

Nested ontology and causal options:

A paradigm for consciousness

Bjorn Merker

Gamla Kyrkv. 44, Segeltorp, Sweden

Copyright © Bjorn Merker, 2010

Address correspondence to:

Dr. Bjorn Merker

Gamla Kyrkv. 44

SE-14171 Segeltorp

Sweden

E-mail: gyr694c@tninet.se

Abstract

A brain charged with guiding its body through a complex and lively world from a position of solitary confinement inside its opaque skull faces a set of functional problems whose solution may account for the existence and nature of consciousness. An analysis of the more general and basic of these problems, sensory as well as motor, suggests the utility of implementing a high-level mutual interface between sensory target selection, motor action selection, and motivational ranking of needs at a late stage in the run-up to the brain's decision about the very next action to take. The three selection processes are subject to a number of mutual dependencies such that a regimen of constraint satisfaction among them would yield gains in behavioral efficiency. The logistics of implementing such a regimen can be simplified by casting the interface in a particular nested, analog format. It would host a running synthetic summary of the rest of the brain's interpretive labors, reflecting best estimates of the veridical current state of world, body, and needs for purposes of real-time decision making. Detailed scrutiny of the design requirements for such a mechanism discloses that it would be functionally partitioned in a way that defines a conscious mode of operation. Moreover, the design of the mechanism mandates a specific departure from veridicality at a point that makes its functional format match the assumptions of naive realism. Consciousness itself thus introduces a significant, though not insuperable, psychological obstacle to the development of a veridical account of its nature.

Keywords: Architecture of consciousness, constraint satisfaction, control theory, egocenter, naive realism

Nested ontology and causal options: A paradigm for consciousness

1 INTRODUCTION

The fact that we find ourselves surrounded by a world of complex objects and events directly accessible to our inspection and manipulation might seem too trivial a commonplace to merit scientific attention. Yet here as elsewhere familiarity may mask underlying complexities, as we discover when we try to unravel the appearances of our experience in causal terms. Consider, for example, that the visual impression of our surroundings originates in the pattern of light and darkness projected from the world through our pupils onto the lightsensitive tissue at the back of our eyes. On the retina a given sudden displacement of that projected image behaves the same whether caused by a voluntary eye movement, a passive displacement of the eye by external impact, or an actual displacement of the world before the still eye. Yet only in the latter two cases do we experience any movement of the world at all. In the first case the world remains perfectly still and stable before us, though the retinal image has undergone the selfsame sudden displacement in all three cases. But that means that somewhere between our retina and our experience, the facts of self-motion have been brought to bear on retinal information to determine our experience. That in turn implies that the reality we experience is more of a derivative and synthetic product than we ordinarily take it to be.

That implication only grows as we pursue the fate of retinal patterns into the brain. There visual neuroscience discloses not only a diverse set of subcortical destinations of the optic tract, but an elaborate cortical system for visual analysis and synthesis. Its hierarchical multi-map organization for scene analysis and visuospatial orientation features functional specialization by area (Lennie, 1998) and functional integration through a pattern of topographically organized bidirectional connectivity that links each area directly with typically a dozen others or more (Felleman & Van Essen, 1991).

From the point of view of our experience, a remarkable fact about this elaborate system is the extent to which we are oblivious to much of its busy traffic. As we go about our affairs in a complex environment we never face half-analyzed objects at partial way stations of the system, and we never have to wait even for a moment while a scene segmentation is being finished for us. We have no awareness of the multiple partial operations that allow us to see the world we inhabit. Instead it is only the final, finished products of those operations that make their way into our consciousness. They do so as fully formed objects and events, in the aggregate making up the interpreted and typically well understood visual scene we happen to find ourselves in.

So compelling is this "finishdness" of the visual world we inhabit that we tend to take it to be the physical universe itself, though everything we know about the processes of vision tells us that what we confront in visual experience cannot be the physical world itself but rather must be an image of it. That image conveys veridical information about the world and presents some of the world's properties to us in striking and vivid forms, but only to the extent that those properties are reflected in that tiny sliver of the electromagnetic spectrum to which our photoreceptors are sensitive, and which we therefore call visible light. The fact that this tiny part of the spectrum serves as medium for the *entirety* of our visual world suggests that somehow that world lies on "our side" of the photoreceptors, and that what we experience directly is an image of the world built up as an irremediably indirect and synthetic internal occurrence in the brain. But where then is that brain itself, inside of which our experienced world is supposedly synthesized on this account of things? And indeed, does not the location of our retinae appear to lie inside this world we experience rather than beyond it?

These legitimate questions bring us face to face with the problem of consciousness in its full scope. That problem, they remind us, is not confined to accounting for things like the stirrings of thoughts in our heads and feelings in our breasts – what we might call our "inner life." It

extends, rather, to everything that fills our experience, among which this rich and lively world that surrounds us is not the least. In fact, as we shall see, there are attributes of this experienced world that provide essential clues to the nature of consciousness itself. It may even be that short of coming to grips in these terms with the problem of the world that surrounds us, the key to the facts of our inner life will elude us.

2 STRATAGEMS OF SOLITARY CONFINEMENT

Our visual system not only provides us with robust conscious percepts such as the sight of a chair or of storm clouds gathering on the horizon, it presents them to us in a magnificently organized macro-structure, the format of our ordinary conscious waking reality. Our mobile body is its ever-central object, surrounded by the stable world on all sides, on which we look out from a point inside our body through a cyclopean aperture in the upper part of the face region of our head. This truly remarkable nested geometry in three dimensions around a central perspective point is a fundamental organizing principle of adult human consciousness (Hering, 1879; Mach, 1897; Roelofs, 1959; Merker, 2007a, pp. 72-73). It requires explanation despite – or rather exactly because of – its ever-present familiarity as the framework or format of our experience. As such it provides unique opportunities for analysis, because it offers *specificities of structure* whose arrangement simply cries out for functional interpretation.

The key to that interpretation, I suggest, is the predicament of a brain charged with guiding a physical body through a complex physical world from a position of solitary confinement inside an opaque and sturdy skull. There it has no direct access to either body or world. From inside its bony prison, the brain can inform itself about surrounding objects and events only indirectly, by remote sensing of the surface distribution of the world's impact on a variety of receptor arrays built into the body wall. Being fixed to the body, those sensors move with it, occasioning the already mentioned contamination of sensory information about the world by

the sensory consequences of self-motion. But even under stable, stationary circumstances, primary sensory information is not uniquely determined by its causes in the world. Thus an ellipsoid retinal image may reflect an oval object, or a circular one tilted with respect to our line of sight, to give but one example of many such problems occasioned by the brain's indirect access to the world (Helmholtz, 1867; Witkin, 1981, pp. 29-36).

Nor is the brain's control of its body any more direct than is its access to circumstances in the world on which that control must be based. Between the brain and the body movements it must control lie sets of linked skeletal joints each supplied by many muscles to be squirted with acetylcholine through motor nerves in a sequence and in amounts requisite to match the resultant movement to targets in the world. In such multi-joint systems, degrees of freedom accumulate across linked joints (not to speak of muscles). A given desired targetting movement accordingly does not have a unique specification either in terms of the joint kinematics or the muscle dynamics to be employed in its execution (Bernstein, 1967; Gallistel, 1999, pp. 6-7).

On both the sensory and motor sides of its operations the brain is faced, in other words, with under-specified (or ill-posed, under-constrained) problems in the sensing and control tasks it must discharge. We know, nevertheless, that somehow it has managed to finesse these so called inverse problems, because we manifestly are able to function and get about competently even in quite complex circumstances. The brain has in fact mastered its problems in this regard to such an extent that it allows us to remain oblivious to the difficulties, to proceed with our daily activities in a habitual stance of naive realism. We look, and appear to confront the objects of the world directly. We decide to reach for one or another of them, and our arm moves as if by magic to land our hand and fingers on the target. Much must be happening "behind the scenes" of our awareness to make such apparent magic possible.

Reliable performance in inherently underconstrained circumstances is only possible on the

basis of the kind of inferential, estimation, and optimization approaches to which engineers resort when faced with similar problems in building remote controllers for power grids or plant automation (McFarland, 1977). In such approaches a prominent role is played by so called forward and inverse models of the problem domain to be sensed or controlled, and they have been proposed to play a number of roles in the brain as well (Kawato, 1999; Kawato, Hayakawa, & Inui, 1993; Wolpert, Ghahramani, & Jordan, 1995). In effect they move the problem domain "inside the brain" (note: this does not mean "into our 'inner life'") in the form of a neural model, in keeping with a theorem from the heyday of cybernetics stating that an optimal controller must model the system it controls (Conant & Ashby, 1967).

There is every reason to believe that a number of these neural models contribute crucially to shaping the contents of our experience through functions such as cancellation of sensory consequences of self-produced movement to give us the stability of our experienced world despite movements of our eyes, head and body (Dean, Porrill, & Stone, 2004). At the same time there is no a priori reason to think that these neural models themselves are organized in a conscious manner. The proposal that they are has been made on the basis of their essential contribution to sophisticated neural processes, rather than by reference to some principle that would make these and not other sophisticated neural processes conscious (Kawato, 1997). There is, however, a potential functional niche for casting one such neural modelling device in a format yielding a conscious mode of operation, partial versions of which have been sketched in previous publications of mine (Merker, 2005, 2007a).

3 A DUAL-PURPOSE SIMULATOR

Two different functional constructs are joined in the present proposal regarding the role and organization of consciousness. One introduces a comprehensive central solution to the captive brain's sensors-in-motion problems in the form of a dedicated "orienting domain." The other

achieves savings in behavioral resource expenditure by exploiting dependencies among the brain's three principal task clusters (viz. target selection, action selection, and motivational ranking) through constraint satisfaction among them in a "decision domain." Each of the functional problem constellations addressed by these two domains may be amenable to a variety of piece-meal neural solutions, in which case neither of them requires a conscious mode of organization. The central claim of this paper, and the key defining concept of the paradigm for consciousness it is proposing, is that *when comprehensive analog solutions to both are combined in a single neural mechanism an arrangement results which defines a conscious mode of operation.* The two functional domains have not been treated explicitly as such in the technical literature, so I begin by giving a thumb-nail sketch of each before outlining the prospects and consequences of combining the two. Both concern movement and the immediate run-up to the brain's decision about the very next overt action to take, but they relate to it in very different ways.

3.1 The orienting domain: A nested remedy for the liabilities of mobility

The already mentioned contamination of information about the world by the sensory consequences of self-motion is not the brain's only problem caused by bodily mobility. The body moves not only with respect to the world, but relative to itself as well. The brain's sensor arrays come in several modalities differently distributed on the body and move with its movements. Its twisting and flexing cause sensors to move with respect to one another, bringing the spatial information they convey out of mutual alignment. In the typical case of gaze shifts employing combined eye and head movements, vision and audition are displaced with respect to the somatosensory representation of the rest of the body. To this is added misalignment between vision and audition when the eyes deviate in their orbits (Sparks, 1999). The brain's sensors-in-motion problem combines, in other words, aspects of sensor fusion (Mitchell, 2007) with those of movement contamination of sensor output (von Holst &

Mittlestaedt, 1950).

A number of local solutions or piece-meal remedies for one or another part of this problem complex are conceivable. Insects, for example, rely on a variety of mechanisms of local feedback, gating, efference copy, inter-modal coordinate transformations, and perhaps even forward models to this end (see examples reviewed by Webb, 2004; also Altman & Kien, 1989). More centralized brains than those of insects offer the possibility of re-casting the entire sensors-in-motion problem in the form of a comprehensive, multimodal solution. In so doing, the fundamental role of gaze displacements in the orchestration of behavior can be exploited to simplify the design of the requisite neural mechanism.

The first sign of evolving action in the logistics of the brain's control of behavior is typically a gaze movement. Peripheral vision suffices for many purposes of ambient orientation and obstacle avoidance (Marigold, 2008; Trevarthen, 1968; Zettel, Holbeche, McIlroy, & Maki, 2005), yet when locomotion is initiated or redirected towards new targets or planned while traversing complex terrain, the gaze leads the rest of the body by fixating strategic locations ahead (Marigold & Patla, 2007). This is even more so for reaching and manipulative activity, down to their finely staged details. Fine-grained behavioral monitoring of eyes, hand and fingers during reaching and manipulation in the laboratory has disclosed the lead given by the gaze in such behavior (Johansson, Westling, Bäckström, & Flanagan, 2001). Arm and fingers *follow* the gaze as if attached to it by an elastic band. In fact the coupling of arm or hand to the gaze appears to be the brain's default mode of operation (Chang, Papadimitriou, & Snyder, 2010; Gorbet & Sergio, 2009).

The centrality of gaze shifts, also called orienting movements, in the orchestration of behavior makes them the brain's primary and ubiquitous output. The gaze moves through combinations of eye and head movements and these can be modelled – to a first approximation – as rotatory displacements of eyes in orbit and head on its cervical pivot,

using a convenient rotation-based geometry (see Masino, 1992; Merker, 2007a, p. 72; Smith, 1997).¹ This opens the possibility of simplifying the transformations needed to manage a good portion of movement-induced sensory displacement and misalignment of sensory maps during movement by casting them in an analog format adapted to such a rotation-based geometry. The orienting domain is a term introduced here for a hypothetical format that does so by actually nesting a map of the body within a map of the world, concentrically around an egocentric conjoint origin.

Let the brain, then, equip itself with an analog multi-modal central model (neural simulation space) in which an analog (spatially mapped) model of the body is spatially nested within an analog (spatially mapped) model of the enclosing world. Within this framework, let all *global* sensory displacement be introduced into that simulation as *body movement* (of one kind or another) relative to a stationary model world (Brandt, Dichgans, & Koenig, 1973), and let all global mismatches between sensory maps be reflected as displacement relative to one another of sensor-bearing parts of the model body (such as eyes relative to head/ears). This artificial allocation of movement between model world and model body has its ultimate foundation in the clustering of correlated sensory variances during random effector movement (for which see Philipona, O'Regan, Nadal, & Coenen, 2004; Dean et al., 2004). It presupposes a common geometric space for body and world within which their separation is defined, but not necessarily a metrical implementation of its mode of operation (Thaler & Goodale, 2010).

To exploit the simplifying geometry of nested rotation-based transformations, the origin of the geometric space shared by body and world must be lodged inside the model body's head representation, that is, the space must be egocentric.² During the ubiquitous gaze shifts of orienting this egocentric origin remains fixed with respect to the world, while the head map turns around it, *interposed* between egocenter and stabilized model world. Translatory and other locomotion-related global sensory effects would be registered in the simulation space as

displacements of the egocentric origin with respect to the world map, the origin carrying the body map with it relative to world. Note that these stipulations are all couched in geometric terms, and imply no commitment regarding the manner in which they might be implemented neurally, whether through gain-fields or other means (see footnote 1). Figure 1 illustrates the principle of the orienting domain in minimal outline.

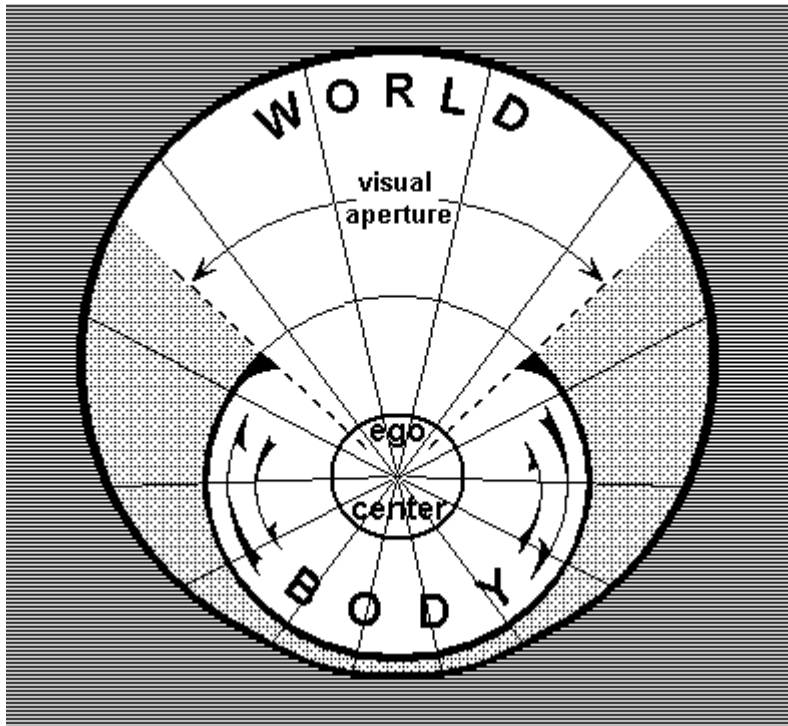


Figure 1. A minimal sketch of the orienting domain. A neural space organized as a rotation-based geometry is partitioned into three nested zones around the egocentric origin: egocenter, body zone and world zone. The latter two house spatial maps supplied with veridical content reflecting, in analog fashion, circumstances pertaining to the physical body and its surrounding world, respectively. In this mapping, global sensory motion is reflected as movement of the body alone relative to world, indicated by curved arrows (such as – in this case – might map, and hence compensate for, sensory effects of a gaze movement). The rotation-based transformations that supply the means for such stabilization of the sensory world during body movement require the geometric space to be anchored to an origin inside the head representation of the body zone. The device "visual aperture" marks the cyclopean aperture discussed in the penultimate section of the text. Auditory space, in contrast to that of vision, encompasses the entire "world" zone, including its shaded portion.

So far this model scheme is only an attempt to manage the sensors-in-motion problem by segregating its natural clusters of correlated variances into the separate zones of mobile and deformable body on the one hand and enclosing stable world on the other. Needless to say this

model sketch employing rotation-based nesting is a bare-bones minimum only. To accommodate realistic features such as limb movements it must be extended through means such as supplemental reference frames centered on, say, shoulder or trunk (see McGuire & Sabes, 2009). These might be implemented by yoking so-called gain-fields associated with limbs to those of the gaze (Chang et al., 2009), which would directly exploit the leading behavioral role of the gaze emphasized in the foregoing. However implemented, the centrality and ubiquity of gaze movements in the orchestration of behavior means that a simplification of the brain's sensors-in-motion problem is available in the nested format proposed for the orienting domain.

3.2 The decision domain: Triple play in the behavioral final common path

Given a workable solution to the sensors-in-motion problem, the brain as controller faces the over-arching task of ensuring that behavior – the time series of bodily locations and conformations driven by skeletal muscle contractions – comes to serve the many and fluctuating needs of the body inhabited by the brain. In so doing it must engage the causal structure of a world whose branching probability trees guarantee that certain options will come to exclude others (Shafer, 1996). With branches trailing off into an unknown future there is, moreover, no sure way of determining the ultimate consequences of choosing one option over another. Potential pay-offs and costs, and the various trade-offs involved in alternate courses of action, are therefore encumbered with a measure of inherent uncertainty. Yet metabolic facts alone dictate that action must be taken, and choices therefore made, necessarily on a probabilistic basis.

The world we inhabit is not only spatially heterogeneous in the sense that things like shelter, food and mates are often not to be found in the same place, it is temporally lively such that the opportunities it affords come and go, often unpredictably so. Since the needs to be filled are diverse and cannot all be met continuously, they change in relative strength and therefore

priority over time, and compete with one another for control over behavior (McFarland & Sibly, 1975). Few situations are entirely devoid of opportunities for meeting alternate needs, and one or more alternatives may present themselves at any time in the course of progress toward a prior target. The utility of switching depends in part on when in that progress an option makes its appearance. Close to a goal it often makes sense to discount the utility of switching (McFarland & Sibly, 1975), unless, of course, a "windfall" is on offer. Keeping options open can pay, but the capacity for doing so must not result in dithering.

The liveliness of the world sets the pace of the controller's need to update assessments, and saddles it with a perpetual moment-to-moment decision process regarding "what to do next." In the pioneering analysis just cited, McFarland and Sibly introduced the term "behavioral final common path" for a hypothetical interface between perceptual, motor, and motivational systems engaged in a final competitive decision process determining moment-to-moment behavioral expression. In a previous publication I sketched how savings in behavioral resource expenditure are available by exploiting inherent functional dependencies among the brain's three principal "task clusters" (Merker, 2007a), and to do so the brain needs such an interface, as we shall see.

The three "task clusters" consist of selection of targets of approach or avoidance in the world ("target selection"), the selection of the appropriate action for a given situation and purpose ("action selection"), and the ranking of needs by motivational priority ("motivational ranking"). Though typically treated as separate functional problems in neuroscience and robotics, the three are in fact bound together by intimate mutual dependencies, such that a decision regarding any one of them seldom is independent of the state of the others (Merker, 2007a, p. 70). As an obvious instance, consider prevailing needs in their bearing on target selection. More generally, bodily action is the mediator between bodily needs and opportunities in the world. This introduces the on-going position, trajectory, and energy

reserves of the body and its parts as factors bearing not only on target selection (cf. Körding & Wolpert, 2006) but on the ranking of needs as well. Thus the three task clusters are locked into mutual dependencies.

The existence of these dependencies means that savings are available by subjecting them to an optimizing regimen. To do so, they must be brought together in a joint decision space in which to settle trade-offs, conflicts and synergies among them through a process amounting to multiple constraint satisfaction in a multi-objective optimization framework (for which see: Pearl 1988; Tsang, 1993). Each of the three task clusters is multi-variate in its own right, and must be interfaced with the others without compromising the functional specificities on which their mutual dependencies turn. Those specificities include, for sensory systems, the need to be represented at full resolution of sensory detail, since on occasion subtle sensory cues harbor momentous implications for the very next action (say, a hairline crack in one of the bars of a cage housing a hungry carnivore).

Moreover, constraint-settling interactions among the three must occur with swiftness in real time, since it is over the ever-shifting combination of states of the world with the time series of bodily locations and conformations under ranked and changing needs that efficiency gains are achievable. Hence the need for a *high-level* interface *late* in the run-up to behavioral expression, i.e. McFarland & Sibly's "behavioral final common path." The interactions must also contend with the inconvenience that its contributors originally adhere to different data formats. Not only do the sensory systems differ among themselves in this regard, but they must be related to body conformations, and both must interact with motivational signals reflecting needs.

In the aggregate these diverse requirements for a constraint satisfaction interface between target selection, action selection, and motivational ranking may appear daunting, but they need not *in principle* be beyond the possibility of neural implementation. As first suggested

by Geoffrey Hinton (Hinton, 1977), a large class of artificial so called neural networks are in effect performing multiple constraint satisfaction (Rumelhart, Smolensky, McClelland, & Hinton, 1986). Algorithmically considered, procedures for constraint satisfaction that rely on *local* exchange of information between variables and constraints (such as "survey propagation") are the ones that excel on the most difficult problems (Achlioptas, Naor, & Peres, 2005). They are accordingly amenable to parallel implementation (Mézard & Mora, 2009), suggesting that the problem we are considering is neurally tractable.

A number of circumstances conspire to make the geometry of the orienting domain a promising framework for parallel implementation of constraint satisfaction among the brain's three principal task clusters. Two of these – target selection and action selection – are directly matched by the "world" and "body" zones of the orienting domain, already cast in parallel, analog (spatial) formats. Even on their own, their nested arrangement must satisfy a number of mutual constraints to set contents of the analog body map properly turning and translating inside those of the stabilized world map, using information derived from sources such as cerebellar "decorrelation" (Dean et al., 2004), vestibular head movement signals, eye movements, and visual flow patterns (Brandt et al., 1973). Moreover, implemented in the setting of the orienting domain, constraint satisfaction for the decision domain would be spared the need to address discrepant sensory alignment and data formats already managed in the arrangement of the orienting domain.

Add to this the circumstance, already noted, that the gaze leads the rest of the body in the execution of behavior (Johansson et al., 2001). This means that the *decision* domain's most immediate and primary output – the "very next action" – typically is a gaze shift, i.e. the very same combined eye and head movements that furnish the rationale for casting the *orienting* domain in nested analog format. Taken together these indications suggest that the two domains lend themselves to joint implementation in a single unitary mechanism.

3.3 A forum for the brain's final labors

The fact that both the decision domain and the orienting domain find themselves perpetually perched on the verge of a gaze movement means that there is no time to conduct constraint satisfaction by serial assessment of the relative utility of alternative target-action-motivation combinations. To be accomplished between gaze shifts, constraints must be settled by a dynamic process implemented in parallel fashion, as already mentioned. In this process, the "late" position of both domains in the run-up to overt behavior (i.e. the behavioral final common path) is a major asset: at that point all the brain's interpretive and inferential work has been carried as far as it will go before an eye movement introduces yet another change in its assumptions. Image segmentation, object identity, and scene analysis are as complete as they will be before a new movement is due. Since the neural labor has been done, there is every reason to supply its results in update form to the orienting domain mechanism involved in preparing that movement.

The orienting domain would thus host a real-time synthetic summary of the brain's interpretive labors, and is ideally disposed for the purpose. The spatial nesting of analog neural maps of body and world in itself imposes a requirement for three spatial dimensions on their shared geometric space. Extending such a space to accommodate arbitrarily rich three-dimensional object detail is a matter of expanding the resolution of this three-dimensional space, but involves no new principle. While the brain's many preliminaries require vast neural resources, a space hosting a summary synthesis of their outcomes does not. A few million neurons presumably should suffice to compactly represent our visual world at full resolution in its seven basic dimensions of variation (Adelson & Bergen, 1991; Lennie, 1998, pp. 900-901; Rojer & Schwartz, 1990, p. 284; see also Watson, 1987). Vision is by far our most demanding modality in this regard, so a full multimodal synthesis might be realizable in a neural mechanism composed of fewer than 10 million neurons, and possibly far fewer than

that.

In and by itself such a mechanism would supply no more than a format for whatever facts of the matter the rest of the brain extracts from its many sources of information by inferential means. Within its confines these results would be concretely instantiated in analog fashion through updates to a fully elaborated multimodal neural model. Its contents would simulate, in other words, the current situation in three-space with full object-constellation spatial detail. It would provide the brain with a stand-in for direct access to body and world, denied it by its solitary confinement inside its skull.³ Every veridical detail entered into the neural simulation space from the rest of the brain would have the effect of fixing parameters of decision domain variables, foreclosing some action options, and opening others. The decision domain's options are then those that remain unforeclosed by the simulator's typically massive supply of veridical constraints. Real-time constraint satisfaction accordingly would be concerned only with residual free parameters in this rich setting of settled concrete fact. Think of the latter as "reality," of the action possibilities latent in residual free parameters as "options" within it, and of the neural mechanism as a whole as the brain's "reality simulator." Its principal task is essentially no more than to determine "how to continue" efficiently, given the rich set of constraints already at hand, determined by convergent input from the rest of the brain. In the language of the gestalt psychologist we might say that the task of the reality simulator is to educe "global good continuation" *given* the brain's best estimates of current states of world, body and needs. The reality simulator does not determine these best estimates. It is given them, and must decide only how best to continue in their light.

At this penultimate stage of the brain's final labors, the options reside in whatever combinations of targets in the world and bodily actions are still available for filling motivational needs. The process of selecting among alternate such combinations saddles a decision mechanism with a set of functional requirements, among which the most basic is

global and parallel access to both world and body zones. Much remains to be learned by formal modelling about how these requirements might be fulfilled in the proposed format, but one of its features is bound to figure in any plausible solution, namely the convenience that the orienting domain contains a location with equal and ubiquitous access to both body and world zones. That location is the egocenter serving as origin for the geometry that holds the two zones together. As such it maintains a perpetually central position vis-a-vis the ever shifting constellations of world and body, and accordingly offers an ideal nexus for implementing decision making regarding their combined states.

Let its place in the system be taken, then, by a "miniature analog map," spatially extended and housing an intrinsic connectivity – presumably inhibitory – dedicated to competitive decision making (Merker, 2007a, p. 76; see also Richards, Seung, & Pickard, 2006). This central map would be connected with both world and body zones, in parallel and in tandem as far as its afferents go, but with principal efference directed to eye and head movement control circuitry associated with the body zone (cf. Deubel & Schneider, 1996). Whether to regard the competitive mechanism thus lodged at the system's egocentric origin as a decision maker or a monitoring function would seem to be a matter more of taste than of principle. It would depend in part on what stage in its progress towards its principal output – triggering the next gaze shift – is under consideration, as well as on the level of situational pressure on decision making under which it happens to be operating at the time. It is marked "e" in Fig. 2.

So far this decision nexus lacks one of the three principal sources of efference it needs in order to settle residual options among the brain's three task clusters, namely efference from the composite realm of motivational variables. Again, the orienting domain offers a convenient topological space, so far unused, through which a variety of biasing signals of this kind can be brought to bear on the decision mechanism. While the world zone must extend up to the very surface of the analog body, there is nothing so far occupying the space inside that

boundary, between the analog body surface and the decision nexus occupying the egocenter inside its head representation. That space can be rendered functional through the introduction of a variety of biasing signals, motivational ones among them, along and across the connectivity by which the decision nexus is interfaced with the body and world zones. This would embed the miniature analog decision map in a system of afferents of diverse origin injecting bias into its decision process, as depicted schematically in Fig. 2.




Figure 2. Two constituents of the "decision domain" embedded in the schematism of the orienting domain of Fig. 2. The egocenter zone (e) has been equipped with a decision making mechanism based on global mutual inhibition or other suitable connectivity. It is connectively interfaced with both the world and body maps, as before, but in addition with a zone of afferents interposed between the analog body surface and the decision mechanism. This zone introduces biasing signals derived from a set of motivational systems and other sources into the operation of the central decision mechanism. These biases are depicted as "sliding" annular sectors, each representing a different motivational system (hunger, fear, etc.). Each has a default position from which it deviates to an extent reflecting the urgency of need signalled by a corresponding motivational system outside of the orienting domain itself (collectively depicted in the lower portion of the figure). The central decision mechanism "e" is assumed to direct its output first and foremost to circuitry for eye and head movement control associated with the body zone.

Interposed between the analog body surface and the egocentric decision nexus, this multifaceted system of extrinsically derived signals would interact with those derived from body and world zones in their influence on the central decision nexus, introducing current values of motivational state variables into the constraint satisfaction regimen as a whole (see Sibly & McFarland, 1974 for a state space treatment of motivation). In keeping with their biasing function they would not assume the spatially articulated forms characterizing the contents of world and body zones, but something more akin to forces, fields, and tensional states. Even then, each would have some specificity reflecting the source it represents, along with more or less pronounced differential localization within the space enclosed by the analog body surface. They would in effect supply the neural body with what amounts to "agency vectors," animating it from within, as it were. Motivational needs undoubtedly represent the principal source of such signals, but the same logic allows other signals of global import, including some memory-related ones, to enter the scheme in this manner. More particular memory-derived information, such as the learned significance of an object, would plausibly enter the scheme in association with that object more specifically.

The introduction of the biasing mechanism into the body interior of the reality simulator completes in outline the mechanism as a whole. Each of its components – three separate "content zones" hosting bias, body and world contents nested and arrayed in tandem around a central egocentrically positioned decision nexus – is essential to its mode of operation. The mechanism is unlikely to serve any useful purpose in the absence of any one of them. It is thus a functionally unitary mechanism which allows the highly diverse signals reflecting needs, bodily conformations and world circumstances to interact directly. Its operation supplies, in effect, a means of mutual interaction among them (a functional "common currency": see McFarland & Sibly, 1975; Cabanac, 1992) that allows the brain to harvest in real time the savings hidden among their multiple mutual dependencies.

The substantial and mutually reinforcing advantages of implementing constraint satisfaction among the brain's principal task clusters in the setting of the orienting domain suggest that the brain may in fact have equipped itself with an optimizing mechanism of this kind.⁴ Whether it actually has done so can only be determined by empirically canvassing the brain for a candidate instantiation of a neural system that matches the functional properties of the proposed mechanism point for point. Such a search would have to be conducted without prejudging the level or levels of the neuraxis that might host the dedicated mechanism (Merker, 2007a), and with cognizance of the possibility of its quite compact implementation, as outlined above. Here, however, we are concerned only with the functional implications and consequences of such a hypothetical arrangement, to one of which we turn next.

3.4 A curious consequence of combined implementation

All three content categories featured in the joint mechanism, whether as biases, body conformations or world constellations, have this in common that their specific momentary states and contents are determined by brain mechanisms outside the confines of the reality simulator itself. Its own structural arrangement shorn of its externally supplied contents is no more than a nested neural scaffolding designed to provide as efficient a format as possible for hosting a running synthetic summary of the interpretive work of the rest of the brain, for the sole benefit of the centrally placed decision process. In its own terms this entire mechanism is, in other words, pure *functional format* designed for decision-making over all possible combinations of motivated engagement between body and world, in the absence of any specific motive, body conformation or constellation of the world. It is a format, in other words, of ubiquitously *open options*.

This aspect of the mechanism follows directly from its postulated implementation of a running constraint satisfaction regimen. To serve in this capacity the simulator must be capable of representing every possible motivated body-world combination in order to settle

optimality among them. We can call the mechanism that hosts such a process a "relational plenum." As we saw in the previous section, the entry of veridical content from the rest of the brain into this framework forecloses open option status for the corresponding portions of this plenum, leaving the remainder causally open. In the end, what remains of these open causal options at the point where the next action is imminent is a matter only for the centrally placed decision mechanism. They are the set of options with which its decision process engages to settle on a choice that triggers the next gaze movement. We might figuratively picture this in terms of residual options being "interposed," functionally speaking, between the decision nexus and the currently specified contents of the mechanism as a whole.

The role of residual options as mediators between decision nexus and contents is played out in a decidedly asymmetric context: the egocentric format of the mechanism as a whole ensures that the decision nexus itself is categorically excluded from the mechanism's contents. It occupies the central location *from which* all contents are accessed in such a geometry and its implementing connectivity. Taken together these circumstances amount to a global functional bipartition of the simulator into a central decision or monitoring nexus on the one hand and, on the other, a relational plenum (of biases, body, and world) whose given contents are interfaced with the centrally placed decision nexus *across* intervening portions of the plenum that remain in open option status. Such a partitioning into a decision nexus or monitor on the one hand, and monitored content on the other, via an intervening functional space of causal options, in fact supplies the most fundamental condition for a conscious mode of operation, to be considered more fully in the section that follows.

4 INADVERTENTLY CONSCIOUS

It is time to step back from this conjectural model construction to see what has been done. A number of neural mechanisms with particular design features have been proposed as solutions

to problems encountered by a brain charged with acting as body controller from a position of solitary confinement inside its skull. The design features were introduced not as means to make the brain conscious, but rather to provide solutions for functional problems the brain inevitably faces in informing itself about the world and controlling behavior. The logistical advantages of implementing a mechanism of multiple constraint satisfaction for optimal "task fusion" within a mechanism serving "sensor fusion" suggested an integrated solution in the form of a reality simulator. Within its nested analog format the brain's interpretive labors would receive a perpetually updated synthetic summary reflecting its best estimate of the veridical state of its surroundings, of its body, and of its needs. The core of its egocentric frame of reference would host a decision nexus spatially interfaced with a world map from inside a body map, and subject to action dynamics driven, ultimately, by motivational needs and serving behavioral optimization.

A conservative version of the foundational conjecture for the paradigm of consciousness proposed here states that such an implementation of a comprehensive regimen of constraint satisfaction among the brain's principal task clusters within the framework of an egocentrically nested analog solution to its sensors-in-motion problem defines a conscious mode of operation on the part of the joint mechanism. This claim rests, ultimately, on the nature of a fundamental functional asymmetry that inheres in the reality simulator's mode of operation. That asymmetry is geometric as well as causal. In geometric terms, the decision nexus, by virtue of its location at the origin of the simulator's egocentric geometry, stands in an inherently asymmetric (perspectival) relation to any content defined by means of that geometry. In causal terms, open options intervene between the decision nexus and the simulator's veridically specified contents. Options are inherently asymmetric: they are options only *for that* which has them, which in this case is the decision mechanism at the egocentric core of the reality simulator. Causally speaking, the options are that *about which* the decision

nexus makes its decisions, and geometrically speaking the decision nexus is that *for which* the veridically specified contents of the rest of the mechanism are in fact contents. In the setting of the simulator's spatial organization, this places the decision nexus "in the presence of" its monitored contents in the sense that its egocentrically defined "here" is separated from the "there" of any and all contents, and relates to them via relationally unidirectional causal options. Such a state of affairs, I submit, defines a conscious mode of operation by virtue of these facts alone (see Merker, 1997, for additional detail)

To be clear on this crucial point of attribution of conscious status: it is not a matter of decision making itself nor of its occurrence in a neural medium. Decisions are continually being made in numerous subsystems of the brain – in fact wherever outcomes are settled competitively – without for that reason implying anything regarding conscious status. It is only in the setting of the orienting domain, on account of its egocentric and spatially nested analog format, that decision making of the particular kind we have considered entails a global partitioning of the decision space into an asymmetric relation between monitor and monitored, marked by an intervening set of causal options. It is on account of this functional format, and of it alone, that such a mechanism can be said to operate consciously, according to the present account. That is, it operates consciously by virtue of this functional format, and not by virtue of anything that has been or needs to be added to it "in order to make it conscious."

The only natural setting in which such a format is likely to arise would seem to be the evolution of centralized brains, given the numerous specific and interlocking functional requirements that must conspire in order to fashion such a format. Its functional utility is predicated solely on the predicament of a brain captive in a skull and under pressure to optimize its efficiency as controller. Since the proposed mechanism would generate savings in behavioral resource expenditure, it would hardly be surprising if some lineages of animals, our own included, had in fact evolved such a mechanism. If, therefore, the claim that such a

functional format defines a conscious mode of operation is sound, it would be worth examining the thesis that it is the so far hypothetical presence of such a mechanism in our own brains that accounts for our status as conscious beings. For that to be the case we ourselves would have to be a part of – and a specific part of – that mechanism. This follows from the fact that the functional asymmetry at the heart of the mechanism insures that the only way to attain to consciousness on its terms is to occupy the position of egocentrically placed decision maker within it. Let us examine, therefore, the fit of that possibility with some of what we know about our own conscious functioning.


Consider, first, the curious format of our visual waking experience, that namely, by which we face, from a position inside our head, a surrounding panoramic visual world through an empty, open, cyclopean aperture in our upper face region. Anyone can perform the Hering triangulation to convince themselves that the egocentric perspective point of their visual access to the world is actually located inside their head, some four centimeters right behind the bridge of the nose (Howard & Templeton, 1966; Roelofs, 1959). That, in other words, is the point "from where we look." But that point, biology tells us, is occupied and *surrounded on all sides* by biological tissues rather than by empty space. How then can it be that looking from that point we do not see body tissues, but rather the vacant space through which in fact we confront our visual world in experience?

From the present perspective such an arrangement would be the brain's elegant solution to the problem of implementing egocentric nesting of body and world, given that the body in fact is visually opaque, but the egocenter must lodge inside the analog body's head, for simplicity of compensatory transformations between body and world. The problem the brain faces in this regard is the following: the body must be included in the *visual* mapping of things accessible from the egocentric origin inside the body's head. Such inclusion is unproblematic for distal parts of the body, which can be mapped into the simulator as any

other visual objects in the world. However, in the vicinity of the egocenter itself, persistence in the veridical representation of the body as visually opaque would block egocenter visual access to the world, given the egocenter's location inside the head representation of the body.

In this mapping quandary the brain has an option regarding the design of an analog neural body that is not realizable in a physical body, namely the introduction of a neural fiction – cast in the format of naive realism – of a cyclopean aperture through which the egocenter is interfaced with the visual world from its position inside the head (Merker, 2007a, p. 73). But this is exactly the format in which our visual experience demonstrably comes to us. We actually find ourselves looking out at the world from inside our heads through an oval empty aperture. This view, though for one eye only, is captured in Ernst Mach's classical drawing, reproduced here as Figure 3 (Mach, 1897). When both eyes are open the aperture is a symmetrical ovoid within which the nose typically disappears from view. What then is this paramount fact of our conscious visual macrostructure other than direct, *prima facie*, evidence that our brain in fact is equipped with a mechanism along the lines proposed here, and that we do in fact form part of this mechanism by supplying its egocentric perspectival origin?

Figure 3. Ernst Mach's classical rendition of the view through his left eye. Inspection of the drawing discloses the dark fringe of his eyebrow beneath the shading in the upper part of the figure, the edge of his moustache at the bottom, and the silhouette of his nose at the right-hand edge of the drawing. These close-range details framing his view are available to our visual experience, particularly with one eye closed, though not as crisply defined as in Mach's drawing. In a full cyclopean view with both eyes open the scene is framed by an ovoid within which the nose typically disappears from view (see Harding, 1961, for a detailed first person account). Apparently Mach was a smoker, as indicated by the cigarette extending forward beneath his nose. The original drawing appears as Figure 1 in Mach (1897, p. 16). It is in the public domain, and is reproduced here in a digitally retouched version, courtesy of Wikimedia (http://commons.wikimedia.org/wiki/File:Ernst_Mach_Innenperspektive.png).



This body which we can see and touch and which always is present wherever we are and obediently follows our every whim would accordingly be the model body contrived as part of the brain's reality simulator. And this rich and complex world that extends all around us and stays perfectly still even when we engage in acrobatic contortions would be the brain model's simulated world, stabilized at considerable neural expense. How else could that world remain

unaffected by those contortions, given that visual information about it comes to us through signals from our retinae, organs which flail along with the body during those contortions?

From this point of view a number of phenomena pertaining to the nature and contents of our consciousness can be interpreted as products of the workings of the proposed reality simulator and of our suggested place within its logistics. In our capacity of egocentric perspective point and decision nexus in this mechanism we would, on the present account and in agreement with commonplace observation, be maneuvering our body in accordance with a variety of impulses and tensional states – the biasing signals – playing in the interior of the space enclosed by our body surface. Motivational signals such as hunger, fear, bladder distension, and so forth, do in fact enter our consciousness as occurrences in various parts of our body interior (such as our chest region). Each of these variously distributed biases feels a bit different and makes us want to do different things (Izard, 1991, 2007; Sachs, 1967). Thus bladder distension is not only experienced in a different body location than is hunger or anger, it feels different from them, and each impels us to do different things. Common to them all is their general, if vague, phenomenological localization to the body interior, in keeping with what was proposed in the section devoted to the decision domain.

Far from all bodily systems and physiological mechanisms are thus able to intrude on our consciousness, or have any reason to do so. As cogently analyzed by Morsella and colleagues (Morsella, 2005; Morsella & Bargh, 2007; Morsella, Berger, & Krieger, 2010), those among them that do so involve, in one way or another, action on the environment (or on the body itself) by musculoskeletal means. This principle fits well with the present perspective, which traces the very existence and nature of consciousness to functional attributes of a mechanism designed to optimize exactly such behavioral deployment. Thus, the regulation of respiration is normally automatic and unconscious, but intrudes on consciousness in the form of a feeling of suffocation when blood titres of oxygen and carbon dioxide go out of bounds (Liotti,

Brennan, Egan, et al., 2001). Correcting the cause of such blood gas deviations may require urgent action on the environment (say, to remove a physical obstruction from the airways or to get out of a carbon dioxide filled pit). The critical nature of the objective matches the urgency of the feeling that invades our consciousness in such emergencies. For additional considerations and examples bearing on the relation between motivational factors and consciousness, see Cabanac (1992, 1995) and Denton, McKinley, Farrell, & Egan (2009).

Just as many bodily processes lack grounds for being represented in consciousness, so do many neural ones. As noted in the introduction, the busy neural traffic that animates the many way stations along our sensory hierarchies is not accessible to consciousness in its own right. Only its final result – a synthetic product of many such sources conjointly – enters our awareness: the rich and multimodal world that surrounds us. There is no dearth of evidence regarding neural activity unfolding “implicitly” without entering consciousness (for vision alone, see Rees, 2008, and references therein). This includes activity at advanced stages of semantic interpretation, motor preparation at the cortical level, and even instances of prefrontal executive activity (Dehaene et al., 1998; Eimer & Schlaghecken, 1998; Frith, Perry, & Lumer, 1999; Gaillard, Del Cul, Naccache, et al., 2006; Lau & Passingham, 2007; Luck, Vogel, & Shapiro, 1996; van Gaal, Ridderinkhof, Fahrenfort, Scholte, & Lamme, 2008).

One straightforward interpretation of this kind of evidence assigns conscious contents to a separate and dedicated neural mechanism, as proposed under the name of “conscious awareness system” by Daniel Schacter (Schacter, 1989, 1990). The present conception of a dedicated reality simulator is in good agreement with that proposal in its postulation of a unitary neural mechanism hosting conscious contents. In fact, on the present account, the simulator must exclude much of the brain’s on-going activity – sensory as well as motor – in order to protect the integrity of its constraint satisfaction operations. To serve their purpose those operations must range solely over the simulator’s internal contents with respect to one

another, and they should occur exclusively on terms dictated by the nested format that hosts them in the simulator. Such functional independence is presumably most readily achieved through local and compact implementation of the simulator in a dedicated neural mechanism of its own.⁵

This need to keep the constraint satisfaction operations of the reality simulator free of external interference has a crucial consequence bearing on the operation of our consciousness. The need to update the simulator's contents has been repeatedly mentioned, but never specified. As we have seen, the reality simulator is given the brain's best estimates of current states of world, body and needs as input to deciding, through constraint satisfaction among them, on an efficient "next move." With pressure on decision making at times extending down to subsecond levels (think of fencing, for example) constraint settling would typically fill the entire interval from one decision to the next. Externally derived parameter changes in the course of this might prolong constraint settling indefinitely. This means that contents must *not* be updated until a decision has been reached, and then updating must occur wholesale and precipitously. Wholesale replacement of contents is feasible, because the sources that delivered prior content are available at any time an update is opportune. The ideal time to conduct such a "refresh" or "reset" operation is at the time of the gaze shift (Henderson, Brockmole, & Gajewski, 2008; Singer, Zihl, & Pöppel, 1977) or body movement (Svarverud, Gilson, & Glennerster, 2010) that issues from a completed decision process. As already noted, such movements in and of themselves alter the presuppositions of decision making, rendering prior simulator content obsolete.

The same logic applies to instances in which sudden change is detected, signalled by a transient that attracts attention – and typically (but not necessarily) a gaze shift – to the change. That change also alters the presuppositions of the simulator's current operation, again favoring wholesale resetting of its contents. When transients are eliminated by stimulus

manipulation in the laboratory, a change that otherwise would be conspicuous goes undetected (Turatto, Bettella, Umiltà, & Bridgeman, 2003). Assuming, then, that the reality simulator's contents are subject to punctate and wholesale replacement, the reality it hosts has no other definition than the constellation of its current content alone, maintained as such only till the next reset. Since in present terms the simulator contents are the contents of our consciousness, this means that we should be oblivious to veridical sensory changes introduced at the exact time of a reset, and so we are, indeed, as demonstrated by the well documented phenomena of change blindness (Simons & Ambinder, 2005).

The only conscious contents that appear to reliably escape oblivion in the reset are those held in focal attention (Rensink, O'Regan, & Clark, 1997), a circumstance most readily interpretable as indicating a privileged relation between the contents of focal attention and memory (Turatto et al., 2003; see also Merker, 2007a, p. 77), allowing pre- and post-reset focal content to be compared. Focal attention and its contents accordingly would be the key factor maintaining continuity of consciousness across frequent resets of simulator content, as might be expected from the central role of a competitive decision mechanism of limited capacity at its operational core. The more intense and demanding the focal attentional engagement, the higher the barrier against its capture by an alternate content, as dramatically demonstrated in inattention blindness (Mack, 2003; Most, Scholl, Clifford, & Simons, 2005; see also Cavanagh, Hunt, Afraz, & Rolfs, 2010, for further considerations relevant to update operations).

More generally, our limited capacity to keep in mind or track independent items or objects simultaneously (Cowan, 2001; Mandler, 1975; Miller, 1956) presumably reflects the "late" position of the reality simulator (consciousness) in the brain's functional economy. As emphasized in the foregoing, the decision nexus engages only the final options to be settled in order to trigger the next gaze movement (and reset), and as such forms the ultimate

informational bottleneck of the brain as a whole. It receives convergent afference from the more extensive (though still compact) world, body and bias maps of the simulator, and they in turn are convergently supplied by the rest of the brain. Some such arrangement is necessary if the brain's distributed labors are to issue, as they must, in unitary coherent behavioral acts (McFarland & Sibly, 1975). Moreover, in its capacity as final "convergence zone" the decision nexus brings the informational advantages of the quite general 'bottleneck principle' to bear on the simulator's global optimization task (Damasio, 1989; Moll & Miikulainen, 1997; Kirby, 2002).

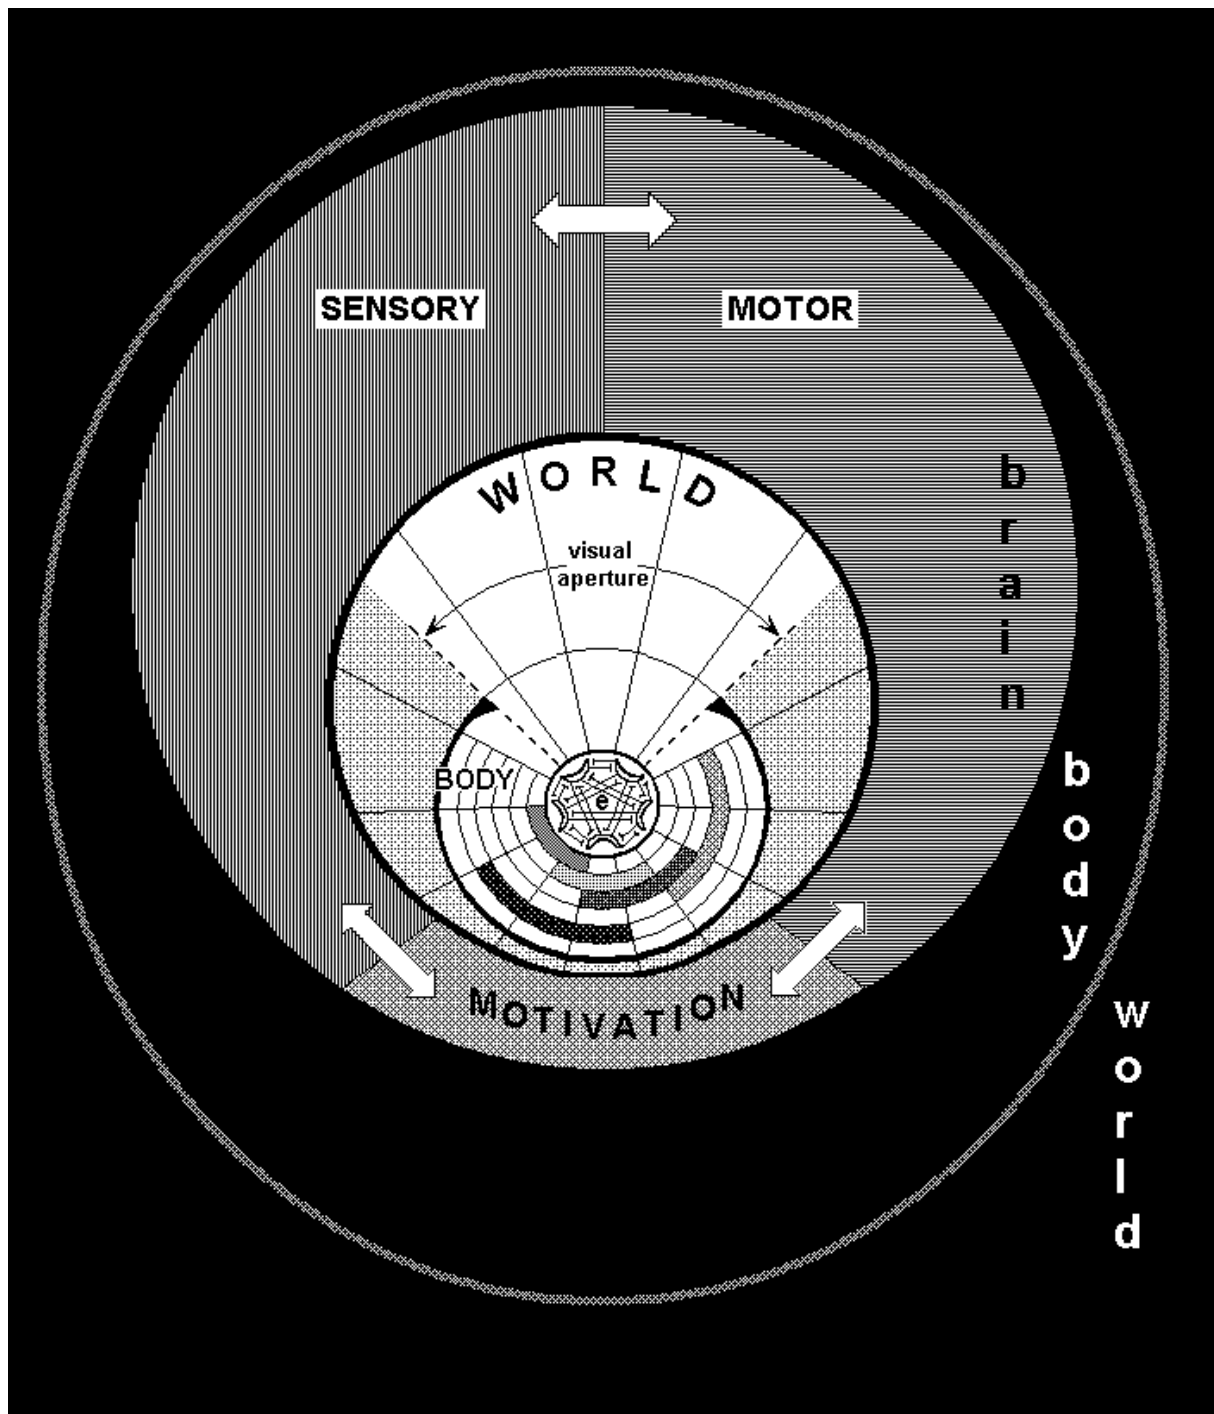
The crucial position of the decision nexus at the core of the reality simulator may account for a further, quite general, aspect of our consciousness as well: our sense of facing the world as an arena of possibilities within which we exercise choice among alternative actions. As we have seen, the simulator's operational logic interposes causal options between the decision nexus and the veridical contents of the simulator. Our sense of having options and making choices – a sense that presumably underlies the concept of free will – thus corresponds, on this account, to a reality. It follows as a direct corollary of our hypothesized status as decision nexus at the core of the brain's reality simulator, determining what to do next under the joint influence of a rich set of veridical constraints and residual causal options. This sense, moreover, would be expected to vary inversely with the pressure of events on decision-making, as indeed it appears to do on intuitive grounds.

Finally, to conclude this brief survey of aspects of the proposed mechanism that can be directly related to the nature of our conscious functioning, consider the fact that the objects of our sensory awareness have "sidedness" and "handedness." This fact cannot be captured by any set of measurements on the objects themselves, and was noted by William James as a puzzle requiring explanation (James, 1890, vol. II, p. 150). In present terms, it follows from the fact that in our position at the egocentric core of the nested geometry framing our

conscious contents we have no choice but to relate to those contents perspectively. It is a direct consequence of the egocentric arrangement of the reality simulator as a whole (see Green, 2003, for a concerted treatment of the issue).

Fig. 4 provides a summary diagram of the present conception as a whole, depicting the nested *phenomenal* space of the neural reality simulator in its nesting within a larger *noumenal* setting (to use Kant's precise terminology) of physical brain, physical body and physical world. All our experience moves within the confines of the phenomenal space alone, and it is the nesting of this entire phenomenal space (featuring egocenter, body and world) inside a physical (*noumenal*) brain, body and world that bears the name "nested ontology" in the title of this paper.

Figure 4. The full ontology of the consciousness paradigm introduced in the text. The joint orienting and decision mechanism illustrated in Fig. 2 supplies a neural format defining a conscious mode of operation by virtue of its internal organization, as explicated in the text. It forms a unitary and dedicated subsystem set off from the rest of the brain by the heavy black line, beyond which neural traffic is assumed to take place without consciousness. Broad white arrows mark interfaces across which neural information may pass without entering consciousness. The physical brain, which – except for its phenomenal subsystem – is devoid of consciousness is part of the physical body, in turn part of the physical world, both depicted in black in the figure. This embedding of a phenomenal space reflecting the circumstances of a physical body in a physical world for purposes of remote control of that body on the part of a physical brain lacking direct access to body and world is the "nested ontology" referred to in the title of this paper.



5 CONCLUSION: A LONE BUT REAL STUMBLING BLOCK ON THE ROAD TO A SCIENCE OF CONSCIOUSNESS

As an exercise intended to bring out the ontological implications of the paradigm for consciousness introduced here, consider being asked to indicate the approximate location of

the brain which – on present terms – hosts the neural model that furnishes you with the reality you are at that moment experiencing. Consider trying to indicate its location in the experienced space it makes available to you, extending from inside your body, to its surface and beyond it through the world all the way to the horizon (Lehar, 2003; Trehub, 1991).

Where in that space of your conscious reality is the brain that I claim synthesizes that conscious reality located? The answer provided by the present perspective is, of course, that there is no such location because *that entire perceived space* is a neural artifice contrived for control purposes in a dedicated neural mechanism *inside* the brain you are asked to localize. To ask you to localize the *containing* brain *inside the space it contains* is to ask for the impossible, obviously.

Strictly speaking there is no unique location to point to, but if one nevertheless were to insist on trying to point, all directions in which one might do so would be equally valid. In particular, pointing to one's head would be no more valid than pointing in any random direction. That head, visually perceptible only at the margin of its cyclopean aperture from which one looks out at the world, is but a part of the model body inside the model world of the brain's neural reality simulator, as Schopenhauer clearly recognized in declaring this familiar body, which we can see and move, to be a picture in the brain (Schopenhauer, 1844, vol. II, p. 271).

I hasten to offer my assurances that I am no less a slave to the ineradicable illusion of naive realism than anyone else. I cannot shake the impression that in perception I am directly confronting physical reality itself in the form of a mind-independent material universe, rather than a neurally generated model of it. Far from constituting counter-evidence to the perspective I have outlined, this unshakeable sense of the reality of experienced body and world supports it, because it is exactly what the brain's reality simulator would have to produce in order to work, or rather, *that is how it works*. It defines the kind of reality we can

experience, and its format is that of naive realism: through the cyclopean aperture in our head we find ourselves directly confrontating the visual world that surrounds us. It is our reality by default. This elaborate neural contrivance repays, or rather generates, our trust in it by working exceedingly well: the brain's model world is veridical (Frith, 2007). It faithfully reflects those aspects of the physical universe that matter to our fortunes within it, while sparing us the distraction of having to deal with the innumerable irrelevant forms of matter and energy that the physical universe actually has on offer, just as it spares us distraction by the busy neural traffic of most of the brain's activity.

Thus, for all practical purposes the deliverances of the reality simulator provide a reliable guide to the physical world beyond the skull of the physical body, in a manner similar to that of a "situation room" serving a general staff during wartime, from which operations on far away battle fields are conducted by remote control on the basis of the veridical model assembled in the situation room itself (Lehar, 2003). The corresponding neural model's usefulness in that regard is the reason it exists, if my account has any merit. And that usefulness extends beyond our everyday life to our endeavors in every area of science so far developed, because those endeavors have been concerned with defining in increasingly precise and deep ways the causal relations behind the surface phenomena of our world. Our reality model is an asset in these as in our other dealings with the world, because typically it reflects circumstances in the physical universe faithfully, and for aspects of the world that lie beyond our sensory capacities, suitable measuring instruments have been devised to bring them within range of those capacities. Accordingly, even a *conceptual* commitment to naive realism is perfectly compatible with these scientific endeavors. Conceptions of the ontological status of our *experience* does not affect their outcomes, because they are not concerned with the nature of our experience but with explaining the world, regarding which our experience supplies reliable guidance in most respects.

This is no longer so the moment the scientific searchlight is turned to the nature of experience itself, as in a prospective 'science of consciousness'. Here, the ontological status of experience itself is the principal question under investigation, along with its various characteristics, such as the scope and organization of its potential contents, its genesis, and its relation to the rest of the picture of reality science has pieced together for us with its help. Now, suddenly and uniquely, adherence to naive realism as a *conceptual* commitment, even in the form of a lingering tacit influence, becomes an impediment and a stumbling block. By its lights the world we experience is the physical universe itself rather than a model of it. Such a stance seriously misconstrues the scope and nature of the problem of consciousness, most directly by excluding from its compass the presence of the world we see around us. When the latter is taken for granted as mind-independent physical reality rather than recognized as a principal content of consciousness requiring explanation, the problem of consciousness contracts to less than its full scope. Under those circumstances, inquiry tends to identify consciousness with some of its subdomains, contents, or aspects, such as thinking, subjectivity, self-consciousness, an "inner life", "qualia" and the like.

Any such narrowing of the scope of the problem of consciousness allows the primary task of accounting for the very existence of experience itself to be bypassed, and promotes attempts to account for the world's "experienced qualities" (hence qualia) even before addressing the prior question of why there is a world present in our experience at all, or indeed why experience itself exists. "Experienced qualities" can be referred to our "inner life", the stirrings of thoughts in our heads and feelings in our breasts, and so might seem exempt from the problem of the external world. Yet our experience is not thus confined to our "inner life." It extends beyond it to encompass a wide and varied external world, bounded by the horizon, the dome of the sky, and the ground on which we walk (Frith, 2007; Lehar, 2003; Lopez, 2006; Revonsuo, 2006; Trehub, 1991; Velmans, 2008). The objects and events of the

world, whose attributes provide many an occasion for the events of our inner life, are no less contents of consciousness than the stirrings those attributes may occasion in us.

When consciousness is identified with our "inner life" the concept of "simulation" tends to be used in the sense of the manifest power of our imagination to create scenarios "in the mind's eye" for purposes of planning or fantasy. This is not the sense in which the concept of simulation has been used in this article, of course. The concept as used here refers to neural simulation of the entire synthetic framework and content of our reality, including the rich and detailed world that surrounds us when we open our eyes in the morning and which stays with us till we close them at night. From the present point of view, our additional capacity for imaginative simulation is derivative of this prior, massive, and more basic reality simulation, as are the scenes enacted in our dreams.

The distinctions of the past few paragraphs are made by way of clarification of the present perspective, and should not be taken to imply that these "inner life" topics lack interest or validity as objects of scientific scrutiny. As contents of consciousness they provide worthy topics of study in their own right, but their place in the subject matter of consciousness theory as a whole will remain moot until at least a provisional explanation has been advanced for the fact and existence of experience itself, the circumstances under which it arises, and the manner of its arrangement where it is present, as in our own case.

In this essay I have presented my approach to such questions, questions I take to define the primary subject matter of a prospective science of consciousness. It may seem ironic that pursuing those questions without heeding the alarms rung by naive realist predispositions should issue in an account of consciousness according to which its global format necessarily bears the stamp of naive realism. That, however, should be a reason to take my account seriously, because that is one of the basic attributes of our consciousness that any valid account of its nature must, in the end, explain. But that is also, according to my account, the

very attribute in which consciousness must abandon veridicality in order to function properly.

A sound theory of consciousness therefore must abandon, in its turn and on this point uniquely, trust in the deliverances of consciousness as guides to the realities we wish to understand.

Acknowledgement: I am indebted to Louise Kennedy for helpful suggestions regarding style and presentation, and also to Henrik Malmgren for helpful comments on matters of content.

Footnotes

1: "Rotation-based geometry" is a non-committal shorthand for a geometry of spatial reference that implements a nested system relating the rotational kinematics of eyes and head to target positions in the world (see Smith, 1997, for a useful introduction to rotational kinematics related to the eye and head movements of the vestibulo-ocular reflex, and Thaler & Goodale, 2010, for alternative geometries that might implement spatial reference). Much remains to be learned about the neural logistics of movement-related spatial reference, and the role of, say, gain fields in its management (Andersen & Mountcastle, 1983; Chang et al., 2009; see also Cavanagh et al., 2010).

2. This basic egocentricity does not exclude an adjunct oculocentric specialization for the eyes, as briefly touched upon later in the text. The center of rotation of the eye does not coincide with that of the head. Therefore the empirically determined visual egocenter – single despite the physical fact of the two eyes – lies in front of that of the auditory egocenter (far more cumbersome to chart empirically than the visual one, see Cox, 1999; Neelon, Brungart, & Simpson, 2004). The ears, of course, are fixed to the head and move with it. The egocenter in Fig. 1, 2 and 3 therefore lies closer to that of audition than to that of vision, a placement motivated by the fact that the limits of the visual aperture are determined largely by the bony orbit of the eyes, which is fixed to the head. Thus, a 45 degree rightward eye deviation extends the visual field by far less than 45 degrees to the right.

3. As such it would be in receipt of signals from any system in the brain, cortical or subcortical, relevant to decision making aimed at the very next action (typically a targetted gaze movement, as we have seen). The decision making in question therefore should by no means be identified with deliberative processes or prefrontal executive activity. Such activities serve as inputs among many others converging on the more basic and late decision process envisaged here (cf. Merker, 2007b, pp. 114, 118).

4. The present account violates two programmatic commitments of the "subsumption architecture" of behavior based robotics introduced by Brooks (1986), namely the stipulations "little sensor fusion" and "no central models." It would therefore seem to have to forego some of the advantages claimed for subsumption architectures, but this is only apparently so. The reality simulator of the present account is assumed to occupy the highest functional level (which is not necessarily cortical, see Merker, 2007a) of such an architecture without being the sole input to behavior control. The initial phase of limb withdrawal on sudden noxious stimulation and the vestibulo-ocular reflex are examples of behaviors which access motor control independently of the reality simulator. See Merker (2007a) pp. 69, 70, and 116 for further detail.

5. The cerebral cortex appears to offer a most inhospitable environment for such an arrangement. The profuse, bidirectional and exclusively excitatory nature of cortical inter-areal connectivity poses a formidable obstacle to any design requiring a modicum of functional independence (see Merker, 2004). There is also no known cortical area (or combination of areas) whose loss will render a patient unconscious (cf. Merker, 2007a). On the present account, the cerebral cortex serves, rather, as a source of much of the sophisticated information utilized by the simulator's reality synthesis, supplied to it by remote and convergent cortical projections. Candidate loci of multisystem convergence are of course available in a number of subcortical locations.

References

- Achlioptas, D., Naor, A., & Peres, Y. (2005). Rigorous location of phase transitions in hard optimization problems. *Nature*, *435*, 759–764.
- Adelson, E. H., & Bergen, J. R. (1991). The plenoptic function and the elements of early vision. In M. S. Landy & J. A. Movshon (Eds.), *Computational models of visual processing* (pp. 1–20). Cambridge, MA: MIT Press.
- Altman, J. S., & Kien, J. (1989). New models for motor control. *Neural Computation*, *1*, 173–183.
- Andersen, R. A., & Mountcastle, V. B. (1983). The influence of the angle of gaze upon the excitability of the light-sensitive neurons of the posterior parietal cortex. *Journal of Neuroscience*, *3*, 532–548.
- Bernshtein, N.A. (1967). *The co-ordination and regulation of movements*. Oxford: Pergamon Press.
- Blanke, O., Ortigue, S., Landis, T., & Seeck, M. (2002). Stimulating illusory own-body perceptions. *Nature*, *419*, 269–270.
- Blanke, O., Landis, T., Spinelli, L., & Seeck, M. (2004). Out-of-body experience and autoscopia of neurological origin. *Brain*, *127*, 243–58.
- Brandt, Th., Dichgans, J., & Koenig, E. (1973). Differential effects of central versus peripheral vision on egocentric and exocentric motion perception. *Experimental Brain Research*, *16*, 476–491.
- Brooks, R. A. (1986) A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation*, *2*, 14–23.
- Brouwer, A.-M., & Knill, D. C. (2007). The role of memory in visually guided reaching. *Journal of Vision*, *7*, 1–12. doi:10.1167/7.5.6.
- Cabanac, M. (1992). Pleasure: The common currency. *Journal of Theoretical Biology*, *155*, 173–200.
- Cabanac, M. (1996). The place of behavior in physiology. In M. J. Fregly & C. Blatteis (Eds.), *Handbook of physiology, environmental physiology* (pp. 1523–1536). Oxford: Oxford University Press.
- Cavanagh, P., Hunt, A. R., Afraz, A., & Rolfs, M. (2010). Visual stability based on remapping of attention pointers. *Trends in Cognitive Sciences*, *14*, 147–153.
- Chang, S. W. C., Papadimitriou, C., & Snyder, L.H. (2009). Using a compound gain field to compute a reach plan. *Neuron*, *64*, 744–755.

Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of System Science*, 1, 89–97.

Cowan, N. (2001). The Magical Number 4 in Short-term Memory: A Reconsideration of Mental Storage Capacity. *Behavioral and Brain Sciences*, 24, 87–185.

Cox, P. H. (1999). *An initial investigation of the auditory egocenter: Evidence for a "cyclopean ear."* Doctoral dissertation, North Carolina State University, Raleigh, NC.

Damasio, A. R. (1989). The brain binds entities and events by multiregional activation from convergence zones. *Neural Computation*, 1, 123–132.

Dean, P., Porrill, J., & Stone, J. V. (2004). Visual awareness and the cerebellum: possible role of decorrelation control. *Progress in Brain Research*, 144, 61–75.

Dehaene, S., Naccache, L., Le Clec'h, G., Koechlin, E., Mueller, M., Dehaene-Lambertz, G., van de Moortele, P.-F., & Le Bihan, D. (1998). Imaging unconscious semantic priming. *Nature*, 395, 597–600.

Denton, D. A., McKinley, M. J., Farrell, M., & Egan, G. F. (2009). The role of primordial emotions in the evolutionary origin of consciousness. *Consciousness and Cognition*, 18, 500–514.

Deubel, H., & Schneider, W. X. (1996). Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Vision Research*, 36, 1827–1837.

Eimer, M., & Schlaghecken, F. (1998). Effects of masked stimuli on motor activation: behavioral and electrophysiological evidence. *Journal of Experimental Psychology: Human Perception and Performance*, 24, 1737–1747.

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1, 1–47.

Frith, C. (2007). *Making up the mind: How the brain creates our mental world*. London: Blackwell.

Frith, C. D., Perry, R., & Lumer, E. (1999). The neural correlates of conscious experience : An experimental framework. *Trends in Cognitive Sciences*, 3, 105–114.

Gaillard, R., Del Cul, A., Naccache, L., Vinckier, F., Cohen, L., & Dehaene, S. (2006). Nonconscious semantic processing of emotional words modulates conscious access. *Proc Natl Acad Sci USA* 103:7524–7529.

Gallistel, C. B. (1999). Coordinate transformations in the genesis of directed action. In B.O.M. Bly & D.E. Rummelhart (Eds.), *Cognitive Science* (pp. 1–42). New York: Academic Press.

Gorbet, D. J., & Sergio L. E. (2009). The behavioural consequences of dissociating the spatial directions of eye and arm movements. *Brain Research*, 1284, 77–88.

Green, A. (2003). *The science and philosophy of consciousness*. Webpage <http://www.users.globalnet.co.uk/~lka/conz.htm>

Harding, D. E. (1961). *On having no head: Zen and the re-discovery of the obvious*. London: Arkana (Penguin).

Harting, J. K., Updyke, B. V., & Van Lieshout, D. P. (1992). Corticotectal projections in the cat: Anterograde transport studies of twentyfive cortical areas. *Journal of Comparative Neurology*, 328, 379–414.

Helmholtz, H. von (1867). *Handbuch der physiologischen Optik*, Vol. 3. Leipzig: Leopold Voss.

Henderson, J. M., Brockmole, J. R., & Gajewski, D. A. (2008). Differential detection of global luminance and contrast changes across saccades and flickers during active scene perception. *Vision Research*, 48, 16–29

Hering, E. (1879). *Spatial sense and movements of the eye*. (English translation by C.A. Radde, 1942). Baltimore: American Academy of Optometry.

Hinton, G. E. (1977). *Relaxation and Its Role in Vision*. PhD thesis, University of Edinburgh.

Howard, I. P., & Templeton, W. B. (1966). *Human spatial orientation*. New York: Wiley.

Izard, C. E. (1991). *The psychology of emotions*. New York: Plenum Press.

Izard, C. E. (2007). Levels of emotion and levels of consciousness. *Behavioral and Brain Sciences*, 30, 96–98.

James, W. (1890). *Principles of Psychology*, London: Macmillan.

Johansson, R. S., Westling, G., Bäckström, A. & Flanagan, J. R. (2001). Eye-hand coordination in object manipulation. *Journal of Neuroscience*, 21, 6917–6932.

Kawato, M. (1997). Bi-directional theory approach to consciousness. In M. Ito (Ed.), *Cognition, computation and consciousness* (pp. 233–248). Oxford: Clarendon Press.

Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9, 718–727.

Kawato, M., Hayakawa, H. & Inui, T. (1993). A forward-inverse optics model of reciprocal connections between visual cortical areas. *Network: Computation in Neural Systems*, 4, 415–422.

Kirby, S. (2002) Learning, bottlenecks and the evolution of recursive syntax. In T. Briscoe (Ed.), *Linguistic evolution through language acquisition: Formal and computational models* (pp.

173–204). Cambridge: Cambridge University Press.

Körding, K. P., & Wolpert, D. M. (2006) Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, *10*, 319–326.

Lau, H. C. & Passingham, R. E. (2007). Unconscious activation of the cognitive control system in the human prefrontal cortex. *Journal of Neuroscience*, *27*, 5805–5811.

Lehar, S. (2003). *The world in your head: A gestalt view of the mechanism of conscious experience*. Mahwah, NJ: Lawrence Erlbaum Associates.

Lennie, P. (1998). Single units and visual cortical organization. *Perception*, *27*, 889–935.

Liotti, M., Brannan, S., Egan, G., Shade, R., Madden, L., Abplanalp, B, Robillard, R., Lancaster, J., Zamarripa, F.E., Fox, P. T., & Denton, D. (2001). Brain responses associated with consciousness of breathlessness (air hunger). *Proceeding of the National Academy of Sciences (USA)*, *98*, 2035–2040.

Lopez, L. C. S. (2006). The phenomenal world inside the noumenal head of the giant: Linking the biological evolution of consciousness with the virtual reality metaphor. *Revista Eletrônica Informação e Cognição*, *5*, 204–228.

Luck, S. J., Vogel, E.K. & Shapiro K.L. (1996). Word meanings can be accessed but not reported during the attentional blink. *Nature*, *383*, 616–618.

Mach, E. (1897). *Contributions to the analysis of the sensations*. La Salle, Illinois: Open Court.

Mack, A. (2003). Inattention blindness: Looking without seeing. *Current Directions in Psychological Science*, *12*, 180-184.

Mandler, G. (1975). Memory storage and retrieval: Some limits on the reach of attention and consciousness. In P. M. A. Rabbitt & S. Dornic (Eds.), *Attention and Performance V* (pp. 499–516). New York: Academic Press.

Marigold, D. S. (2008). Role of peripheral visual cues in online visual guidance of locomotion. *Exercise and Sport Sciences Reviews*, *36*, 145–151.

Marigold, D. S., & Patla, A. E. (2007). Gaze fixation patterns for negotiating complex ground terrain. *Neuroscience*, *144*, 302–313.

Masino, T.. (1992). Brainstem control of orienting movements: Intrinsic coordinate system and underlying circuitry. *Brain, Behavior and Evolution*, *40*, 98–111.

McFarland, D. J. (1977). Decision making in animals. *Nature*, *269*, 15–21.

McFarland, D. J., & Sibly, R. M. (1975). The behavioural final common path. *Philosophical Transactions of the Royal Society (London)*, *270*, 265–293.

- McGuire, L. M. M., & Sabes, P. N. (2009). Sensory transformations and the use of multiple reference frames for reach planning. *Nature Neuroscience*, *12*, 1056–1061.
- Merker, B. (1997). *The common denominator of conscious states: Implications for the biology of consciousness*. Archived preprint: <http://cogprints.soton.ac.uk> (biology; theoretical biology).
- Merker, B. (2004). Cortex, countercurrent context, and dimensional integration of lifetime memory. *Cortex*, *40*, 559–576.
- Merker, B. (2005). The liabilities of mobility: A selection pressure for the transition to consciousness in animal evolution. *Consciousness and Cognition*, *14*, 89–114.
- Merker, B. (2007a). Consciousness without a cerebral cortex: A challenge for neuroscience and medicine. Target article & peer commentary. *Behavioral and Brain Sciences*, *30*, 63–110.
- Merker, B. (2007b). Grounding consciousness: The mesodiencephalon as thalamocortical base. Author's response. *Behavioral and Brain Sciences*, *30*, 110–134.
- Mézard, M., & Mora, T. (2009). Constraint satisfaction problems and neural networks: a statistical physics perspective. *Journal of Physiology (Paris)*, *103*, 107–113.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, *63*, 81–97.
- Mitchell, H. B. (2007). *Multi-sensor data fusion: An introduction*. Berlin: Springer-Verlag.
- Moll, M., & Miikkulainen, R. (1997). Convergence-zone episodic memory: Analysis and simulations. *Neural Networks*, *10*, 1017–1036.
- Morsella, E. (2005). The function of phenomenal states: Supramodular interaction theory. *Psychological Review*, *112*, 1000–1021.
- Morsella, E., & Bargh, J. A. (2007). Supracortical consciousness: Insight from temporal dynamics, processing-content, and olfaction. *Behavioral and Brain Sciences*, *30*, 100.
- Morsella, E., Berger, C. C., & Krieger, S. C. (in press). Cognitive and Neural Components of the Phenomenology of Agency: A Meta-Analytic View. ???
- Most, S., Scholl, B., Clifford, E., & Simons, D. (2005). What you see is what you set: Sustained inattentive blindness and the capture of awareness. *Psychological Review*, *112*, 217–242.
- Neelon, M. F., Brungart, D. S. & Simpson, B. D. (2004) The isoazimuthal perception of sounds across distance: A preliminary investigation into the location of the audio egocenter. *Journal of Neuroscience*, *24*, 7640–7647.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (2nd ed.). San Francisco: Morgan Kaufmann.

Philipona, D., O'Regan, J. K., Nadal, J.-P. & Coenen, O. J.-M. D. (2004). Perception of the structure of the physical world using multimodal unknown sensors and effectors. *Advances in Neural Information Processing Systems*, 16, 945–52.

Rees, G. (2007). Neural correlates of the contents of visual awareness in humans. *Philosophical Transactions of the Royal Society (London) B*, 362, 877-886

Rensink, R. A., O'Regan, J. C., & Clark, J. J. (1997). To see or not to see: the need for attention to perceive changes in scenes. *Psychological Science*, 8, 368–373.

Revonsuo, A. (2006). *Inner Presence, Consciousness as a Biological Phenomenon*. Cambridge, MA: The MIT Press.

Richards, W., Seung H. S., & Pickard, G. (2006). Neural voting machines. *Neural Networks*, 19, 1161–1167.

Roelofs, C. O. (1959). Considerations on the visual egocenter. *Acta Psychologica*, 16, 226–234.

Roger, A. S., & Schwartz, E. L. (1990). Design considerations for a space-variant visual sensor with complex logarithmic geometry. In *Proceedings of the 10th International Conference on Pattern Recognition* (pp. 278–285). Washington, DC: IEEE Computer Society Press.

Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. (1986). Schemata and sequential thought processes in PDP models. In D.E. Rumelhart & J.L. McClelland (Eds), *Parallel distributed processing: explorations in the microstructure, vol. 2: psychological and biological models* (pp. 7–57). Computational models of cognition and perception series. Cambridge MA: The MIT Press.

Sachs, E. (1967). Dissociation of learning in rats and its similarities to dissociative states in man. In J. Zubin & H. Hunt (Eds.), *Comparative psychopathology: Animal and human* (pp. 249–304). New York: Grune & Stratton.

Schacter, D. L. (1989). On the relations between memory and consciousness: Dissociable interactions and conscious experience. In H. L. Roediger, III & F. I. M. Craik (Eds.), *Varieties of Memory and Consciousness: Essays in Honour of Endel Tulving* (pp. 356-390). Mahwah, NJ: Lawrence Erlbaum Associates.

Schacter, D. L. (1990). Toward a cognitive neuropsychology of awareness: Implicit knowledge and anosognosia. *Journal of Clinical and Experimental Psychology*, 12, 155 – 178.

DOI: 10.1080/01688639008400962

Schopenhauer, A. (1844/1958). *The world as will and representation*, 2nd ed. Trans. E.F.J. Payne. New York: Dover Publications (paperback, 2 vols).

Shafer, G. (1996). *The art of causal conjecture*. Cambridge, MA: The MIT Press.

- Sibly, R. M., & McFarland, D. J. (1974). A state-space approach to motivation. In J. D. McFarland (Ed.), *Motivational control systems analysis* (pp. 213-250). London & New York: Academic Press.
- Simons, D. J., & Ambinder, M. S. (2005). Change blindness. Theory and consequences. *Current Directions in Psychological Science, 14*, 44–48.
- Singer, W., Zihl, J. & Pöppel, E. (1977). Subcortical control of visual thresholds in humans: Evidence for modality specific and retinotopically organized mechanisms of selective attention. *Experimental Brain Research, 29*, 173–190.
- Smith, M. A. (1997). *Simulating multiplicative neural processes in non-orthogonal coordinate systems: a 3-D tensor model of the VOR*. Master of Arts thesis, Graduate program in Psychology, York University, North York, Ontario, Canada.
- Sparks, D. L. (1999) Conceptual issues related to the role of the superior colliculus in the control of gaze. *Current Opinion in Neurobiology, 9*, 698–707.
- Stein, B. E., & Meredith, M. A. C. (1993). *The merging of the senses*. Cambridge, MA: The MIT Press.
- Svarverud, E., Gilson, S. J., & Glennerster, A. (2010). Cue combination for 3D location judgements. *Journal of Vision, 10*, 1–13. <http://journalofvision.org/10/1/5/>, doi:10.1167/10.1.5.
- Thaler, L., & Goodale, M. A. (2010). Beyond distance and direction: The brain represents target locations non-metrically. *Journal of Vision, 10*, 1–27. Doi:10.1167/10.3.3.
- Trehub, A. (1991). *The cognitive brain*. Cambridge, MA: The MIT Press.
- Trevarthen, C. (1968). Two mechanisms of vision in primates. *Psychologische Forschung, 31*, 229–337.
- Tsang, E. P. K. (1993). *Foundations of Constraint Satisfaction*. London: Academic Press. Downloadable at: <http://cswww.essex.ac.uk/Research/CSP/edward/FCS.html#FCS> Download
- Turatto, M., Bettella, S., Umiltà, C., & Bridgeman, B. (2003). Perceptual conditions necessary to induce change blindness. *Visual Cognition, 10*, 233–255.
- Velmans, M. (2008). Reflexive monism. *Journal of Consciousness Studies, 15*, 5–50.
- van Gaal, S., Ridderinkhof, K. R., Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2008). Frontal cortex mediates unconsciously triggered inhibitory control. *Journal of Neuroscience, 28*, 8053-8062.
- von Holst, E., & Mittelstaedt, H. (1950). Das Reafferenzprinzip (Wechselwirkungen zwischen Zentralnervensystem und Peripherie). *Naturwissenschaften, 37*, 464–476.
- Watson, A. B. (1987). Efficiency of a model human image code. *Journal of the Optical*

Society of America A, 4, 2401–2417.

Webb, B. (2004). Neural mechanisms for prediction: Do insects have forward models? *Trends in Neurosciences*, 27, 278–282.

Witkin, A. P. (1981). Recovering surface shape and orientation from texture. *Artificial Intelligence*, 17, 17–45.

Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, 269, 1880–1882.

Zettel, J. L., Holbeche, A., McIlroy, W. E., & Maki, B. E. (2005). Redirection of gaze and switching of attention during rapid stepping reactions evoked by unpredictable postural perturbation. *Journal Experimental Brain Research*, 165, 392–401.