brain cells of an embryo are transplanted to a different region of another embryo, they are transformed into cells appropriate to that region. This plasticity in the realizability of the mental functions of living beings has been discussed by Bechtel and Mundale (1999), Gillett (2003), Horgan (1992) and Bickle (2006). An example of multiple realization is that the organism's desires intend to remove the inevitability of electrochemical imbalance within the organism, not merely by eliminating this or that electrical imbalance (for example, in cases where to do so would only transmit the imbalance from one part of the nervous system to another, or from one bodily system to another), but rather by changing the context which renders the imbalance inevitable—for example, by spatially removing the entire organism from the disturbing stimulus, by destroying the disturbing stimulus, or by finding or creating a stimulus in relation to which the whole organism's balance can be restored.

Conclusion

The twentieth century philosophy of mind made every effort to remain tenaciously bottom-up. Cognitive functions have been explained as “responses” to incoming “stimuli,” with the stimuli combining in complex ways to mechanically cause the response. The response is thus a purely passive change, brought about by the stimulus. As in the characteristic twentieth century approach to natural science, here too the only inertia is an inertia of passivity; nothing would move or change unless acted upon by an outside force.

In order to overcome the problems we have just outlined, an adequate conception of consciousness must reopen these questions with regard to ontology and the theory of causation. We must develop a theory in which purposeful

processes are able to appropriate their needed substratum elements, rather than merely being passive epiphenomena of them or ontologically identical with them. This in turn will require the development of a workable account of how it is that certain activities can be “purposeful” in a scientifically intelligible universe. The mainstream of twentieth century psychology predominantly turned its back on this problem. Purposeful activity was explained away as a purely mechanical process that only *appears*, anthropomorphically, as if it were purposeful. The standard explanation was that we view a mechanical process as if it were purposeful because we view it as if it were conscious, like ourselves, and we imagine that if we were to engage in that activity, we would be doing so with the consciousness of some purpose in mind. But to characterize a process as purposeful is not to anthropomorphize. The human organism was purposeful before it was conscious—for example, during embryonic development. Consciousness is not necessary to purposefulness, even in the human organism. So purposefulness cannot be explained simply as the addition of consciousness to a process that otherwise could be explained simply as one that displays certain “tendencies” to accomplish certain results, as if the only difference between a purposeful and a non-purposeful process were that, in the latter, there is conscious awareness of an underlying sequence of purely passive reactions. The new theories of self-organization that are now available can be applied to understanding the causal power of a process over its own components, and thus the teleological nature of consciousness as well as many other purpose-directed processes in nature, such as the migration of birds, or the balance of ecosystems.

References

- Asimov, Isaac. 1965. *The Human Brain*. New York: Mentor.
- Aurell, Carl G. 1989. Man's triune conscious mind. *Perceptual and Motor Skills* 68: 747-754.
- Bachmann, Talis. 2000. *The Microgenetic Approach to the Conscious Mind*. Amsterdam: John Benjamins.
- Bechtel, William, and Jennifer Mundale. 1999. Multiple realizability revisited: Linking cognitive and neural states. *Philosophy of Science* 66: 175-207.
- Bickle, John. 2006. Multiple realizability. *The Encyclopedia of Cognitive Science*. New York: Wiley.
- Bickle, John. 2008. *Psychophysical Reduction: The New Wave*. Boston: MIT/Bradford.
- Chalmers, David. 1995. Facing up to the problem of consciousness. *Journal of Consciousness Studies* 2: 5-22.

- Corbetta, M., F.M. Meizen, S. Dobmeyer, G.L. Schulman, and S.E. Petersen. 1990. Selective attention modulates neural processing of shape, color and velocity in humans. *Science* 248: 1556-1559.
- Cytowic, Richard. 1993/1995. *The Man Who Tasted Shapes*. New York: Warner.
- Damasio, Antonio. 1994. *Descartes' Error*. New York: Putnam.
- Damasio, A.R., Grabowski, T.J., Bechara, A., Damasio, H., Ponto, L.L., Parvizi, J. 2000. Subcortical and cortical brain activity during the feeling of self-generated emotions. *Nature Neuroscience* 3: 1049-1056.
- Dascal, Marcelo. 1987. Language and reasoning: Sorting out sociopragmatic and psychopragmatic factors. J.C. Boudreaux, B. W. Hamill, and R. Jernigan, eds. *The Role of Language in Problem Solving* 2. Elsevier: North-Holland, 183-197.
- Dennett, Daniel. 1996. *Kinds of Minds*. New York: Basic Books.
- Ellis, Ralph D. 1986. *An Ontology of Consciousness*. Dordrecht: Springer/Kluwer/Martinus Nijhoff.
- Ellis, Ralph D. 1995. *Questioning Consciousness: The Interplay of Imagery, Cognition and Emotion in the Human Brain*. Amsterdam: John Benjamins.
- Ellis, Ralph D. 2005. *Curious Emotions*. Amsterdam: John Benjamins.
- Ellis, Ralph D., and Natika Newton. 1998. Three paradoxes of phenomenal consciousness. *Journal of Consciousness Studies* 5: 419-42.
- Ellis, Ralph D., and Natika Newton. 2010. *How the Mind Uses the Brain*. Chicago: Open Court.
- Eslinger, Paul, and A.R. Damasio. 1985. Severe disturbance of higher cognition after bilateral frontal lobe ablation. *Neurology* 35: 1731-1741.
- Farah, Martha. 1989. The neural basis of mental imagery. *Trends in Neuroscience* 12: 395-399.
- Freeman, Walter. 2001. *How Brains Make Up Their Minds*. New York: Columbia University Press.
- Gendlin, Eugene. 1992. Thinking beyond patterns. B. den Ouden and M. Moen, eds. *The Presence of Feeling in Thought*. New York: Peter Lang.
- Gibson, J.J. 1986. *An Ecological Approach to Visual Perception*. Hillsdale, N.J.: Lawrence Erlbaum.
- Gillett, Carl. 2003. The metaphysics of realization, multiple realizability, and the special sciences. *The Journal of Philosophy* 100: 591-603.
- Goldstein, Irwin. 1994. Identifying mental states: a celebrated hypothesis refuted. *Australasian Journal of Philosophy* 72: 46-62.
- Globus, Gordon. 1992. Towards a noncomputational cognitive neuroscience. *Journal of Cognitive Neuroscience* 4: 299-310.
- Globus, Gordon. 2003. *Quantum Closures and Disclosures*. Amsterdam: John Benjamins.

- Haines, D., E., Dietrich, E., Mihailoff, G.A., and McDonald, E. F. 1997. Cerebellar-hypothalamic axis: Basic circuits and clinical observations. J. Schmahmann, ed. *The Cerebellum and Cognition* (pp. 84-110). New York: Academic Press.
- Hanze, Martin and Friedrich Hesse. 1993. Emotional influences on semantic priming. *Cognition and Emotion* 7: 195-205.
- Helmholtz, Hermann. 1962. *Helmholtz's Treatise on Physiological Optics*, J.P.C. Southall (trans). New York: Dover.
- Hull, C. L. 1952. *A Behavior System*. New Haven: Yale University Press.
- Ito, Masao. 1993. Movement and thought: Identical control mechanisms by the cerebellum. *Trends in the Neurosciences* 16: 448-450.
- Jackendoff, Ray. 1987. *Consciousness and the Computational Mind*. Cambridge: MIT Press.
- Jeannerod, M. 1994. The representing brain: Neural correlates of motor intention and imagery. *Behavioral and Brain Sciences* 17: 187-244.
- Jeannerod, M. 1997. *The cognitive neuroscience of action*. Oxford: Blackwell.
- Jibu, Mari, and Kunio Yasue. 1995. *Quantum Brain Dynamics and Consciousness*. Amsterdam: John Benjamins.
- Johnson-Laird, Philip N., and R.M.J. Byrne. 1991. *Deduction*. Hillsdale, N.J.: Erlbaum.
- Jouvet, Michael. 1967. Neurophysiology of the states of sleep. *Physiological Review* 47: 117-127.
- Kandel, Eric, and James Schwartz. 1981. *Principles of Neural Science*. New York: Elsevier-North Holland.
- Kauffman, Stuart. 1993. *The Origins of Order*. Oxford: Oxford University Press.
- Kelso, Scott. 1995. *Dynamic Patterns: The Self-organization of Brain and Behavior*. Boston: MIT/Bradford.
- Kim, J. 1992. Multiple realization and the metaphysics of reduction. *Philosophy and Phenomenological Research* 52: 1-26.
- Kim, J. 1993. The nonreductivist's troubles with mental causation. J. Heil and A. Mele, eds. *Mental causation* (pp. 189-210). Oxford: Oxford University Press.
- Kim, J. 1998. *Mind in a Physical World: an Essay on the Mind-body Problem and Mental Causation*. Cambridge, Mass.: MIT Press.
- Lavy, Edith and Marcel van den Hout. 1994. Cognitive avoidance and attentional bias: Causal relationships. *Cognitive Theory and Research* 18: 179-194.
- LeDoux, Joseph. 1996. *The Emotional Brain*. New York: Simon & Schuster.
- Legrenzi, P., V. Girotto, and P.N. Johnson-Laird. 1993. Focussing in reasoning and decision making. *Cognition* 49: 37-66.
- Lethin, Anton. 2002. How do we embody intentionality? *Journal of Consciousness Studies* 9: 36-44.

- Lethin, Anton. 2005. Covert agency with proprioceptive feedback. *Journal of Consciousness Studies* 12:96-114.
- Levine, Joseph. 1983. Materialism and qualia; the explanatory gap. *Pacific Philosophical Quarterly* 64: 354-61.
- Logan, G.D. 1980. Attention and automaticity in stroop and priming tasks: Theory and data. *Cognitive Psychology* 12: 523-553.
- Luria, Alexander R. 1980. *Higher Cortical Functions in Man*, 2nd ed. New York: Basic Books.
- Mack, A., and Rock, I. 1998. *Inattentional Blindness*. Cambridge: MIT/Bradford.
- McHugh, D.E. and A.T. Bahill. 1985. Learning to track predictable target waveforms without a time delay. *Investigative Ophthalmology and Visual Science* 26: 932-937.
- Merleau-Ponty, Maurice. 1942/1967. *The Structure of Behavior*. A. Fischer (trans). Boston: Beacon.
- Merleau-Ponty, Maurice. 1941/1962. *Phenomenology of Perception*. Colin Smith (trans). New York: Humanities Press.
- Monod, Jacques. 1971. *Chance and Necessity*. New York: Random House.
- Moore, G.E. 1900/1956. *Principia Ethica*. Cambridge: Cambridge University Press.
- Natsoulas, Thomas. 1994. What is wrong with appendage theory of consciousness. *Philosophical Psychology* 6: 137-154.
- Nauta, Walle. 1971. The problem of the frontal: A reinterpretation. *Journal of Psychiatric Research* 8:167-187.
- Newton, Natika. 1993. The sensorimotor theory of cognition. *Pragmatics and Cognition* 1: 267-305.
- Newton, Natika. 1996. *Foundations of Understanding*. Amsterdam: John Benjamins.
- Newton, Natika. 1982. Experience and imagery. *Southern Journal of Philosophy* 20: 475-487.
- Newton, Natika. 1989. Visualizing is imagining seeing: A reply to White. *Analysis* 49: 77-81.
- Newton, Natika. 1991. Consciousness, qualia, and reentrant signalling. *Behavior and Philosophy* 19: 21-41.
- Newton, Natika. 1992. Dennett on intrinsic intentionality. *Analysis* 52: 18-23.
- Newton, Natika. 2000. Emotion in a dynamical system. R. Ellis and N. Newton, eds. *The Caldron of Consciousness*. Amsterdam: John Benjamins.
- Noë, Alva. 2006. *Action in Perception*. Boston: MIT/Bradford.
- Panksepp, Jaak. 1998. *Affective Neuroscience*. Oxford: Oxford University Press.
- Panksepp, Jaak. 2012. defense of multiple core affects. P. Zachar and R. Ellis,, eds. *Categorical and Dimensional Models of Affect*. Amsterdam: John Benjamins.

- Pardo, J.V., P.J. Pardo, K.W. Janer, and M.E. Raichle. 1990. The anterior cingulate cortex mediates processing selection in the stroop attentional conflict paradigm. *Proceedings of the National Academy of Sciences* 87: 256-259.
- Posner, Michael I. 1980. Orienting of attention. *Quarterly Journal of Experimental Psychology* 32: 3-25.
- Posner, Michael I. 1990. Hierarchical distributed networks in the neuropsychology of selective attention. A. Caramazza (ed). *Cognitive Neuropsychology and Neurolinguistics: Advances in Models of Cognitive Function and Impairment*. New York: Plenum, 187-210.
- Posner, Michael I., and S.E. Petersen. 1990. The attention system of the human brain. *Annual Review of Neuroscience* 13: 25-42.
- Posner, Michael I. and Mary K. Rothbart. 1992. Attentional mechanisms and conscious experience. A.D. Milner and M.D. Rugg, eds. *The Neuropsychology of Consciousness*. London: Academic Press.
- Restak, Richard. 1984. *The Brain*. New York: Bantam.
- Rhodes, Gillian and Tanya Tremewan. 1993. The Simon then Garfunkel effect: Semantic priming, sensitivity, and the modularity of face recognition. *Cognitive Psychology* 25: 147-87.
- Richardson, John. 1991. Imagery and the brain. C. Cornaldi and M. McDaniels, eds. *Imagery and Cognition*. New York: Springer, 1-46.
- Rolls, E. 1999. *The Brain and Emotion*. Oxford: Oxford University Press.
- Runeson, Sverker. 1974. Constant velocity—not perceived as such. *Psychological Research* 37: 3-23.
- Schiffman, Harvey Richard. 2001. *Sensation and Perception: An Integrated Approach*. New York: John Wiley.
- Smart, J.J.C. 1970. Sensations and brain processes. *Philosophical Review* 68, Reprinted in C.V. Borst, ed. *The mind-brain identity theory* (pp. 141-56). London: Macmillan.
- Spence, K. W. 1956. *Behavior Theory and Conditioning*. New Haven: Yale University Press.
- Spence, S.A., and Frith, C.D. 1999. Towards a functional anatomy of volition. *Journal of Consciousness Studies* 6: 11-29.
- Springer, Sally, and Georg Deutch. 1989. *Left Brain, Right Brain*. New York: W.H. Freeman.
- Srebro, Richard. 1985. Localization of visually evoked cortical activity in humans. *Journal of Physiology* 360: 233-246.
- Thelen, E., and Smith, L. 1994. *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, Massachusetts: MIT Bradford.
- Thompson, Evan. 2007. *Mind in Life*. New York: Belknap.
- Valenstein, Elliot. 2002. *Blaming the Brain*. Chicago: Free Press.

- Varela, Francisco, Evan Thompson, and Eleanor Rosch. 1991/1993. *The Embodied Mind*. Cambridge: The MIT Press.
- Winson, Jonathan. 1986. *Brain and Psyche*. New York: Random House.
- Wrightson, Philip. 1989. Management of disability and rehabilitation services after mild head injury. In H.S. Levin, H.M. Eisenberg, and A.L. Benton, eds. *Mild Head Injury*: 245-256. Oxford: Oxford University Press.