The Journal of Mind and Behavior
Autumn 2005, Volume 26, Number 4
Pages 255–280
ISSN 0271–0137

The Unity of Consciousness: An Enactivist Approach

Ralph D. Ellis

Clark Atlanta University

and

Natika Newton

Nassau Community College

The enactivist account of consciousness posits that motivated activation of sensorimotor action imagery (through efferent activity) anticipates possible action affordances of environmental situations, resulting in representation of the environment with a conscious "feel" associated with the valences motivating the anticipations. This approach makes the mind–body problem and the problem of mental causation easier to resolve, and offers promise for understanding how consciousness results from natural processes. Given a process-oriented understanding of the way many systems in non-conscious nature are "proto-motivated" toward realizing unactualized possibilities, and can use symbolic objects to "proto-represent" unactualized possibilities, it becomes more clear how self-organizing systems can subserve subjective consciousness. If a system executes, in a unified way, both a proto-desire and a proto-representation of the same unactualized possibility — in order to provide a kind of causal power for the unactualized possibility — then the result is the familiar experience of phenomenal consciousness.

Key words: consciousness, dynamical systems, embodiment

Ideally, a complete explanation of conscious experience would show not only how certain processes in nature, describable in third-person terms, can occur without consciousness; such an explanation would also show why it turns out that those physical processes, when occurring in combination or interaction in the same system, would exhibit all the hallmarks of the familiar *first-person* phenomenon that we experience as consciousness. This paper aims to move us a little closer to such an understanding of how certain objec-

Requests for reprints should be sent to Prof. Ralph D. Ellis, Department of Religion and Philosophy, Campus P.O. Box 1832, Clark Atlanta University, Atlanta, Georgia 30314 or to Prof. Natika Newton, Philosophy Department, Nassau Community College, Garden City, New York 11530–6793. Email: ralphellis@ mindspring.com or natika@worldnet.att.net

tively describable physical phenomena can add up to a subjective experience, with all its first-person ineffability.

We believe that the most fundamental starting point for such an understanding is the notion that consciousness is a type of activity of a self-organizing creature, as many philosophers recently have suggested (Humphrey, 2000; O'Regan and Noë, 2001). Viewing consciousness as an activity, rather than as a passive state in which the subject is affected by afferent stimuli while being somehow aware that it is so affected, seems to offer hope of avoiding certain difficulties of the latter view, which is the traditional one in recent information processing theory, and ultimately traces to the Cartesian passively-observing subject. We shall refer to this traditional information-receiving model of consciousness as the "perceptual model." It thinks of consciousness as being like a recipient of perceptual inputs, observing intentional contents as a perceiver observes that which is perceived. In this model, consciousness tends to end up in the role of the end point of a causal chain, and often seems to be a ghostly observer of information that has already been processed within the physical brain–body system.

By contrast, if consciousness is an activity, or as we shall say, *enacted*, then its inaccessibility to objective observers is no longer a problem: only the agent of an activity can perform it. On the traditional model, consciousness is a way of experiencing things, a special way of receiving information; the question immediately arises as to why only the conscious subject is able to receive this information. As Jackson (1986) points out, an experimenter's empirical knowledge of someone's brain, even if complete, would not by itself reveal what the person's subjective experience "is like," or what it feels like for that subject — or indeed that it feels like anything at all. On the enactivist approach, asking why only the subject can enjoy her conscious experiences would be like asking why only a dancer can perform a particular instance or token of the act of dancing. It is conceptually meaningless to suggest that a second person can simultaneously perform a given agent's bodily action.

It has been objected, however, that the move from a passive to an enactive view is of no help in crossing the explanatory gap between physiology and phenomenology: " 'ways of acting' are by no means closer to experiential features than 'internal representations' are" (Kurthen, 2001). The zombie argument concerning phenomenal consciousness can easily be applied to a crude claim that consciousness is an activity: it seems possible that a creature could perform this, or any, sort of activity with no associated phenomenal experience.

In this paper we sketch an enactivist theory of consciousness that addresses, in addition to the functional aspects of consciousness, the question of phenomenal experience. We hold that consciousness is emergent in sufficiently complex self-organizing beings, and that its function is to enable the being to anticipate the future effects of possible actions, so as to make an optimal

selection. The account is enactivist in holding that anticipation is an essential component of the more global activity of self-maintenance and growth or development in the face of challenging environmental conditions. Complex creatures like human beings are able to test possible actions in advance of performing them by means of sensorimotor imagery — activated memory images of the experiences of such actions and their past consequences, understood not in passive perceptual terms, but in terms of sensorimotor and proprioceptive imagery (a kind of "feeling forward") of the consequences of actions for the organism (Newton, 1982, 1993, 1996). The consequences in question are in turn understood, not primarily in external terms, but in terms of potential future organismic actions that they would or would not afford. Images of objective events in the world are secondary in this context, and they are not presupposed by the sensorimotor imagery needed for this process; on the contrary, conscious imagery of the world, in our view, is a *consequence* of the sensorimotor understanding of the environment's action affordances, which in turn is a consequence of imagining ourselves executing those possible actions through sensorimotor action imagery. The activation of the sensorimotor action imagery is subserved by various brain areas, such as the motor cortex and the cerebellum, and is normally the precursor of actual performance. In imaginatively contemplating the actions, their initial stages are activated and then inhibited from overt realization by effector systems (Jeannerod, 1994, 1997). Our account holds that conscious experience is fundamentally and essentially the imagined performance of sensorimotor activity, explicitly or implicitly. That is, much of conscious experience is the explicit contemplation, through proprioceptive and action imagery, of things that we might do (Newton, 1993, 1996, 2000, 2001). Conscious experience can also involve *implicit* contemplation of action, in cases where action imagery underlies reasoning about abstract matters, such as logical or mathematical questions (Ito, 1993; Johnson, 1987; Newton, 1996).

The above claims do not address the question of why such activity "is like" something for the organism — why it results in phenomenal experience — but they do provide a basis for such an explanation. We propose an account of phenomenal experience at the conclusion of the paper. In anticipating actions and their possible outcomes, subjects must make use of working memory of previous action images, in which multiple sensorimotor states are represented. It is necessary for these alternative states to be held "on line" simultaneously, or nearly so, so that the subject can compare them for selection purposes (Damasio, 1994, 1999). Not only are possible future states of the active organism represented, but past and present ones are as well. When we think about what to do next, we must be aware of what we are doing now, and what we were doing in the immediate past, since future states are constrained by the past and present ones. It is this simultaneous representing of

different temporal moments that, we shall argue, is a crucial part of the explanation of the ineffable experience of phenomenal consciousness.

We begin with a general discussion of the ontology of consciousness as an enactive process, and we outline a way the process emerges from its component processes. Then we look in more detail at the type of emergent property we hold consciousness to be, and we show how it is that this kind of emergence can provide the familiar but ineffable phenomenal experiences.

Some Basic Concepts

Our general view of the ontology of consciousness is that it is "emergent" out of physical realities in the sense that it is a higher-order dynamical pattern of organization of physical components which, because of its self-organizing nature, is capable of exhibiting "downward causation" once it has emerged. That is, the overall pattern of organization can rearrange the background conditions so that a variety of different lower-level mechanisms may be appropriated to serve as the needed components for the maintenance of the larger pattern. We do not think of this emergent consciousness as a novel substance in the universe, exerting a force on a passive subject; on the contrary, it is a novel *activity* of the subject considered holistically. Structural patterns in the activity of a dynamical system can make a causal difference to the efficacy of various micro-constituents in the holistic context, and sufficiently complex systems can select, appropriate and replace micro-components as needed to maintain structural continuity across replacements and recombinations of the micro-constituents of the structure (Ellis, 1986, 1999, 2000, 2001a, 2001b; Kauffman, 1993). Consciousness correlates with such higher-order processes, and thus has certain causal powers over the parts that make up the whole. We shall argue that this self-organizational structural stability across various micro-level causal constituents provides a unique phenomenal state that alters experience, and hence behavior, and in that sense is a source of causal power. Phenomenal consciousness is a unifying process in sufficiently complex creatures — a process continuous with the way any living organism can sense or have access to proprioceptive and environmental stimuli, and have its actions influenced by them.

The "ineffability" of phenomenal consciousness results from the necessary failure of a conscious subject to identify, in her conscious experience, the components of physical reality that have formed the components of this organizational unity. Since these components can be both predicted and identified from an objective standpoint, it is less useful to think of consciousness as a novel form of existence than as a novel configuration of prior states. Such configurations can be inductively predicted when the subject is understood well enough to know how a subject would react to greater degrees of

complexity, or to more layers of represented experience. In this section, we argue that *dynamical* emergence — the result of a novel combination of quite lawlike physical phenomena in a self-organizational structure — should be distinguished from *metaphysical* emergence — the notion that consciousness cannot be a combination of physical processes that, in isolation from each other, could have existed without consciousness. In our view, consciousness is a combination of several other processes that in some natural contexts do occur in isolation from each other and in the absence of consciousness, yet when these processes interact within a certain structure, this interaction necessarily results in consciousness; it is for this reason that metaphysical dualism can be avoided.

In essence, our argument is that a kind of proto-desire can exist in some self-organizing systems that lack consciousness; similarly, in some living and non-living systems, there can be nonconscious proto-representation of non-present logical possibilities. But when the same living and self-organizing system both proto-desires and proto-represents an unactualized logical possibility, within the same framework of purposeful action, consciousness is *necessarily* the result. In this case, desire and representation occur in their full conscious sense, not just in their "proto" and non-conscious senses.

Explaining Consciousness as a Self-organizing Activity

When we ask for an "explanation" of consciousness, what are we asking for? Do we need only a description of the physical conditions such that, whenever those conditions are met, and only when they are met, there is consciousness? This listing of necessary and sufficient physical conditions alone would not be a satisfactory explanation, because as Levine (1983) points out, we still would not have explained *why* (or *how*) those physical conditions are necessarily accompanied by consciousness. For a description of the physical conditions to explain consciousness in a way that would provide real philosophical understanding, it would ideally be possible for someone receiving the explanation to relate it to her own experience; to recognize that *that* description fits what she knows her conscious states to be like.

There is only one useful alternative to an explanation that merely lists necessary and sufficient physical conditions: one that would enable someone hearing the explanation to understand not only the physical conditions themselves, but also why and how those conditions would create an experience whose subjective character is not captured by the description of the physical components that subserve consciousness; while the explanations of the components do not entail consciousness, their combination must be unimaginable as unaccompanied by consciousness. In other words, we should recognize at the outset that the ineffability of conscious experience might be

an unavoidable feature. It might be impossible for a conscious subject to analyze her own experience, while she is having it, into a combination of states that are the inevitable result of certain underlying brain mechanisms. But if that is impossible, then that impossibility itself should be explicable in a way that is intuitively acceptable, subjectively as well as objectively. Our view is that the latter state of affairs is the case with our own (human) type of consciousness.

Our goal is to explain why certain physical features, when combined in a particular way, will inevitably entail the familiar experience of consciousness. The entailment here will be an inductive entailment between physical events and phenomenological experiences, not a deductive inference from the physical events to the experiences, because to attempt to deduce consciousness from physical elements would require providing a metaphysical explanation as to why the universe is constituted in such a way that consciousness is an inevitable feature of any possible world. Such a sweeping explanation is not called for in philosophy of mind. If we can show, more modestly, that consciousness as we phenomenologically experience it would necessarily have to result whenever certain third-person physical ingredients are combined, then we would provide an explanation that bridges the explanatory gap, but without entailing the idea that the existence of such a thing as consciousness is deductively entailed by some particular physical description (Kelso, 1995, p. 260). Given this aim, it will be necessary to begin with a definition of "consciousness" that is framed in first-person terms — that is, a phenomenological description of what consciousness is like. Only then can we hope to bridge a scientific story to what is really meant by "consciousness."

A phenomenological description of what consciousness "is like" must distinguish it from the non-conscious processing of information that is accomplished by digital computers, robots and the like. With this contrast in mind, a phenomenological characterization must include at least the following elements:

1. Consciousness is colored by an ongoing *affective* dimension, and this affectivity seems to motivate us to "anticipate" experience in the sense that a perceptual event cannot be conscious unless we "look to see" what is there, as Merleau–Ponty (1941/1962, p. 232) says. This point is consistent with the Mack and Rock perceptual experiments in which subjects are not conscious of perceptual data in their visual fields when their attention is distracted by an irrelevant mental task. Receiving input is not enough to create consciousness; as Mack and Rock (1998) put it, "attention, when otherwise engaged, must be captured *before* perception can occur" (p. 18, italics added). We experience things consciously only if we are engaged in a "looking for . . ." and not just a "looking at . . ." (Ellis, 1995).

2. This first point may not be enough by itself to distinguish consciousness from the information processing of a cleverly designed nonconscious com-

puter or robot, but the distinction becomes more sharp when we see that it also implies our second and third points. The second point is that consciousness includes an "imaginative" dimension in the sense that it requires that we contemplate some sort of vague notion of what might be presented *before* it is presented; and this means that consciousness, at least in part, is always imagining possibilities independently of whether they are actually presented to us. It is true that our expectations might be surprised, and they frequently are; but the important point is that, in either case, we are contemplating (anticipating, imagining, wanting, fearing, interested in . . .) a *non-present* event, as an as-yet-unactualized possibility, and only as a consequence of this *awareness of unactualized possibilities* can any actual state of affairs enter our consciousness when we are aware of it. Computers can register events without being aware of unactualized possibilities, but consciousness of events always requires that we also motivatedly anticipate data in a goal-directed manner by imagining, desiring, etc., that which is not present. We are always comparing the actual with the possible when we are conscious of it. While computers can and do compute logical possibilities independent of the actual, they do not require a motivated anticipation of the imaginary in order for the computations to occur.

In Husserl's terms, every experience must include a "meaning intention" if the corresponding "meaning-fulfilling intention" is to register in consciousness (Husserl, 1900/1970, see especially Investigation V; see also Husserl, 1913/1969, 1962). The meaning intention is simply the entertaining of a possibility, which then becomes actualized when we perceive, remember, or further articulate the intentional meaning in terms of which the meaning-fulfilling event (for example, perceptual input) achieves its experienced meaning for consciousness. The meaning intention thus "anticipates" an imagined possible fulfillment.

Christina Schües (1994) describes the anticipatory element in consciousness this way:

> The continual course of experiences takes place as a process of actual anticipations followed by subsequent assurances in which the same object of perception, remembering, etc., is held in awareness and is determined more closely. If an anticipation is not assured, but dissapointed, then I might be surprised and a modalization of my experience takes place . . . . We find here the possibility of a meaning which may even retroactively *inhibit* the already constituted meaning and overlie it with a new one, and hence, transform the experience accordingly . . . . When an object of experiences affects me, it seems to have an affective power which motivates me to perceive it . . . a movement of interest aiming at the object . . . . If the anticipation is disappointed and its motivational power diminishes, then the original mode of the experience is transformed into negation. (pp. 12–14)

If consciousness includes an anticipatory imaginative element of this sort, then we can also say that it always contemplates the possible, not just the actual.

3. In addition to its affective, anticipatory, and imaginative capacities, phenomenal consciousness is also experienced as having a paradoxical relationship to the experienced objective world in at least three main ways. (A) As Locke discussed at length (and probably failed to explain adequately), many perceptual properties seem to be physical properties of their objects or of the spatial array of objects, but in reality are created by consciousness itself. For example, the color "red" seems to be pasted to the surfaces of objects, but in fact does not exist until our nervous systems operate on the perceptual signals that the light frequencies produce only after they reach us. Thus properties like "red" involve the eerie and ephemeral feature of seeming to be "in here" in our consciousness, yet at the same time "out there," pasted to the surface of objects.

(B) Consciousness is similarly paradoxical in the way it seems to be *temporally* extended. Any experience seems to be temporally "thick," in that its experience requires some brief segment of time in which to occur; yet we think of our experience as including an infinitesimal present moment like a point from which past and present stretch out in the form of "memory" and "anticipation": as Husserl (1905/1966) describes this temporal paradox, we seem to be experiencing a minimal amount of anticipation of the continuation of a momentary experience, however brief, in the form of a "protention," along with a minimal "retention" of a just-past dimension of the experience, all rolled into the one "present" experience. Memory and anticipation seem on the one hand to be included *in* the present experience (no matter how brief), and yet also seem to stretch out *from* it in both directions. Just as in the paradox of spatial experience (in which red seems pasted to the surfaces of objects yet also "in here"), the paradox of temporal experience seems to locate aspects of the retained and protended content both "within" and "external to" the present experience.

(C) Conscious experience also presents another paradox: we seem able to direct our movements, yet we also experience our bodies as physical objects, which we observe are subject to regular patterns of causal determination. So we experience ourselves as both free and caused at the same time. Even our decision to direct our attention seems this way: the object seems to "pull" our attention, as if we could not help seeing it, yet we also feel that we are in control of our own attention, and even that we would not have seen the object if we had *not* directed our attention. We also experience our attention as directed by *motivational* concerns, yet the motivation itself is an aspect of our bodies, which again appear to us as both free and determined.

In an earlier paper, we attempted to encapsulate these mysterious phenomenological qualities that characterize conscious experience in the following summary description:

> Conscious experience (by contrast to unconscious information processing) entails an emotionally 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. Incoming data, when combined with anticipatory imagery in conformity with structures specific to the capacities of that organism, becomes conscious experience. 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 or conceptual data. (Ellis and Newton, 1998, p. 432)

Given this characterization of phenomenal experience, we can now ask ourselves: Is there some combination of physical events that are describable in third-person terms, and which could occur *without* consciousness when they occur in isolation from each other, yet whose *combination* cannot be imagined as occurring in the absence of all of the features of conscious experience that we have just described phenomenologically? If there is some such combination, then it will bridge the explanatory gap, at least to the extent that our phenomenological description of consciousness has been adequate. That is, we will have specified a combination of third-person events that cannot be imagined as not accompanied by the phenomenal experience of consciousness as we have described it in first-person terms.

Conscious States and Their "Proto" Counterparts

The explanation we offer is one in which both the juxtaposition of certain physical components of consciousness and the components themselves can be explained by known physical principles. For the purposes of this explanation, we shall make certain ontological assumptions about the physical realm. These assumptions will be aspects of process philosophy (e.g., Whitehead, 1925) and of the theory of self-organization (e.g., Kauffman, 1993; Kelso, 1995) that we think are minimal enough to be relatively noncontroversial. The essence of this argument has three parts.

1. *Proto-desire.* Many processes in nature show a robust and resilient tendency to maintain themselves across replacements of their components; these tendencies in nature we shall call "*proto-desires*," because they are functionally similar to our conscious desires in all physical respects except for those that entail consciousness. Proto-desires, that is, impel an organism to alter its states and environment until a goal state is reached, and then to relax those particular efforts. All self-organizing living systems seek out replacement components on the basis of "proto-desire." Like desires in the psychological sense, which can be conscious, they tend to seek out conditions favorable to a certain homeostatic outcome for the systems that proto-

desire them, and to react in such ways as to incorporate the proto-desired conditions when those are available. Especially important for our purposes are the following properties:

> Proto-desires have an active tendency to maintain not only inertias of motion, but also inertias of organization.
>
> Proto-desires reflect the tendency of a system to seek out substratum elements needed in order to maintain the system's definitive patterns.
>
> Proto-desires tend to convert available substrata to their ongoing patterns of organization.

We do not claim that a proto-desire is a desire. But when both the second and third elements of our explanation are added to proto-desire, the result is desire in the conscious sense.

2. *Proto-representation.* Consciousness entails that a currently non-actualized possibility is "represented" by a system that has desires and aversions in relation to those possibilities, where the process that desires and averts is the medium in which the possibility is represented. There are also "*proto-representations*" — situations in which a copy, replica, or isomorphic encoding of an event occurs. But these proto-representations lack "intentional" reference, in the traditional philosophical sense of being about the states of affairs that they proto-represent. There is nothing about the system in which they occur to make them any more than a very similar or analogous version of the original, as when a photocopy is made from an original.

Even proto-representation, however, can occur either in a trivial or in a more meaningful sense. For example, a digital computer can non-trivially proto-represent something without consciousness. The proto-representation occurs in a non-trivial sense because the proto-representational state does more than just be isomorphic to an input to which the state is causally related. In addition to mere isomorphism and causal relatedness to an input, the computer also uses the isomorphic state to achieve a "goal" that has been defined for the computer (but not *by* the computer, since digital computers are not living, self-organizing systems capable of defining their own goals).

For example, a map on the Mapquest website proto-represents a geographical area in a non-trivial sense, because the map can be used in the context of solving a problem for the user of the computer. This proto-representation is not trivial, in the way that, say, a billiard ball struck by another billiard ball might be said to "represent" the force and trajectory of the first ball. But in order for the map to become a representation in the full sense, rather than a non-trivial *proto*-representation, a designer or user of the computer must

understand how the isomorphism could be used in a purposeful context. A digital computer, then, can have proto-representations in a non-trivial sense, but does not have representations in the full sense. Representation in the full sense requires that the system doing the representing should at least be able to form proto-desires, so as to establish purposeful goals in a non-trivial sense of "purposeful."

Proto-representations are often referred to as if they were representations in the full sense, but most such references fail to make the crucial distinction between the two senses. Philosophers frequently use the term "representation" to refer to a state that in some physical way reflects, or maintains a trace of, its cause. This account is grossly inadequate for intentional representation. Not only will any effect of any cause be a representation on this account (Thelen, Schoner, and Scheier, 2001), but it offers no explanation of how a physical state can represent one of its causes that is picked out from among all the rest. It fails to provide, moreover, any role for the *activity of meaning*, which is at the root of the traditional sense of intentionality. Symbols refer to something when they are meant to do so by their users. Physical causation and isomorphism alone cannot explain this relation.

Notice that this traditional account fails for this same reason across many different usages of the term "representation" — for example, whether one thing represents another by "re-presenting" it, by "standing in for it," or by "symbolically portraying" it. A repeat showing of a TV show may "re-present" the earlier airing of the show, but it does not "represent" it in a non-trivially intentional sense unless someone — the TV producers or the viewers — can remember the original airing of the show. If no one remembers, then the "representation" fails to intentionally represent the earlier one.

Similarly, the body map in the parietal lobe forms images that "stand in for" parts of the body, but the sheer isomorphism of the parietal activity to the body part it stands for is not enough to make it intentionally represent that body part. The isomorphic parietal activity must occur in a broader context of organismic activity in order to be experienced as the image of that body part.

Even when we use terms in a symbolic way — for example, in the statement "the pen is mightier than the sword" — it is not the isomorphism of the pen and the sword that makes them represent what they do. The representational relation results from the way the pen or the sword could be used by a purpose-directed subject.

It is necessary to develop a non-trivial concept not only of representation, but even of proto-representation. If a cancerous cell replicates itself in my brain, the one cancerous cell does not "represent" the other one, or even "proto-represent" it in a non-trivial sense. The reason is that my system is not purposefully *using* this activity to represent the original cancerous cell with

which the new ones are isomorphic. There must be a purposeful activity whose aims are served by some similarity between the original and the copy in order for the copy to bear an intentional relation to the original — to represent it in the full sense. Thus proto-desire is presupposed both by genuine representation in the full sense and by proto-representation in any non-trivial sense. If there were no proto-desire to establish the purposefulness of the activity for which the isomorphic activity is being used, isomorphism alone, even in the presence of a causal relation, would not constitute an intentional relation of representation. We shall argue that when proto-desire and proto-representation (in a non-trivial sense) *of an unactualized possibility* are executed by *the same elements* of a physical system, intentionality and consciousness can result.

Proto-representation in a non-trivial sense differs from representation (in the full sense) in that proto-representation uses the isomorphic state to achieve goals defined *for* the system, but not *by* the system. A similar point can be made with regard to knowledge and proto-knowledge in general. Proto-representation is a specific type of proto-knowledge. Many systems in nature include the property of proto-knowledge. Proto-knowledge occurs when a system gathers information about its environment, but does not entail the "what it's like" component exhibited by the knowledge that conscious beings have. Proto-knowledge occurs any time information is transmitted, as in a digital computer. But such proto-knowledge can occur in either a trivial or a more significant sense. Proto-knowledge occurs in a trivial sense when any causal force is transmitted: the causally affected system "knows" that it has been affected in a certain way, because the information that it has been so affected is "encoded" in the output of its behavior (Jackendoff, 1996). This is trivial proto-knowledge, however, because any outcome can have a number of possible causal antecedents, so the system does not know which possible antecedent has caused the behavioral outcome. More importantly, to extend the notion of "knowledge" to include all causal processes would render the term virtually meaningless. For example, to say that a jar lid "knows" (or even "proto-knows") the temperature of the hot water that we run onto it, and then "knows" (or even "proto-knows") how much to expand in response, would trivialize the meaning of "knowledge."

A more significant form of proto-knowledge occurs when the system forms a proto-representation of the causal antecedents that have affected it in the relevant ways. A system proto-represents a causal antecedent when it forms a pattern isomorphic to that which it proto-represents, while at the same time associating that pattern with the outcomes that it produces for the system. Digital computers, for example, have proto-knowledge in this more significant sense, because not only are they *causally affected* by that which they "proto-know" about, but they also *proto-represent* that which causes these

effects by contextualizing the knowledge in terms of purposes that the system is trying to achieve.

In general, proto-knowledge occurs anytime there is non-conscious information processing. But non-conscious knowing is only "proto"-knowledge, because, as in the case of proto-desire, it does not intentionally refer to its object in a non-trivial sense. If an input causes a computer screen to display a picture, P, of the environmental object, EO, from which the input is derived, the fact that P is a representation of EO, rather than EO being a representation of P, or both P and EO being representations of some third thing, is known by us, the users of the computer, because we know enough about the structure of causation that we can establish that EO is the cause while P is the effect. Thus the difference between knowledge and proto-knowledge (in the non-trivial sense) is that in proto-knowledge a system (a) is causally affected by its object and (b) proto-represents the object that affects it in this way in a non-trivial sense, and also (c) knows that the object in the picture it forms refers to the one exerting the causal effect. Knowledge in the full sense is similar to this except for item (b); rather than merely proto-representing its object, knowledge in the full sense represents its objects in the full sense. But even to *proto*-represent the object in a non-trivial sense requires more than just isomorphism plus causal relatedness.

As we have already suggested, to say that something is represented in the full sense requires more than that it proto-knows that an input has caused a reaction. It requires that the subject of the representation also forms conscious or preconscious sensorimotor imagery of possible actions that could be performed in relation to the environmental object or situation, and this also entails enacting the motivations for those actions. Moreover, it requires that the subject imagine actions toward a symbolizing element that are connected to those that could be performed toward actual environmental elements. For example, in thinking the word "blue," I not only imagine myself saying the word (Joseph, 1982), but also imagine performing the attentive act that could prepare me to form the visual image of blue (Ellis, 1995). That is, I motivatedly "look for" or "anticipate" blue.

It might seem superficially that computers could have true knowledge rather than mere proto-knowledge. A computer can be programmed not only to represent a causal input, but also to know that the representation resulted from the actual causal input and not merely from the logical possibility of the environmental conditions in question. But this would be to assume that computers do actually represent rather than proto-representing their objects. As we saw earlier, digital computers only proto-represent, because the goals they use the isomorphic states to achieve are defined *for* the computer, but not *by* the computer. So it would seem more appropriate to characterize a digital computer as proto-knowing (in a non-trivial sense), although whether

we call it knowledge in the full sense would have no effect on the present argument.

3. *Unactualized possibilities.* We hinted earlier that consciousness must not only desire and represent, but must also desire and represent unactualized possibilities. This additional essential ingredient of "unactualized possibility" must be stressed. It is not enough merely to add proto-desire to proto-representation in order for consciousness to result; what the desire and representation are about must include not only physical environmental objects, but also the unactualized logical possibilities that can be anticipated, sought out, or imagined, against which existing states can contrast and in terms of which they can be used and understood (Ellis, 1995).

In a sense, consciousness is one of the ways some natural phenomena have developed of allowing unactualized logical possibilities to have causal power in the real world. However, this causal power must be understood in a completely physical way if we are to avoid an untenable Cartesian dualism. Our thesis, then, is that consciousness can occur only when the elements of proto-representation and proto-desire (or -aversion) of unactualized possibilities are co-present in the same pattern of activity. At that point, they become both a representation and a desire (or aversion) in the full senses of those terms.

Since the elements of proto-desire and proto-representation are present throughout nature, with or without consciousness, and since the essence of consciousness is that it executes a unified activity that both represents and desires a currently unactualized possibility, we can say that the existence of consciousness necessarily follows from the unified execution within the same system of proto-desire and proto-representation of an unactualized possibility.

The notion of a desire that represents an unactualized possibility can be described in purely objective terms, and the three elements whose combination in the way just mentioned realizes consciousness — proto-desire, proto-representation and unactualized possibility — are *third-person* phenomena. If the explanation captures *both* the functional purpose of consciousness as described above, *and* the phenomenological, what-it's-like character of the third-person phenomena involved, for the subject of the phenomena, then the explanatory gap will have been bridged. That is, we will have physically explained components that, when combined, inevitably must equal consciousness. This claim is asserted, however, not as an analytic one, but as an inductive one that allows all relevant evidence and argument to cohere in a consistent theory.

Given these definitions, we can now clarify what is needed for an explanation of consciousness that bridges the explanatory gap. Essentially, we can begin with certain elements — proto-desire and proto-knowledge, which exist in many places throughout nature and can be given purely objective

descriptions — and then show that when these elements come together in a certain way, all the elements in the definition of consciousness are satisfied.

Fitting Consciousness into the Ontological Picture

Consciousness must involve more than mere computation, in the sense of information processing, knowledge, or proto-knowledge. We can approach the explanation of consciousness by asking what is needed for consciousness above mere computation, while at the same time computation seems to be an element in consciousness.

Kinds of Computation

First let's be clear as to what is meant by "*computation.*" Computation can occur in an important or a trivial sense. In the *trivial* sense, computation is based on the correspondingly trivial type of proto-knowledge. In this sense, any transfer of causal effect from one physical event to another involves "computation" or a processing of information. For example, when we run hot water onto a jar lid, the jar lid "computes" the temperature of warm water we run onto it and "responds" to this "input" by expanding to the appropriate extent. In the same trivial sense, a thermostat "computes" the desired temperature and "responds" by throwing a switch for the necessary amount of time. To equate "computation" in this trivial sense with thinking per se would constitute what we might call "jar lid computationalism" — i.e., a view which would ultimately equate any cause-and-effect mechanism in nature (such as the expansion of a jar lid in response to a temperature increase or the thermostat's throwing of the switch) with a "computation," and thus would fail to distinguish genuinely "mental" events like thinking from completely non-mental ones like the response of the jar lid to the temperature increase (Ellis, 1995).

But we must also work toward a more significantly *mental* sense in which the brain "computes" (which is more reasonably equated with "thinking"); in some cases this mental computation is indeed unconscious, but these unconscious computations are often habituated or sedimented from past conscious processes, and others are straightforwardly conscious, as when we purposely use a logical principle such as *modus tollens* to help us make an inference — although this purposeful use of logical principles may often be minimally conscious (like the deciphering of written music) because it too is habituated and sedimented. Of course, very simple computer systems "compute" only in the same sense in which jar lids and thermostats "compute"; but more elaborately programmed ones, like the human brain, compute in a more important sense, and it is this more important sense that is of concern for our purposes.

The more *important* kind of computation occurs when not only does the computing system exhibit behavioral outputs resulting from transformation of the input (as occurs in the trivial example of the jar lid), but the system accomplishes the transformations by using not only its knowledge or proto-knowledge of the inputs, but also by applying a knowledge or proto-knowledge of unrealized logical possibilities to a transformed representation of the inputs. For example, a digital computer, after transforming the input that we key in to an isomorphic proto-representation vis-à-vis one of the computer's "languages," it then operates on that information by making use of its proto-knowledge of a system of *logical possibilities* which are "understood" by the system independently of the specific input received on that occasion — for example, by means of a proto-knowledge of certain logical principles such as *modus tollens* or *modus ponens*, or mathematical principles such as the theorems of algebra or calculus.

Halfway between the important and the trivial sense, there is an *intermediate* case. A computing system may make use of logical or mathematical principles on a "rote memorization" basis rather than with a knowledge of the epistemic derivation of them (as unfortunately many math students are encouraged to do). Although the theorems are not known to be true in this case, the system does still use them *as if* they were known to be true. To the extent that the epistemic bases of logic and math are still in doubt, that is what we humans do as well; but we can feel confident in using the theorems as if they were known to be true, because we also know that they and the assumptions on which they are based have proven somewhat reliable. So there can be various degrees to which "computation" occurs in the significant rather than the trivial sense, but in all of them, there are fundamental elements and relations between them that are not found in computation in the most trivial cases.

To the extent that computation occurs in a non-trivial sense, the system is doing more than simply reacting to inputs, and what it knows about the input is known in relation to more than just the input. The input is known in relation to knowledge or proto-knowledge of logical possibilities. To know (or proto-know) something about a logical possibility is to know more than is exhausted by any current input. It involves knowing about *uninstantiated* possibilities as well as those that are instantiated. And it involves comparing any current input with a previously known system of possibilities which the specific input may instantiate, fail to instantiate, or approximately instantiate.

In some computational systems, the uninstantiated possibilities themselves are not represented, or are not represented very vividly. In others, the uninstantiated possibilities are entertained as phenomena in their own right, independently of the instantiated possibilities to which they are being com-

pared. As a hint to where these distinctions are leading, we shall soon see that the element of proto-desire is the crucial determinant of whether the entertaining of an uninstantiated possibility is fully representational, or only proto-representational.

Proto-desire and the Computation of Unactualized Possibilities

There are some instances in nature where an unactualized possibility is computed (in a significant sense of "compute") as an essential part of the sequence of events that leads to the actualization of that possibility or prevention of its actualization, or to the actualization of some other possibility that is related to the one computed in a more complex way (for example, its approximate actualization, the closest feasible approximation given other constraints, etc.). Consciousness is one such instance. We shall now argue that the computation is a conscious one if the unactualized possibility is represented through the medium of the self-organizational need (a type of proto-desire) which anticipates the unrealized possibility as desirable or undesirable. That is, the representation and the self-organizational need are not two separate processes, but rather the self-organizational need itself enacts a pattern that is intentionally isomorphic to the represented state of affairs.

In order to accomplish such a unified activity, the organism also must be able to include quasi-perceptual imaginative events (representations or proto-representations of unactualized possibilities) and anticipatory emotional ones (desires or proto-desires) into one process, and this requires reentry of perceptual processing into a feedback loop with the emotion-laden anticipation of events, in such a way that the just past phase of the flow of proto-representational activity is superimposed over the present anticipatory phase of this same activity, resulting in the two phases being fully represented as if they were one. Thus a temporal thickness is necessary to consciousness (Ellis and Newton, 1998; Newton, 2000, 2001).

Non-actualized logical possibilities acquire the ability to affect the actual world when the computation that X is possible leads to a different outcome from any that otherwise would have occurred. For example, the realization on the part of Magellan that it might be possible to circumnavigate the globe led him to make the attempt, with subsequent effects on the course of history. This causal impact of unactualized possibilities is reminiscent of the Popper and Eccles (1977) "World III objects" such as the logical processes in the minds of scientists that causally affect experimental and technological events; but in light of our causal analysis of the role of consciousness, we must reject the Popper–Eccles dualistic conclusions. The important point for our purposes is that only through computation of this sort (i.e., in the impor-

tant sense of "computation") does the fact that something is possible have a direct effect on the actual course of affairs, except in the trivial sense that things cannot be actual unless they are also possible.

Although computations involving unactualized possibilities, as their computation affects actual ones, originally occurred in self-organizing systems, they can also occur in non-self-organizing ones, such as digital computers. A computer's "realization" that a certain chess strategy is possible may lead it to actually try the move, thus causing an effect in the realm of the actual. (For example, the computer may defeat a grand master, thus changing the future flow of A.I. grant monies.)

We may therefore think of computation in the important sense as nature's way of allowing consideration of what is possible to have a causal effect on what is actual. Why would nature "proto-want" to do this? Because, as we have seen, proto-desires exist in many places throughout nature, and seek ways to satisfy themselves. This is why the first computational systems (in the important sense of "computation") were also biological ones driven by proto-desire. These proto-desires in turn were merely tendencies for self-organizational processes to seek out the material components needed to sustain their inertias of motion and, at a higher level, to sustain their patterns of activity.

Proto-desire Does Not Necessarily Proto-represent a Non-actualized Possibility

Although very primitive biological systems may compute in a trivial sense, they do not compute in the important sense. They do not even non-trivially proto-represent the unactualized states of affairs that they proto-desire (let alone represent them in the full sense involving intentional referral). As biological systems become less primitive, they do proto-represent *actual* states of affairs in their environments, but they do not proto-represent the *unactualized possibilities* that they desire to a significant extent. Instead, they simply proto-desire an outcome, and wait (perhaps moving around in the environment while they wait and automatically executing certain foraging patterns hard-wired or learned through a blind natural selection process) until the right actualized possibility occurs. The actualized possibility may at this point be non-trivially proto-represented in a perceptual system, but computation in the important sense is still not occurring because the organism did not consider the actualized possibility in relation to a system of knowledge or proto-knowledge of *unactualized possibilities*. Nor did it proto-represent the anticipated but as-yet unactualized possibility in terms of perceptual imagery of the not-yet-actualized possibility.

In these cases, we cannot say that the conditions needed for the existence of consciousness have been met. In principle, nothing has happened to distinguish such a system's mode of information processing from the informa-

tion processing of completely inorganic self-organizing systems such as ecosystems or forming crystal structures. There is a proto-desire (i.e., a mere tendency to maintain a certain pattern across replacements of substratum elements) and a naturally-selected method of registering the presence of the proto-desired outcome, *when it does present itself*, through a proto-representational but not fully representational system. The organism is unaware of the relation of intentionality between the image in its perceptual system and the corresponding image in the environment, because it has no understanding of what the object is, nor any feeling toward the object qua intentional object. To be sure, it has affective responses — further proto-desires (or aversions) or satiation of proto-desires (or aversions) — as *effects* produced *by* the object, but to respond in this way with satiety or further proto-desires or aversions is not to feel intentionally toward the object. If I feel depressed because of chemical imbalances in my brain, but have no idea that the feelings occur with reference to the chemical imbalances, then it is obvious that the feelings do not take the chemical imbalances as their *intentional object*. Similarly, if there is no intentional relation between the feeling and its environmental causes, then the feeling is not a consciousness of that intentional object, even if the object is proto-represented in a nervous system devoid of conscious status.

Mere co-existence of proto-desire and proto-representation does not equal consciousness. We can easily imagine a mechanical device implanted into a very primitive organism, such that the mechanical device non-consciously proto-represents environmental objects in the device's "nervous system," and helps the organism direct its movements toward the object in question, with no more consciousness occurring in the organism than if the device had not been implanted. Nor is there consciousness in the mechanical device in this case, because it is not the device that proto-desires the object, but rather the organism. So in this case, the proto-representation of the device plus the proto-desire of the organism clearly do not add up to consciousness of the object. In what case, then, would it be clear that consciousness is present?

Consciousness: Representation-Through-Desire of Unactualized Possibilities in One Unified Self-organizing System

We defined phenomenal consciousness above in terms of its functions — as a process in which we have some (proto-) desire (or aversion) toward an as yet unactualized state of affairs, and also non-trivially represent unactualized possibilities relevant to the (proto-) desire. In this case, we argued that the proto-desire becomes simply a desire in the familiar sense that conscious beings have. To non-trivially represent in relation to unactualized desires (or aversions) is what it is like to be conscious.

If this definition of consciousness is applied, we can see that consciousness can result from a computational process when a system purposely intends to introduce certain possibilities into the realm of the actual. In primitive self-organizing systems (which we admit do make primitive computations), the system does not *intend* to introduce a non-actual possibility into the world, but merely *tends* to do so. That is, the system has no advance knowledge or intentional anticipation of the desired possible state of affairs. The non-actual possibility is not represented by the system as a possibility that it desires (or averts). In a primitive organism, the actual is represented, but the possible is not; and while the possible may be desired (or averted) in such organisms, it is not even proto-represented in any significant sense. Only when the same non-actual possibility is both desired (or averted) and represented-as-desired (or averted) is there consciousness.

In order to desire and represent the non-actual possibility, an organism must compare the actual with the envisioned possibility, and be aware of the difference. Thus it must cognitively represent an actuality simultaneously with the representation of a non-actual possibility, and bring the two into juxtaposition. It is this juxtaposition that produces the ineffable experience of phenomenal consciousness.

The Phenomenal Feel of Consciousness

We have presented a functional analysis of consciousness, in terms of brain mechanisms that coordinate and combine component states — proto-desire, proto-representation, and proto-knowledge of unactualized possibilities — into a higher-order unified organismic state that is capable of meeting complex and long-range goals. But we have not yet explained why that state should feel the way it does, or, indeed, feel any way at all. We have argued that phenomenal conscious experience plays a crucial role in the attainment of goals, by presenting the organism with anticipations of the experiential outcomes of the various possibilities. But we have not examined the mechanism by which such an experiential state could come into being with its familiar but ineffable qualities. Only an account that explains such a mechanism can bridge the explanatory gap.

We want to argue that phenomenal experience is characterized by a kind of illusion — the illusion of a "temporally thick specious present" — produced by simultaneous representational activities involving distinct temporal moments. What we experience in phenomenal consciousness is time itself, as it affects our self-organizing activity. The experience is unique, in that it occurs only in consciousness, and consciousness is the only state in which we have phenomenal awareness. But while the phenomenon of experienced

temporal thickness is unique, it is not incomparable; within conscious experience, an analogous phenomenon is the visual experience of spatial depth perception.

Binocular vision is necessary for seeing our spatial environment in a single, unified visual experience. Not only the breadth and width of a visual scene is portrayed as existing at a single time, but so is depth, stretching out before the perceiver. Our physical existence in three-dimensional space is functionally relevant to us only with respect to bodily movement: we can extend our limbs, and propel our bodies, indefinitely in three dimensions. The ability to do this, however, is represented not only in bodily action imagery, but also in visual experience. Depth perception presents us with almost infinite possibilities for movement, even when we are relatively inactive: the information necessary for planning actions is given in vision, even though vision itself does not move about in the world. The experience of depth perception is not normally experienced as mysterious, but a little thought will remind us of how remarkable is our ability to project possible movements into our environment by means of a sensory modality that does not itself, literally, project into the environment. What is remarkable is not that we receive information that allows us to calculate possible movements, but rather that the information we receive visually allows us to experience, in a visual modality, what various movements would be like. Vision can do this because the overlapping of two slightly offset visual fields creates an emergent property in our experience. It does not create an ontological novelty, but rather an experiential novelty: we experience or "see" possible movement without actually moving. (The necessary eye movements, of which we are not conscious, are a different matter; here we are referring to gross and deliberate bodily motions.)

The mystery of depth perception is highlighted by the fact that, when we combine two slightly different perspectives on a cylinder to give it depth, this means that we can see slightly *more* than 50 percent of the circumference of the cylinder *all at once*; so what we represent to ourselves is one side of the cylinder — 50 percent of it — which nonetheless shows *more* than 50 percent of it (because we have combined two different perspectives on it). So what we see the cylinder as looking like is actually a geometrical impossibility — a cylinder with more than 50 percent of its circumference showing from one single perspective. Yet all the while, we feel that what we experience ourselves as seeing is a geometrically possible state of affairs.

Another way to notice the paradoxical nature of depth perception is to try the following simple experiment: close one eye and hold up two cylindrical objects such as dry-erase markers, of equal size but different colors, in such a way that one marker completely blocks the other one from view. Now open both eyes, and you notice that the edges of the "hidden" marker cannot be

completely blocked by the one in front of it, no matter how much you try. Your depth perception is allowing you to see "around" the edges of the front marker, making the marker in back appear larger than the one in front, even though they are the same size. If the red marker is in front of the blue one, the blue one appears larger; but if the blue one is in front of the red one, then it is the red one that appears larger. As long as we keep both eyes open, the front marker will always appear *smaller* than the one in back of it, even though it is *closer* to us.

In spite of the paradoxical nature of what depth perception presents to us, the ability to see depth does not normally strike us as an ineffable mystery, but rather as the natural outcome of mechanisms of sight. Why should the same not be true of the experience of temporal extension?

Phenomenal consciousness, in which we experience past, present, and future in a single "specious present," need be no more mysterious than visual depth perception. Arguably the most striking aspect of phenomenal consciousness is the fact that we can linger over it and savor the experience. If phenomenal consciousness lasted only during an instantaneous "now," forever vanishing and being replaced by a completely new "now," consciousness would not be the rich experience it is, and consequently would not be mysterious. We would not have time to know that we were conscious. It might be that such a fragmented form of experience occurs in other species. What is striking in us is that we can dwell on our present conscious experience; it is this ability that allows us to say that consciousness is "like" something for us. We cannot articulate what it is like, except in terms of the sensory qualities of the objects that we consciously experience — properties of external objects and of our own bodies. But there is nothing more to articulate than the very fact that those properties have enough endurance and stability for us that we can know about them, anticipate them, and relate them to prior anticipations that are still vivid for us, while we are sensing them: that is the natural result of the combination of past, present and future in a single representation of the "present."

The experience of seeing spatial distance cannot be articulated except in terms of the *possible motions* that this depth makes possible for us. The fact that it is vision that presents us with those possibilities is a striking and important fact, but it is not an ontologically emergent kind of entity that can be verbally characterized independently of its components. This fact does not bother people, but our inability to nail down consciousness with a precise verbal definition does bother people. Not only should the ineffability of consciousness not bother us any more than the ineffability of depth perception, we argue, but rather this ineffability is exactly what one should expect, given the self-organizing processes that go into presenting us with both our desires

and the means for bringing them to fruition. If a type of organismic state is unique in its effect on the organism, then it stands to reason that it cannot be characterized in exactly the same terms that also refer to other types of organismic states. But that does not mean that it cannot be fully comprehended, from an objective, scientific standpoint.

Notice that the same solution also applies to the paradox of freedom and determinism. If we *are* self-organizing systems that represent possibilities for our future action, then it makes sense that we would feel ourselves as having a choice. In fact, we do make choices, and in self-organizing systems these choices may be free from currently external determination, although this does not mean they are free from the *past* causal determinations that may have led the self-organizing system to have the particular kind of structure it now has, which in turn frees it *from* many *current* determining influences, so that it can act somewhat freely of them. Moreover, if consciousness *is* an aspect of a self-organizing system, this would also imply that consciousness can have every bit of the same *causal powers* as the physical system of which it is an ontological feature. To say that consciousness can have no causal powers is actually an incoherent claim, unless one wishes to embrace dualism, because if consciousness had no causal powers, then it would have to be separate from any physical system, since all physical systems *do* have causal powers.

Conclusion

Our proposal is that phenomenal experience, the subject of the "hard problem," can occur as a natural by-product of the juxtaposition or superposition of distinct temporal moments of a single experienced event. This thesis does not in itself require an enactive view of consciousness. If we are right in our theory of the phenomenal result of this mechanism, then in principle such an explanation could be part of a variety of different accounts of brain activity, not just self-organizing ones. But in any sort of theory that is *not* enactive and self-organizational, the explanation would be ad hoc; it would not follow naturally from the goal-seeking activity of a self-organizing creature. It is because a conscious being must anticipate future experiences, and compare them to past and present ones for purposes of evaluation, that the superposition must occur. When this happens, then the world around us "lights up" in the familiar way of conscious experience. The explanation of this "lighting up" is not complicated or mysterious. It is the inevitable consequence of the prolonging of motivated anticipation, sensory stimulation and imaginative activity for a subject, so that the qualitative properties of the objects stimulating the sensory systems can interact with the imagining of unactualized sensory activity and with the sensorimotor imagining of desired (or averted) actions — and

thus inform the subject's conceptual activities. In this way, they acquire a substantive presence that is familiar to us as conscious experience.

The inevitable objection to such a proposal raises questions about computers. If a computer were designed so that "sensory" stimuli affected it in the same way that has been described here, would we have to grant it consciousness? We see no reason why not, if all the complex functions required by self-organizing conscious entities were present. It would be easy enough to say "I can imagine a machine doing all the things described, without being conscious." But what is really being imagined here? If someone were to take the trouble to describe every function performed by a living conscious being, and propose a machine counterpart for it, then a simple claim that the machine still would not be conscious would not carry much weight. Consciousness might "feel different" to the computer from the way it feels to a flesh and blood creature, but that fact alone is not sufficient to deny it consciousness. And to claim that such a machine would feel nothing at all, in spite of all its animal-like functions, would require more argument than has yet been offered in any literature that we have seen. At the degrees of complexity we are now discussing, intuitions about imaginability provide feeble evidence indeed.

At the same time, a computer meeting all the requirements of a self-organizationally structured being with proto-desire and significant proto-representation of unactualized possibilities would be very different from the nuts-and-bolts computers with which we are now familiar. In effect, they would be human-made biological organisms. And this point is merely a corollary of our overall argument, that "computing" is not the essence of consciousness, but rather a component of it. Consciousness occurs when a self-organizing system represents (imagines) itself as acting toward unactualized possibilities that play a role in terms of organismic desires or aversions.

References

Damasio, A. (1994). *Descartes' error*. New York: Putnam.

Damasio, A. (1999). *The feeling of what happens*. New York: Harcourt Brace.

Ellis, R.D. (1986). *An ontology of consciousness*. Dordrecht: Kluwer/Martinus Nijhoff.

Ellis, R.D. (1995). *Questioning consciousness: The interplay of imagery, cognition and emotion in the human brain*. Amsterdam: John Benjamins.

Ellis, R.D. (1999). Why isn't consciousness empirically observable? Emotion, self-organization, and nonreductive physicalism. *Journal of Mind and Behavior, 20*, 391–402.

Ellis, R.D. (2000). Consciousness, self-organization, and the process-substratum relation: Rethinking nonreductive physicalism. *Philosophical Psychology, 13*, 173–190.

Ellis, R.D. (2001a). Can dynamical systems explain mental causation? *Journal of Mind and Behavior, 22*, 311–334.

Ellis, R.D. (2001b). A theoretical model of the role of the cerebellum in cognition, attention and consciousness. *Consciousness and Emotion, 2*, 300–309.

Ellis, R.D., and Newton, N. (1998). Three paradoxes of phenomenal consciousness: Bridging the explanatory gap. *Journal of Consciousness Studies, 5*, 419–442.
Humphrey, N. (2000). How to solve the mind–body problem. *Journal of Consciousness Studies, 7*, 5–20.
Husserl, E. (1962). *Phänomenologische Psychologie*. Den Haag: Martinus Nijhoff.
Husserl, E. (1966). *The phenomenology of internal time consciousness* [J. Churchill, Trans.]. Bloomington: Indiana University Press. (Based on lectures delivered in 1905)
Husserl, E. (1969). *Ideas* [W.R. Boyce Gibson, Trans.]. London: Collier. (From "Ideen zu einer reinen Phänomenologie und phänomenologischen Philosophie," 1913)
Husserl, E. (1970). *Logical investigations* [J.N. Findlay, Trans.]. New York: Humanities Press. (Originally published 1900)
Ito, M. (1993). Movement and thought: Identical control mechanisms by the cerebellum. *Trends in Neurosciences, 16*, 448–450.
Jackendoff, R. (1996). How language helps us think. *Pragmatics and Cognition, 4*, 1–34.
Jackson, F. (1986). What Mary didn't know. *Journal of Philosophy, 83*, 291–295.
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.
Johnson, M. (1987). *The body in the mind*. Chicago: University of Chicago Press.
Joseph, R. (1982). The neuropsychology of development: Hemispheric laterality, limbic language and the origin of thought. *Journal of Clinical Psychology, 38*, 4–33.
Kauffman, S. (1993). *The origins of order*. Oxford: Oxford University Press.
Kelso, J.A. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge, Massachusetts: MIT/Bradford.
Kurthen, M. (2001). Consciousness as action: The eliminativist sirens are calling. *Behavioral and Brain Sciences, 24*, 990–991.
Levine, J. (1983). Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly, 64*, 354–361.
Mack, A., and Rock, I. (1998). *Inattentional blindness*. Cambridge, Massachusetts: MIT/Bradford.
Merleau-Ponty, M. (1962). *Phenomenology of perception* [C. Smith, Trans.]. New York: Humanities Press. (Originally published 1941)
Newton, N. (1982). Experience and imagery. *Southern Journal of Philosophy, 20*, 475–487.
Newton, N. (1993). The sensorimotor theory of cognition. *Pragmatics and Cognition, 1*, 267–305.
Newton, N. (1996). *Foundations of understanding*. Amsterdam: John Benjamins.
Newton, N. (2000). Conscious emotion in a dynamic system: How I can know how I feel. In R. Ellis and N. Newton (Eds.), *The caldron of consciousness: Motivation, affect, and self-organization* (pp. 91–108). Amsterdam: John Benjamins.
Newton, N. (2001). Emergence and the uniqueness of consciousness. *Journal of Consciousness Studies, 8*, 47–59.
O'Regan, J.K., and Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences, 24*, 883–917.
Popper, K., and Eccles., J. (1977). *The self and its brain*. Berlin: Springer Verlag.
Schües, C. (1994). The anonymous powers of the habitus. *Study Project in the Phenomenology of the Body Newsletter, 7*, 12–25.
Thelen, E., Schoner, G., and Scheier, C. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences, 24*, 1–86.
Whitehead, A.N. (1925). *An inquiry concerning the principles of natural knowledge*. New York: Dover.