Article

How the Brain Makes Up the Mind: An Approach to the Hard Problem of Consciousness

Dan Bruiger*

Abstract

While a strictly scientific solution to the 'hard problem of consciousness' is not possible within the terms of a science that excludes the very subjectivity it seeks to explain, the explanatory gulf between mind and matter can be indirectly bridged by taking the point of view of the organism and its sub-agents. The proposal is that the organism constructs phenomenal experience through a process that resembles how meaning is created in language. For humans, conscious experience involves internal representation—to, for, and by an executive agent (the conscious self). This agent is tasked with monitoring the state of the organism and its environment, planning future action, and coordinating various sub-agencies.

Keywords: Brain, mind, matter, consciousness, hard problem, organism, agency, phenomenal experience, subjectivity, qualia, epiphenomenal, embodiment, intentionality.

1. Introduction

The 'hard problem of consciousness' [Chalmers 1995] is the challenge to explain phenomenality,[1] especially in terms of neural processes. Beginning with Leibniz, it has seemed to many that there is an unbridgeable category gulf between one's conscious subjective experience and physiological events in the brain that presumably cause it. Even though many neural correlates of consciousness have been identified, it remains unclear in principle how physical processes could cause or bring about vivid conscious mental states such as feelings, sensory experience of color, smell, sound, and touch, as well as more subtle experiences such as memories, mental images, volition, thoughts, and dreams. Among the many factors which conspire to make the problem hard, one in particular stands out: the fact that science has long banished the subjective and the mental from its ontology, while insisting that mind must ultimately be reduced to such an ontology.

In recent decades the mental has re-entered scientific discussion through the computational metaphor. The significance of the metaphor lies not so much in assimilating mind to machine as in an engineering or design strategy that puts the designer or programmer in the shoes, so to speak, of the brain.[2] The best approach to understanding how the natural brain constructs

* Correspondence: Dan Bruiger, Independent Researcher, Canada. Email: dbruiger [at] gmail [dot] com

[1] *Phenomenality* is used here to mean phenomenal experience. It refers to the entire domain of what can enter consciousness, which includes sensory experience, emotion, imagination and memory, dreams, thought, etc. In short, everything for which there is "something it is like" to be in that state [Nagel 1974].

[2] *Brain* is used here as shorthand for the entire organism and its embodied relationships to an environment.

phenomenality may not be causal explanation, but rather an engineering approach that puts one creatively in its place. This is rather the inverse of Dennett's [1987/1998] approach of assimilating intentional states to natural design. Here we assimilate natural design to human intentions.

2. The disembodied Cartesian theater

It may be objected that to suppose an already-conscious agent within the head (a homunculus) does nothing to explain consciousness and simply regresses the problem—which is to explain how an agent can be conscious at all, and under what circumstances. It crucially makes a difference whether *agency* is considered from an observer's extrinsic point of view or from the point of view of the agent itself, operating on its own intentionality and behalf. Here we will argue that regression poses no problem and that such a "homuncular" approach can bring us as close as logically possible to understanding how the brain produces experience, or how there can be a first-person point of view at all [MacLennan 2005, sec 3A]. Toward that end, we will explore the active nature of mental agency as original cause—the very opposite of the passivity of the physical, as traditionally conceived in terms of "efficient" causation.

Dennett [1978] articulated objections to the regression and dualism that seem to attend the notion of internal representation, which might apply also to the notion of internal communication. Yet, dualism and regression are involved only when representation or internal communication are conceived as ready-made contents of (human) consciousness—reified as the output of a process. We consider rather the process itself, which is not a thing but an action. The unwanted implication of an inner audience for representations, in some sort of inner theater, comes of thinking of the representation as something already presented to consciousness, which an inner witness encounters as it would an external object. There is no dualism, however, when there is nothing but the agent and its internal actions involved. There is no regression when the homunculus is the conscious investigator and not some element of the system investigated.

There are also legitimate objections to considering the brain in isolation from the body, and to the computational metaphor because it likens the brain literally to a disembodied computer. These are not the issues addressed here. Rather, we assume from the outset that the brain is an organ of the body, embedded in an environment with which its present interactions and evolutionary history are crucial for the organism's behavior and perception, and therefore for any conscious experience.

3. The homuncular strategy

Instead of considering the brain a causal system, let it (like its subsystems) be considered an *agent*: an original first cause with a point of view of its own. The strategy then is to identify with that point of view in order to engage the challenges facing it. While such a strategy need not involve a regression of agents within agents, it may well point to a special executive function within the brain. This function could be likened to the CEO of a corporation, who is responsible for decisions on the highest level, based on "reports" provided by conditionally reliable subalterns. The job of this executive is to monitor and trouble-shoot the overall operation of the system, to plan ahead on various time scales, and to intervene in situations where established protocols are inadequate. One might wonder why this executive function could not be performed *unconsciously*, by some automated sub-routine. What is there about the executive function that requires phenomenality and conscious volition? What is the relationship between these latter? What is it about phenomenality that makes conscious processing sometimes more effective than unconscious? Bluntly, what are qualia good for? And to put it yet another way, why must there be "something it is like" to be this executive (Nagel 1974)? We shall try to answer these questions.

Let us suppose, to the contrary, that an AI (e.g., an "expert" program) performs the executive function adequately, and there is nothing it is like to be that program, which consists of definite algorithms, as do the sub-routines it is supposed to supervise. Yet, we are looking specifically for whatever sets the executive function uniquely apart from such pre-defined systems, and which seems to require consciousness. The question is whether the executive function *can* be fully automated—that is, whether a system that is presumed *not* to involve consciousness can produce all the behavior associated with it. That is a question that may have to be resolved empirically.

Yet, the executive function *could* stand apart from its unconscious counterpart if its algorithms are *not* definite or pre-existing, but provisional and updated in real time. While the immediate attention demanded by some alarum is often accompanied by a pre-programmed quick first response, the executive function is thereby put on notice for future short-term and long-term planning. The monitoring capacity of the executive entails *directing* attention where it is not necessarily demanded, and on a time scale that permits reason and reflection.

An external representation of some sort is often useful for conscious problem solving. In our corporate metaphor, the CEO relies upon various visual aids, in addition to verbal and written reports, to summarize "at a glance" complex information supplied by various departments. That is, the executive makes use of both symbolic and iconic forms of representation in preparing the next action or next step in planning. The key feature of an interface used for real-time guidance is feedback concerning the state of the system in relation to its environment. As an interface, each sense modality offers its own form of representation, with the qualia specific to it.

4. The virtual reality metaphor

Computation has served as a fruitful metaphor to explore the nature of mind. The computer is a virtual machine, running on an actual machine [Boden 2006; Sloman 1993]. The virtual machine can produce a virtual reality—a simulation to be experienced by a human user, often for recreational purposes. This aspect of computation has been proposed [Lehar 2003; Metzinger 2009] as a metaphor for how the brain generates phenomenality. If the mind is a virtual machine running on the wet-ware of the brain, then perhaps what we know as conscious experience is a virtual reality (VR) it produces.

However, an actual VR is a simulation made outside the brain for a human user to experience, via a technological interface with the user's real senses. In the metaphor, the *brain* is both the creator and the user of the experience. To apply the VR metaphor to consciousness, we must switch from thinking of it as an external artifact for human use to thinking of it as a simulation the brain creates for its own use. Moreover, this simulation does not *imitate* external reality, although it is experienced *as* external reality. We do not experience this natural VR as happening inside the head, or as generated by the brain, but as happening in real external space. In other words, we quite normally have, so to speak, an "out-of-brain experience" [Revonsuo 2006]. Indeed, the brain's VR, like the technological artifact, would not be convincing otherwise. By definition, metaphors are never quite literally true, which does not detract from their usefulness. An actual VR may be regarded two ways. Viewed third-personally, it is a *program* running on a machine. On the other hand, it is an *experience*, by definition first-personal.

The brain's simulation is not a *copy* of the external world. The notion of simulation as replication implies comparing two things: the copy and the original. But the brain cannot get outside the skull to directly view the external world and copy it, or to verify the copy by comparing it to the original. Indeed, the brain's VR is an original creation that serves to guide the organism in navigating the external world. It functions like a map that is based on data collected through interaction with whatever lies "outside." It is a *theory* of the external world, whose accuracy can only be measured by its success in permitting the survival of the organism. There is no question of resemblance to a domain to which there is no access! In sum, the brain creates phenomenality as a VR for its own use, as an efficient way to monitor the world, the body, and their relationship. The metaphor can help us accomplish two things: first, to show why consciousness is required for such monitoring; secondly, to come as close as possible to a solution to the hard problem by showing how a physical system that is an agent generates phenomenality.

Metzinger [2009, p108] likens the situation of the brain to that of a "total flight simulator, a self-modelling airplane that, rather than being flown by a pilot, generates a complex internal image of itself within its own internal flight simulator." Since it remains humanly challenging to imagine a "self-modelling airplane," we propose here to have it provisionally flown by a human pilot, who serves as homunculus to stand in for the brain. In this situation of "flying by instrument" [Oatley 1978],[3] the pilot must invent a model of what lies outside a closed cockpit that offers no direct view.

To literally put ourselves in the place of the brain, let us imagine a subject sealed within a windowless compartment, with no permitted exit, connected to the outside world via diverse "sensors" that supply only streams of digital information. The chamber is also equipped with diverse "effectors" of unknown function. The subject finds, through trial and error, that activating an effector changes the sensory input. In real life, "mistakes" in this process may end in the termination of the experiment; for the sake of economy, we treat this as a repeatable game in which the player can "die," and reset to begin again. The subject's challenge is to remain in the game, by correlating incoming data streams with outgoing commands in such a way as to avoid termination.

Trial and error would reveal correlations that can appropriately be encapsulated in a program or hard-wired. Greater flexibility could be achieved by then constructing a theoretical model to

[3] Oatley's simile may predate computerized flight simulators.

account for observed patterns, so that predictions from this model can be tested by trying specific outputs. The model can be consulted to compare its predictions with current input and can be continually updated. A final step would be to construct a VR version of this theoretical model, which becomes an efficient interface with the unknown external world. The model is "true," or "resembles" the outside environment, only in the sense—and to the statistical extent—that it permits continuance in the game. Moreover, the subject in this interactive situation will inevitably come to experience the model *as* a real external world!

Simulation involves a loss of detail, since it runs on a finite program, whereas natural reality may be indefinitely detailed. Perhaps one reason why virtual reality is engaging—and its coarse-graining *works* as entertainment—is that sensory experience (the brain's VR) itself is largely impressionistic. When attending to a visual scene, for example, it is the overall impression that constitutes the experience, not a comprehensive registration of all detail—which is taken in selectively and sequentially, mainly in the foveal area. The impression that we fully and uniformly *see* all that is before the eyes is an illusion. As we will explore further later on, we see only *that* there is detail, often without putting a finer point on it. To the extent that a VR can present fine enough detail to give that sort of impression, it can pass as a substitute for sensory experience.

Virtual reality provides a metaphor to grasp how the brain, encapsulated in the skull, can produce the spectacle we call phenomenality. Ultimately, the engineering task would be to write a virtual reality program for the use of the brain rather than for the human user. As in the Cartesian theater, the metaphor suggests an audience for this spectacle. It would miss the point, however, to think of the brain's virtual reality as "goggles" it puts on to have an experience, or to imagine the vista seen as dancing inside the brain like a holographic image or a miniature replica of the external world. The very point of the metaphor is that the brain *is* the device that generates the virtual reality—for its own benefit, which is not entertainment but survival. Since the brain cannot get outside its compartment to make a comparison, adequacy of the model can only be measured by the evolutionary success the VR facilitates. One may choose to speak of it as illusion or as reality, but there is no standard of comparison except the truths of science, which ultimately also can only be measured by evolutionary success.

5. The role of consciousness

Conscious attention seems to be required in novel, complex, or otherwise mobilizing situations. If the action cannot be adequately performed by rote, if it confronts a novel or demanding situation or requires planning and forethought, conscious attention is brought into play. This suggests that phenomenality makes real-time sensory input available to higher centers for planning or dealing with novelty or emergency. As a function, one role of consciousness is to muster additional resources to deal with situations that are not already handled by existing automatisms, which could then result in improved algorithms. (Ironically, in effect, role of consciousness is to restore a state no longer requiring consciousness [Solms 2014, p179; Lamme 2015, sec11].) Formal processes of reasoning can be more precise (as in mathematical calculation) than informal estimation or guessing, which is faster but probabilistic. The non-conscious visual perception involved in blindsight, for example, is better than random guessing but far from certain. Conscious problem solving is slow because its processing involves many steps and coordinates more areas of the brain that are farther apart. A visual process, for example, seems to be conscious if it involves information shared over a key distance of 10 cm in the human brain [Lamme 2015, sec9].

Consciousness plays a different role than behavioral responses that can occur without it, and presumably involves different or additional neural processes. Some motor tasks are executed "automatically" (though one is sometimes conscious *of* performing them). While awareness of deliberately *initiating* such activity comes after the neural processes that have actually caused it, this awareness serves as the basis for choosing future action, or action in a larger context [Frith and Metzinger 2013]. To put it in more general terms, consciousness plays a specific useful role and is not epiphenomenal.

Walking, or even driving a car, does not require conscious guidance if the route is well known—as sleepwalking dramatically demonstrates. On the other hand—though they may lead to no immediate action—planning, forethought, and reasoning involve conscious attention, just as do alarums that mobilize the body and do require immediate response. Such facts suggest that consciousness serves in an executive role, to monitor behavior or processing by sub-systems that must be coordinated and integrated. It seems to involve a synchronous broadcasting of information globally, especially to parts of the parietal and prefrontal cortices [Kandel 2012] and also seems to be closely associated with particular features of the brain, including the claustrum, a sheet of neural tissue in contact with most of the cortical areas, which might even serve as an "on-off" switch for consciousness [Koubeissi et al 2014].

The agent in this executive role is known to itself and others as the "self" of that individual [Metzinger 2010, p29]. Though it cannot fully take responsibility for the actions of the organism as a whole, it appropriates an identity with the body and is often held accountable by others. It is more than a figurehead but less than an absolute monarch. It is a sort of virtual representative (avatar) of the organism, tasked with specific responsibilities, which include monitoring a real-time inner representation of external reality and the organism's relationship to it. This suggests an interface between this agent and other agencies within the organism, at which conscious experience takes place [Hoffman 2009; Hirstein 2011]. Consciousness is also closely bound up with memory [Lockwood 1998, p84]; it serves as a registration system for information entering the system, which can then be tagged for future retrieval.

Yet, one may still ask: Why could there not be an *unconscious* executive function, and *unconscious* monitoring, just as there can be unconscious perception? A possible answer lies in the mandate of the executive function as decision maker. Information available to the organism is essentially probabilistic and ambiguous. As in a corporation or government, sub-agencies passing such information on to the executive must package it for decisive action. This means that what the conscious self perceives will generally be *un*ambiguous even when wrong. Unconscious models may be "characterized by a high degree of Bayes optimality. But only misrepresenting the probability of a hypothesis as 1.0 and simultaneously flagging it as a fact holding *now*… turns a possibility (or a likelihood) into a reality" [Frith & Metzinger 2013]. To be unambiguous is the very nature of phenomenal experience, which reflects the decisive *action by fiat*[4] that underlies all cognitive judgment. (Consider the Necker cube and other ambiguous figures, which at a given moment appear definitely one way or another.) What renders something certain comes down to an executive decision. Moreover, there should be only one boss; it would not be functional to have multiple seats of consciousness competing in the organism. Sub-agencies might also be restricted to unconscious processing simply because they lack critical resources and complexity required for the executive job.

6. Internal agency

The present theory is that the brain *makes* internal connections on its own initiative by acts of fiat. It thus programs its own VR. Its processes are intentional rather than causal. These constitute *mappings*, in the mathematical sense, and a "language" in an informal sense. This internal language is not merely syntactic, for the organism is motivated to survive. All meaning—including the meaning implicit in qualia—reflects the active relationship of the creature to a world upon which it critically depends. Some of these internal connections an observer will view as representing elements of the external world as the observer perceives it—in other words, as having semantic content. But such reference resides in the organism making the connections, not in the symbols themselves, which have no "pre-existing intentionality" [Kuipers 2007, p86]. They may not mean the same to the outside observer.

For an organism, the meaning to itself of its internal communications is analogous to the meaning that emerges for a human language user—in the act of reading or writing, for example, or of speaking or listening to speech—in which the communicator translates linguistic symbols (written or aural) into mental images, thoughts and feelings, or vice versa. The brain assigns meaning to its own internally symbolic elements, thereby evoking phenomenality in the way that words evoke mental images in particular. The agent translates mere ciphers into conscious experience. Of course, just *how* it does this is the very thing we want to know. Otherwise, we do no more than trade one mystery for another.

In natural language, sounds and symbols carry meaning as words through a constructive process. Phenomenal qualities are comparable to intelligible meanings that emerge from the babble of spoken syllables or the squiggles on a written page. (Similarly, conventional algebraic symbols gain numerical significance by the mathematician's fiat: '*let x stand for such and such…*') Pain, for example, *stands* for something (e.g., tissue damage) as well as compelling a response. We do not normally question the reasons for our own internal connections, to which we do not have conscious access. Yet, it is only from a third-person perspective that they appear arbitrary, merely conventional, or uncompelling, because the observer is then not in the position of being the agent that makes the connection. It then appears mysterious that such "meaningless" connectivity (syntax) carries or implies any meaning.

The creation of sensory experience is *like* the creation of meaning from abstract symbols. Specifically, it is like the evocation of mental imagery in response to language. There is a

[4] That is, by decree, as in "Let there be light!"

resemblance between actual sensory images and their subtler cousins, mental images. Of course, simply pointing to the similarity does not explain what they do and do not have in common, yet their very differences may provide another clue to what is required for sensory phenomenality. One difference stands out: mental images convey only the detail they already embody from prior input. Unlike a sensory image, a memory or visual imagining cannot be searched for more propositional information than it already graphically displays. A retinal image, in contrast, is constantly updated in real time (or nearly), and thus is an ongoing source of new data. The visual field changes as the world changes, but is also continually refreshed through eye saccades. This constant renewal of an external source of sensory input (somehow) provides a vivid experience of reality, as distinguished from paler imagination and memory. But there is another factor besides live input that is required for meaning, and that is its real-time significance to the organism.

7. Embodiment

Embodiment is more than physical presence. It is a relationship of an agent with the world, whether established through the present interaction of the individual or the interaction of the species over generations of natural selection. The meaning to the organism of its internal communications refers ultimately to its evolutionary history. It involves potential consequences, whose significance implies priorities according to which inputs are evaluated, particularly with respect to potential action. To the fly, the configuration of inputs that correspond to the descending swatter *means* "get away" before you are squished or eaten. The fly does not need to have phenomenal experience, think, or make a conscious decision to perform this evasive action, since the reaction can be automatic (and indeed must be, for the sake of speed); yet the action expresses an intention to survive and, accordingly, an interpretation of sensory input. The human consciously intends to kill the fly, or chase it away, but the fly does not need to know the human's intention in order to take the evasive action. A social creature, on the other hand, often *does* need to be able to read the intentions of its conspecifics, who may even issue deliberate warnings as a substitute or prelude for more serious actions. Either way, meaning refers to possible consequence and response, which presupposes an implicit system of values (e.g., survival is good) and the possibility of effective action (e.g., quick take off).

Homeostasis is the ultimate premise of the valuation implicit in all cognition. The creature is motivated to seek external conditions that permit internal conditions to remain within tolerable limits, and to reject external conditions that do not. This is true even for single cells. The cell wall is the primordial boundary between internal and external, and the basis of the distinction between self and other—hence, also, of the distinction between what is good or bad for the organism [de Quincey 2000]. While there is little basis for something to have significance if there is nothing you can do about it or if you don't care, evaluating significance doesn't have to be conscious. It is enough that those creatures survive that are programmed in such a way that they take appropriate action with regard to various stimuli. When it *is* conscious, valuation is experienced as feeling and emotion. The extent to which cognitive science downplays affective states likely reflects a prejudice inherited from the exclusion of the subject, the body, and "secondary" qualities from relevance in physics [Jack et al, p12].

If the organism could do no more than receive physical energy and operate on principles of efficient causation, it would be reasonable to wonder how mere physical stimulation can be imbued with meaning [Dreyfus 2007]. The answer, of course, is that the organism does far more than passively receive energy from a stimulus. It interacts with an environment full of consequence for it and actively assigns significance and meaning according to its own values and criteria, which have been conditioned by natural selection. Its values reflect its evolutionary success—hardly truth as conceived by a human observer.[5] Its nervous system evolved to guide adaptive behavior in a specific niche and for specific purposes, which is not the same as producing veridical perception [Hoffman 2011]. Color perception, for example, provides useful information at a certain scale—such as which fruits are ripe—and less directly information about properties of light (wavelength) and surfaces (reflectance). While the world undeniably has fruit and light and surfaces, the question of whether it has *color* is misguided if not meaningless. For, color is not a property of the world per se, but of the interaction of the world with perceiving organisms, according to their nature and needs as much as the nature of the world perceived. Yet, the same may be said even of so-called primary qualities, such as shape, size, position and time [Edwards & Wilcox 1980], and even the very existence of "objects."

It is unsurprising that we share cognitive biases with other creatures with which we also share most of our genetic makeup. Experiments with monkeys, showing their responses to various colors, shed light on some human preferences [Humphrey 1976]. They also hint at the evolutionary meanings of color qualia themselves as "affect-laden intentional activity" [Rovane 2000]. Following Darwin, we should expect that—whatever else they are—phenomenal qualities are by-products of brain processes conducive to survival [Jackson 1982]. As such, far from being epiphenomenal, they are often intimately linked to appropriate behavior.

8. Phenomenality and behavior

You experience a tickle in the throat and begin to cough. What is the relation between the tickle sensation and the behavior of coughing? Does the experience cause the behavior? Certainly not in the usual sense of efficient causation; for, modern science neither allows nor requires anything outside physical processes or entities as causally effective. Is the tickle sensation epiphenomenal, playing no functional role, while the true causal processes are unconscious or merely physical? No, for the tickle sensation *represents* to you a state of your organism: the existence of an irritant along with the impulse to cough to deal with it.

The tickle represents a state of affairs to an agent (in this case, you) capable of acting upon it independently of any reflex. The sensation serves as an *emblem* of the impulse, to inform you about its occurrence and the condition to which it responds. The registering of the sensation bears information about the stimulus and the appropriate response, and also constitutes a separate function independent of automated behavior. The job of this agent is to monitor the state of the organism, its world, and the activities of various sub-agencies; it has (limited) executive powers

[5] As one commentator pithily overstates it: "Perception is not about truth, it's about having kids" [Hoffman and Prakash 2014]. On the other hand, if *all* cognition is merely adaptive behavior, is science merely a strategy to have kids?

to override the cough or to plan some other action to deal with the irritation that registers as a tickle. The executive function itself cannot be automated like the coughing reflex, because automatisms are fixed algorithms that can deal only with well-defined situations to which they are tailored.

Similarly, there are programmed behaviors associated with other sensations, such as itching, pain, hunger and thirst, and the with valuations involved in such sensations as sweetness, bitterness, or pleasant and repulsive odors. These sensations contain information about the proximal stimulus and the response. In some sense, the associated behavior manifests the *meaning* of the sensation—that is, what the organism should do about the stimulus. However, other sensations—such as provided by the distance senses—may lack obvious behavioral concomitants. What is the "meaning," in the above sense, of the color red, for example? The answer may not be readily forthcoming, yet the question is not without sense. Rather, it suggests the general question: how does the organism establish the meaning to itself of its sensory input?

Direct physical contact is the primordial mobilizing disturbance, which a single-celled organism could deal with by a reflex of withdrawal or engulfing, and which organisms more generally deal with by aversion or attraction. For complex creatures, more complex behavior remains grounded in the biologically fundamental responses of approach and withdrawal. While aversion and attraction need not involve consciousness, consciousness necessarily involves the values behind such behaviors. Affect, with the judgment underwriting it, is therefore central to consciousness.

The sense of touch yields a perception of physical contact, whose experienced quality may be inseparable from the force of the contact itself [Jonas 1966, p148]. At a primordial level, in other words, the affect and the behavior form an undifferentiated whole. Herbert Spencer had a basic intuition that the key to the quality of affect lay in the behavioral response associated with it. He posited a primitive "shock," in which feeling and response are undifferentiated, and which forms the basis for various sensations differentiated according to modality [Spencer 1890]. The relative disinterestedness of the distance senses has its evolutionary roots in the affective values of immediate feeling-responses [Dennett 1991, p181].

A paradigm example is pain, which is at once feeling and response. Pain *means* the associated behaviors of withdrawal, avoidance, and protection of an injured part. But the pain response in humans has two phases, corresponding to distinct neural pathways (c-fibers and a-fibers).[6] The first is a reflex reaction—for example, quick removal of the hand in response to contact with a hot surface—corresponding to the "shock." The slower response of lingering painful sensation reflects an ongoing, *internally generated* stimulus, which acknowledges tissue damage during the process of healing. The ongoing sensation of pain reminds us of the need for the associated protective behavior.[7] This conscious experience of pain implies, first of all, that the initial reflex

---

[6] As Dennett [1978, p200-202] points out, the physiology of pain is more complicated than this, involving separate channels through the "old brain" and the "new brain," and also the possibility of other pathways influencing the experience of pain. This does not affect the argument here, which concerns the grounding of the felt quality in the associated behavior.

[7] Eisemann et al [1984] speculate that insects that do not seem to favor damaged parts do not experience pain as a result of those injuries.

was not sufficient to avoid injury; secondly, that the injured part must be favored during healing; and thirdly, that such stimulus should be avoided in future.

We might think of the ongoing nerve signal as consisting of re-iterations of a first "shock" impulse—in a persisting reverberation or reactivation loop [Humphrey 1992, p204-5]. The integration over time of these impulses constitutes the *quality* experienced, just as the reiterations of a sound wave front constitute a tone. On the basis of such admittedly simplistic analysis, can we hope to understand, in a similar way, qualities that do not seem to involve a reflex response or other associated behavior? The challenge is to understand the grounding of phenomenal qualities in affect, with its associated behavior; but what of qualia that that do not seem to involve any behavioral response? How can we understand specific colors or auditory tones in terms of valuations, judgments and affects?

Now, a reflex response to a proximal stimulus can have direct benefit for a creature. However, on the example of pain, it is not this immediate reaction, but mediated valuation, after the fact, which is salient for phenomenal qualities and upon which *subsequent* action may be based. This mediated valuation continues to carry with it the implied (if not actualized) immediate reaction, if only by demanding attention. (There is neurological evidence that phenomenal experience involves efferent as well as afferent nerves [Ellis 2000, p44][8].) But the quality (color or tone) carries an additional implication as well. Being integrated over time, and accommodated so as *not* to produce a fight or flight response, the sensation is interpreted as distant and potentially innocuous. There is no immediate behavior associated with visible light of a given wavelength, which is too weak a stimulus to directly impact the human organism by the force of contact (of photons, for example). On the other hand, organisms not much larger than molecular scale could not experience the integrations involved in qualia, but would only encounter the "shocks" making them up [Jonas 1966, p29]. For the larger organism, especially, such non-invasive stimulus can be *interpreted* as having a significance requiring *considered* action, which is neither just a physical transfer of energy nor a reflex response, but complex behavior.

All nervous activity consists of electrochemical signals, and all exteroceptors derive from the cell membrane of the organism. Visual receptors differentiate from skin cells in such a way that light is no longer (just) a proximal stimulus with an immediate import, but primarily conveys information from a distance that no longer bears the implications of contact, and which engages a different level of response: monitoring and evaluating from afar. Since the distance senses are by definition largely freed from the need for reaction, visual and aural *qualities* are accordingly dissociated from such response. Yet, auditory tone plainly results from repeated iterations of a wave front impinging on the eardrum, while the single wave front is but a minimal shock. The qualitative experience of tone emerges as the encounter with many such impulses is synthesized into an experience of an overall "texture." Similarly, a stable visual world emerges from repeated saccadic "takes" on the retinal surface, without which vision is not possible [Solms 2014, p183]. And, as we have noted, pain emerges from an ongoing stimulus that was initially but a shock

[8] "Consciousness always involves efferent activity, defined as neural activity generated by the organism itself, for purposes of its own survival and well-being, rather than from passive stimulation by incoming sensory signals." Cf. also [Humphrey 2000, p17-18]: "In order to be able to represent 'what's happening to me', the animal must in fact continue to issue commands such as *would* produce an appropriate response at the right place on the body *if* they were to carry through into bodily behavior."

entailing a reflex response. We shall explore the idea that the brain integrates such micro-events into a gestalt by "filling in" a phenomenal quality.

While unconscious mental processing occurs *before* the conscious experience it underlies, this does not render the conscious experience superfluous (as it would seem to be in a causal framework). Rather, the conscious experience indicates recognition by the executive function, after the fact, of that particular unconscious processing. The conscious experience plays a different role, with a different associated behavior, than the pre-programmed response. It also plays a different role than its own underlying causal mechanisms, just as a computer icon plays a different role than the computational processes that result in the pixels that make it up [Hoffman 2008].

A similar divide between quick and slow pathways, which we see in the case of pain, seems to exist in the visual system too, where an initial fast wave of visual processing happens outside consciousness, and is made available to subsystems for immediate responses. This is followed by a slower phase of "recurrent processing" that involves integration of various brain areas leading to conscious experience [Revonsuo 2010]. Here too phenomenal qualities emerge from the reiteration of an original signal and serve a different purpose from it.

9. The language of the senses

Whatever the color red signifies to a contemporary organism may not be the same as what it did a million years ago, just as many words in Shakespeare's plays do not have the same connotations to a modern ear. Shakespearian scholars attempt to make up this deficit, and evolutionary psychologists may attempt to unravel the changing meanings of the language of the senses. The intentionality involved in human color perception, for example, is deeply buried in the evolution of the primate visual sense. We have little current need to be engaged by possible earlier behavioral correlates of color qualia, firstly because those associations are passé; secondly, because the nature of a distance sense is to divorce itself from immediate behavioral implications altogether. We gain the objectivity achieved at the cost of access to the underlying subjectivity. The problem presented by the apparent objectivity of the visual sense is to understand how its qualia nevertheless derive indirectly from affect—in other words, to speculate on the behavioral meanings of colors. Evolutionary advantage explains color *discrimination*—why things appear to be differently colored—but not why a particular wavelength of light appears just so and not otherwise. Yet, if color categories (hue) reflect needs of the organism as much as properties of light or of reflective surfaces, then they must have an evolutionary significance.

The diets of Old-World primates consist significantly of fruits that are yellow, orange or red [Tsou 2013].[9] It would make sense for these food items to stand out in primate perception from a

[9] Trichromancy enables primates and their human relatives to distinguish effectively in a range of wavelengths corresponding to phenomenal red through phenomenal blue. Though idealized, the "primary" colors correspond roughly with particular sensitivities of receptors. As Newton had noted, red and blue are subjectively connected at the farthest ends of the spectrum, so that violet seems to mix blue with red. This forms a cycle (hence the color wheel) and falls just short of covering an "octave" of

background of foliage (subjectively we tend to think of red as contrasting maximally with green). Yellow is closer to green in wavelength (hence with less contrast than red), and can also indicate a less ripe fruit. In the forest context, at least, the color red serves to alert the creature to something singular—whether a ripe fruit or a poisonous creature that has adopted the color code [Dennett 1991, p385]. Other associations are possible, such as blood, or colorings related to sexuality, or the red of dusk when some predators hunt. Perhaps all qualia involve a network of such associations, for which the quality itself summarily stands [Loorits 2014].

This still does not tell us why the chlorophyll of foliage does not appear red and the ripe fruit green, which would maintain the same relative contrast for discrimination. What is it about the qualitative "feel" of greenness that commends it to represent foliage in the vocabulary of the senses; and what about redness commends it to represent things that must stand out against that particular background? The question may be likened to asking why a particular meaning is denoted in a given language by a particular word, written and pronounced its specific way, rather than by some other symbol. For the native speaker, the association of the word with what it represents seems natural and unquestionable, though of course it is actually a social convention and a product of historical accident. The internal language of the organism may be no less arbitrary and accidental in its choice of symbols. *Some* symbol must be chosen, and will inevitably come to seem imbued with the meaning it has been made to convey. So, perhaps it is backward to ask why grass appears green; rather, greenness is what it is by virtue of the association with grass. Greenness is the way we visually experience the totality of associations related primarily to chlorophyll.

Like perception, language is motivated. Originally, language may have served to alert or warn. Today one may discourse dispassionately about almost any topic, but the first human vocal expressions were probably not so different from the excited alarm calls of primates and other animals. What makes fully grammatical language such an invaluable tool is precisely its flexibility, combining words and ideas in inventive new permutations, removed from their original (or any specific) context. Vision especially is detached in a similar way. For the human being, capable of abstraction and bent on transcending the natural condition, the informing significance of most phenomenal qualities is no longer compelled by their original urgent associations, from which they may nevertheless have descended in the way that grammatical language evolved from the urgency of animal calls.

Certainly, some colors continue to have affective values for human beings, as for our primate relatives. They can serve to capture attention, to convey information, and to bear an emotional charge. Monkeys tested for color preferences show a preference for blue and green and an aversion to red, which agitates them. However, as Humphrey [1976] explains, the laboratory context must be taken into account, where the natural context for response cannot be established. In any case, red can signify a variety of quite different things. It may be precisely its ambiguity

---

wavelengths. That is, one end of the visible spectrum is just less than twice the wavelength of the other. If human vision extended over more than one such octave (into infrared and ultraviolet), one might be permitted to speculate that the subjective experience would resemble that of sound octaves, with pitches at multiples of each other seeming qualitatively similar though "higher or lower." This raises the question, however, of what would distinguish the hue in one octave of light from the "same" hue in another octave. On the other hand, some tone-deaf people cannot tell which of two auditory tones is higher in pitch.

as a signal that imbues it with its characteristic power, and creates anxiety in a situation that does not lend a natural interpretation.

The *sensation* of redness (unlike the word) is not merely a linguistic convention subject to social change, but a convention of neuro-logical organization, with the force of long genetic precedent. Indeed, the human cognitive system adapts to distorting lenses or colored filters in such a way that experience of verdant foliage, for example, is eventually restored to its normally experienced greenness [Neitz, et al 2002]. While the words of a natural language have relatively transient reference, the meanings of qualia are more stable, being backed by evolutionary history. The sensation of redness is what it is, and different from the sensation of greenness, precisely because of the real-world things it refers to in our evolutionary history, from which it cannot be arbitrarily dissociated. This is why there can be no "inverted spectrum."[10]

10. Filling in: perception *of* is perception *that*

The qualities immanent in phenomenality are far from "epiphenomenal." In looking at a source of green light, for instance, one should not imagine that the brain assesses a frequency and then "colors in" that information with the superfluous quality of greenness, while the information is the real causal factor. Rather, the experience of greenness comes with that information built in; the experience *itself* stands for the brain's qualitative estimation of frequency (among other things). The notion of "information" derived from it is an afterthought.

Similarly, the perceived quality of a particular musical tone stands for an estimate of sound frequency. Sensory qualities are thus not something gratuitously added to the information they represent, nor are they caused by it. Rather, they are a *version* of that information, synoptically presenting it to the executive function. Qualia, in other words, are how the embodied subject first-personally and "qualitatively" monitors sensory information that an outside observer might also detect by means of laboratory equipment and describe in terms that are third-personal, physical, propositional, and quantitative. If the specific quality of greenness, for example, seems to convey privileged information beyond that involved in the public analysis of light, this is because (along with information about the world) it also bears information about the organism's internal communication—which involves its relationship to the world, its priorities and evolutionary history. It is this internal communication that is "private," which becomes public when it is translated into propositions (facts) asserted among agents.

Propositional knowledge in general, and as conveyed in science, concerns perceived *differences*—that is, differences among qualia, which are construed to be differences in the world. Facts are based on differences within phenomenal experience, but are not the experience itself. However, qualia themselves contain implicit information about differences in the world, which may be rendered explicitly propositional. Qualia seem ineffable because they *already*

[10] The 'inverted spectrum' is the "apparent possibility of two people sharing their color vocabulary and discriminations, although the colors one sees… are systematically different from the colors the other person sees" [Wikpedia: inverted spectrum].

recognize and refer to differences that can potentially *also* be expressed propositionally. And, just as propositional knowledge involves an intentional act of the observer, so qualia involve intentional acts by an internal agent.

To be perceived as a *tone*, there must be a regular repetition of wave fronts hitting the eardrum. A single wave front impinging on it will not be experienced as a tone, if it registers as an experience at all. Similarly, the eye requires a certain number of photons (or waves) to register color or anything at all. Moreover, it must constantly move (with saccades) in order to maintain a stable ongoing visual field. Without this continual shifting of gaze, so that information is constantly updated on the retina, the visual field dissolves. The succession of saccades provides an effect like the succession of wave fronts on the ear. The eye/brain smooths over the jerkiness of the saccades like it does the separate frames of a motion picture, interpolating between them. It "fills in" the continuity and stability of the visual field like it fills in the visual blind spot and fills in the quality of a tone, which is integrated over successive wave fronts.

However, this "filling in" is not a *perception* of continuity but an *assertion* of it. As Dennett [1992] notes, the brain perceives continuity by disregarding the actual discontinuity. If the brain can create such an appearance by simply asserting that it is so, then why not all appearances? In short, all perception *of* is actually perception *that.* Only in anomalous circumstances do we even notice this. These circumstances include laboratory studies of perceptual completion effects, habituation, perceptual adaptations of various sorts, and phenomena of spatial projection—all of which involve acts of fiat. A feature of such effects is that the subject's phenomenal experience typically seems to go beyond the facts or events noted by an observer. That is, the subject's brain "fills in" what—according to a third-person observer—is "not really there." Let us speculate that qualities in general may be understood as such completion, adaptation, or projection effects, while noting that what is deemed "really there" reflects a fundamental bias in the third-person point of view.

"Filling in" describes a fundamental aspect of all perception, which may alternatively be described as an act of positing, mapping, or fiat. It demonstrates the constructive capability of the brain, a basic process of interpolation [Crick and Koch 1992/2002, p15]. Laboratory experiments demonstrate a taxonomy of visual completion effects of various sorts [Pessoa et al 1998]. Some demonstrate various forms of spatial and temporal projection. Still others (among which the famous experiments of Stratton, oft repeated in variations) show the adaptability of the nervous system to restore perception that corresponds to functional behavior. These can involve quite bizarre experiences, such as the strange feeling of phantom limb that can be induced in the "rubber hand" experiment [Botvinick & Cohen 1998; Metzinger 2010]. Apparent motion effects are well known in modern culture through motion pictures and illuminated signs.

While effects are considered illusory when they involve jumping to a false perceptual conclusion, the general lesson is that *all* perception is naturally a matter of jumping to conclusions, even when warranted. In the case of the blind spot, the experience of continuity of the visual field across it is effectively the brain's way of representing to itself its (true) belief that, despite the absence of enervation, the external visual world is continuous. The broader contention is that qualia *in general* are a result of the same sort of process involved in the blind spot and other completion effects, all of which reflect a positive act of assertion, not simply ignoring an absence. One might then ask—in the general case—what the brain fills in *between.*

In the case of the blind spot, it fills in a phenomenal experience between the *enervated* retinal areas (which normally give rise to experience) on either side of the un-enervated blind spot. If all perception of is perception that, then it must be the case that these adjacent enervated areas *too* are filled in, but on a finer scale (between receptors, for example), and temporally as well as spatially. That is, the brain *generally* asserts continuity across discrete structures or events when their discontinuity is irrelevant, just as it asserts continuity between frames of a motion picture. Phenomenal qualities in general embody such assertions, just as the quality of being fifty years old is "filled in" by fiat between one's fiftieth and fifty-first birthdays.[11]

No matter how complex the neural processes giving rise to them, qualia are essentially simple and integral gestalts [Crane 2000, p188]. This is so first of all because there is no conscious access to the underlying complex processes; but equally, it is because qualia are by nature interpolations or syntheses to summate information. Though qualia are products of neural events and reflect underlying structure, they gloss over and integrate such events and structure, without revealing their etiology. Color or sound experience synthesizes information about structure in the environment to yield a useful appraisal. But the individual wave fronts of light or sound are glossed over and integrated by the nervous system into a seemingly structureless emerging quality.

One may speculate that qualia are built up essentially from primitive *responses* at a lower level, in the way that a digital image is built from pixels [MacLennan 2005, sec3B]. Each "pixel" may represent a simple judgment or assertion (e.g., 1 or 0), but the image that emerges with scale is an integration of that digital information, facilitating judgment and response on another level. The scale of individual impulses (pixels) allows both their miniscule energy and their individual import to be absorbed into a larger synthesis from which the agent concerned can be relatively detached, even though the "raw material is action: impacts, hustlings, clashes on a molecular level" [Jonas 1966, p29].

Mental images are sparsely detailed, compared to perceptual images. But even the latter are only relatively detailed. The impression of dense or complete detail in sensory experience is illusory, since it is actually no more than the assertion *that* there is unlimited detail. This assertion is backed up by the fact that the senses can access additional information about the external world upon demand, whereas mental images cannot. Suppose you close your eyes and conjure a mental image of this printed page. Unless you have perfect eidetic memory, in this image you may recognize *that* there are typed letters in lines, but not necessarily enough detail to read the sentences. It may well be that distinct words or letters are not made out, nor even the exact typeface. Now open your eyes and look at the real printed page. You read the actual words and sentences by directing your (quite narrow) foveal vision at a line of print, one word or phrase at a time. In your peripheral vision, the words are not so clear, though you may anticipate them because of context. Reading even an individual word, however, involves *recognizing that* these

[11] Similarly, the "time" on a digital clock is filled in between numbers displayed, and the "quantity" of a measurement is filled in between defined gradations.

letters constitute a meaningful gestalt, which is a matter of deciding that a particular configuration of marks represents a word in the language you know.

Let's say that now that you notice a painting on the wall in your peripheral vision. Though familiar, you may recognize it without discerning much detail. Peripherally or from a distance, you perceive only *that* it is a particular painting, with its variety of colors and its certain size, and little more until the eyes are directed to center the fovea on some area of it and attention is directed to search out more details within that limited visual area. Noticing a "detail," however, is much the same thing as noticing that it is a painting, but on a finer scale: it is a matter of asserting a fact. "Seeing" a detail is seeing that something is so, just as you saw that it was a painting.

The possibility to identify indefinitely many such details (by moving the fovea about or by paying closer attention) distinguishes sensory perception from remembered imagery, imagination, or dreams. Yet, the impression is illusory that perception reveals reality in all or even many details at a given moment [Crick and Koch 1998, p99]. For example, a glance takes in with clarity only a small portion of a repetitive pattern, such as on wallpaper. Yet, one has the false impression that one uniformly "sees" at once all the individual forms that are repeated [Dennett 1991, p354-5].[12] Let us take this iterative process a step further: it only seems (falsely) as though we see in detail the whole of *even the single isolated form*. In fact, we "see" what we have decided must be there!

We may note further that perception *of* is not even a coherent notion. For it supposes a presentation of what *exists*, as though one passively surveys a panorama given in its entirety. This is a belief made plausible, perhaps, by the objectifications of physics. But perception entails cognizing sensory input as we go along. It is an assertive act of the organism, concerning what is sensed. It may not be a *definitive* assertion (as illustrated by the Necker cube and other ambivalent figures about which the brain cannot make up its mind); but, in the moment, it is made with the tautological finality of all declarations by fiat. Perception is definite even when it is wrong. And it is only the organism's act of fiat that makes it so, without which there could be no presentation in consciousness at all [Frith & Metzinger 2013]. This fundamental action of assertion—which we have also called intention, mapping, or fiat—is the essence of mind and the basis of phenomenality.

11. A note on qualia[13]

[12] Dennett's example is wallpaper by Andy Warhol, with a repeated stylized image of Marilyn Monroe—an artifact with only limited detail. The discrepancy is even more extreme with perception of a natural scene, which contains unlimited detail to be fathomed.

[13] According to Wikipedia, "qualia are individual instances of subjective, conscious experience." The term 'qualia' (singular 'quale') was first used in its modern sense by C. S. Peirce in 1866. But explicit focus on the role of qualia in the mind-body problem probably began in 1958 with Herbert Feigl's *The Mental and the Physical* [Crane 2000, p177-81].

Qualia seem to have aroused more passionate discussion than any other topic in philosophy of mind [Armstrong 1999]. This is scarcely surprising, since explaining qualia effectively *is* the hard problem [Crane 2000, p171]. Since perception is inherently propositional (perception *of* is perception *that*), it has been argued [Dennett 1991] that qualia don't exist. That is tantamount to arguing that only facts or propositional information exist. But, as argued here, qualia are themselves representations of what is otherwise framed as propositional information. This misunderstanding is partly what makes qualia problematic, like the disgraced notion of 'sense-data'. Qualia are not normally *objects* of perception, but are the perceiving itself. One does not naturally see such things as "color patches" nor does one have "raw feels." Yet it is possible, with special effort, to "deconstruct" features of phenomenal experience. For example, one can attend to color patches instead of colored objects. Painters routinely do this, and paint has culturally accustomed us to the idea of "pure color." However, such artifacts are not the normal *input* to cognition but the *result* of special cognitive acts. In experimental situations, qualia are artifacts of the experimental set-up. For example, one does not usually encounter monochromatic colors in nature, in isolation from objects and surfaces.

Qualia are not deducible from physical facts alone, because physical facts are designed to circumvent phenomenal experience in the first place, and because consciousness is not a product of the object alone, but of the subject and object together. This is why Jackson's [1982] artificially sequestered scientist "Mary" can know all there is to know about neural processes without having experienced certain qualia.

12. Summary conclusion

We have argued that a solution to the hard problem requires taking the point of view of the organism and its subagents, and that doing so need not pose any insuperable problem. However, what is necessary may not be sufficient.

We attempt to understand—from the inside, as it were—how the organism constructs its own first-person point of view and phenomenality. The organism does this through acts of fiat that an onlooker can translate as propositional assertions. The brain creates phenomenality like it creates meaning in language, by assigning meaning in the context of an embodied evolutionary history. Phenomenality is an internal representation *for* an executive agent tasked with monitoring the state of the organism and its environment, planning future action, and coordinating various sub-agencies.

Does this sketch provide the basis for a solution to the elusive hard problem? One may still feel that there remains a sense in which even this approach does not tell us why there is phenomenality in the first place. Yet, as Chalmers [1995] notes, "this is the same for any fundamental theory. Nothing in physics tells us why there is matter in the first place, but we do not count this against theories of matter." Despite the advances of science, it is possible that we may never have an answer to the mystery of why anything exists at all. Similarly, we may never have an answer to the mystery of how there can be such a thing as consciousness of it. Indeed, it is only because of the reflexivity of consciousness that the two questions can be separately posed.

Received June 11, 2024; Accepted July 8, 2024

References

Armstrong, D. M. (1999) *The Mind-Body Problem: an opinionated introduction* Westview (Perseus), p121

Boden, Margaret A. (2006) Mind as Machine: *A History of Cognitive Science*, Volume 1&2 Oxford University Press,

Botvinick, M. and Cohen, J. (1998) "Rubber Hands 'Feel' Touch that Eyes See" *Nature* 391, 756; cited in Metzinger (2010)

Chalmers, David J. (1995) "Facing Up to the Problem of Consciousness" http://consc.net/papers/facing.html, p16 (postprint from *Journal of Consciousness Studies* 2(3):200-19

Crane, Tim (2000) "The Origins of Qualia" in *History of the Mind-Body Problem*, ed by Tim Crane and Sarah Patterson, Routledge

Crick, Francis and Koch, Christof (1992/2002) "The Problem of Consciousness" *Scientific American*, updated from September 1992

Crick, Francis and Koch, Christof (1998) "Consciousness and Neuroscience" *Cerebral Cortex* Mar 1998;8:97–107

Dennett, Daniel C. (1978/1981/1996) *Brainstorms* Bradford (MIT Press), p122

Dennett, Daniel C. (1987/1998) *The Intentional Stance* MIT Press, p302-314

Dennett, Daniel C. (1991) *Consciousness Explained* Little Brown & Co

Dennett, Daniel C. (1992) "'Filling in' versus Finding out: A Ubiquitous Confusion in Cognitive Science." In H. Pick, Paulus van den Broek and David C. Knill, eds., *Cognition: Conceptual and Methodological Issues*, American Psychological Association, pp. 33-49

Dreyfus, Hubert L. (2007) "Why Heideggerian AI Failed and how Fixing it would Require making it more Heideggerian", Sec VIII, preprint for *Philosophical Psychology* 171(18):1137-1160

Edwards, David A. and Wilcox, Stephen (1980) "Unity, Disunity and Pluralism in Science" arxiv.org/pdf/1110.6545, sec2 (discussing Berkeley and Kant)

Eisemann, C. H.; Jorgensen, W. K.; Merritt, D. J.; Rice, M. J.; Cribb, B. W.; Webb, P. D.; and Zalucki, M. P. (1984) "Do Insects Feel Pain? – A biological view" *Experientia* 40

Ellis, Ralph (2000) "Efferent Brain Processes and the Enactive Approach to Consciousness" *Journal of*

Consciousness Studies 7, no. 4

Frith, Chris D. and Metzinger, Thomas K. (2013) "What's the Use of Consciousness?" *(https://www.blogs.uni-mainz.de/fb05philosophie/files/2013/04/Frith_Metzinger_Regret_2016_penultimate.pdf)*

Hirstein, W. (2011) The Contribution of Prefrontal Executive Processes to Creating a Sense of Self. In: *Brain, Mind and Consciousness: An International, Interdisciplinary Perspective* (A.R. Singh and S.A. Singh eds.), *MSM,* **9**(1), p153

Hoffman, Donald D. (2008) "Conscious Realism and the Mind-Body Problem" *Mind & Matter* Vol. 6(1), pp. 87–121

Hoffman, Donald D. (2009) "The Interface Theory of Perception" (A chapter for the book "Object Categorization: Computer and Human Vision Perspectives," edited by Sven Dickinson, Michael Tarr, Ales Leonardis and Bernt Schiele. Cambridge University Press)

Hoffman, Donald. D. (2011) "The Construction of Visual Reality" (Chapter preprint for the book "Hallucination: Theory and Practice", edited by J.D. Blom and I. Sommer. Springer Verlag)

Hoffman, Donald D. and Prakash, Chetan (2014) "Objects of consciousness" *Frontiers in Psychology* Volume 5 | Article 577, p3

Humphrey, Nicholas (1976) "The Colour Currency of Nature" First published in *Colour for Architecture*, ed. Tom Porter and Byron Mikellides, pp. 95-98, Studio-Vista, London

Humphrey, Nicholas (1992) *A History of the Mind: evolution and the birth of consciousness* Copernicus (Springer-Verlag)

Humphrey, Nicholas (2000) *How to Solve the Mind-Body Problem* Imprint Academic

Jack, Anthony I.; Robbins, Philip; Roepstorff, Andreas (2007) "The genuine problem of consciousness" Philpapers.org

Jackson, Frank (1982) "Epiphenomenal Qualia" Reprinted from *The Philosophical Quarterly*, 32, No. 127, pp. 127-136

Jonas, Hans *The Phenomenon of Life: toward a philosophical biology* NW University Press, 1966/2001

Kandel, Eric (2012) *The Age of Insight: the quest to understand the unconscious in art, mind, and brain* Random House, p464-5

Koubeissi, Mohamad Z.; Bartolomei, Fabrice; Beltagy, Abdelrahman; Picard, Fabienne (2014). "Electrical stimulation of a small brain area reversibly disrupts consciousness". *Epilepsy & Behavior.* ***37**: 32–35*

Kuipers, Benjamin (2007) "Sneaking Up On the Hard Problem of Consciousness" https://www.aaai.org/Papers/Symposia/Fall/2007/FS-07-01/FS07-01-015.pdf

Lamme, V. (2000) "Neural Mechanisms of Visual Awareness: a linking proposition" *Brain and Mind* 1

Lamme, Victor (2015) “The Crack of Dawn: Perceptual Functions and Neural Mechanisms that Mark the Transition from Unconscious Processing to Conscious Vision”. In T. Metzinger & J.M. Windt (Eds). Open MIND:22(T)

Lehar, Steven (2003) “Gestalt isomorphism and the primacy of subjective conscious experience: A Gestalt Bubble model” *Behavior and Brain Sciences* 26

Lockwood, Michael (1998) “The Enigma of Sentience” in *Toward a Science of Consciousness* ed by Stuart R. Hameroff et al, MIT Press

Loorits, Kristjan (2014) “Structural qualia: a solution to the hard problem of consciousness” Front Psychol. 2014; 5: 237 (citing work of Crick and Koch)

MacLennan, B. J. (2005) “Consciousness in Robots: The Hard Problem and Some Less Hard Problems”, preprint of ROMAN 2005. IEEE International Workshop on Robot and Human Interactive Communication

Metzinger, Thomas (2009) *The Ego Tunnel: the science of mind and the myth of the self* Basic Books

Metzinger, Thomas (2010) “The Self-Model Theory of Subjectivity: A Brief Summary with Examples” Humana.Mente – Issue 14

Nagel, Thomas (1974) “What Is It Like to Be a Bat?” *The Philosophical Review*, Vol. 83, No. 4., pp. 435-450.

Neitz, J., Carroll, J., Yamauchi, Y., Neitz, M. & Williams, D. R. (2002) “Color Perception Is Mediated by a Plastic Neural Mechanism that Is Adjustable in Adults” *Neuron*, 35, pp783–792

Oatley, Keith (1978) *Perceptions and Representations* Methuen

Pessoa, Luiz; Thomson, Evan; and Noë, Alva (1998) “Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception” *Behavioral and Brain Sciences* 21

de Quincey, Christian (2000) “Conceiving the Inconceivable?” discussion of N. Humphrey’s target article in *Journal of Consciousness Studies* 7, no. 4, p72

Revonsuo, Antti (2006) *Inner Presence: Consciousness as a biological phenomenon* MIT Press

Revonsuo, Antti (2010) *Consciousness: the science of subjectivity* Psychology Press, p215, [citing work of Victor Lamme (2000)]

Rovane, Carol (2000) “Not Mind-Body but Mind-Mind”, in discussion of N. Humphrey’s target article in *Journal of Consciousness Studies* 7, no. 4, p89

Sloman, Aaron (1993) “The Mind As a Control System” offprint [2014] from Christopher Hookway & Donald M. Peterson (eds.), Royal Institute of Philosophy Supplement. Cambridge University Press 1993, p69-110

Solms, Mark (2014) “A neuropsychoanalytical approach to the hard problem of consciousness” offprint from *Journal of Integrative Neuroscience*, Vol. 13, No. 2

Spencer, H. (1890) *Principles of Psychology* 3rd Edition, vol 1, London: Williams & Norgate

Tsou, Jonathan Y. (2013) "Origins of the Qualitative Aspects of Consciousness: Evolutionary Answers to Chalmers' Hard Problem" preprint for chapter in Liz Swan (ed.), *Origins of Mind* Springer 259-269