Integrating consciousness

Francis F. Steen, Communication Studies, UCLA, Los Angeles, USA. Email: fsteen@ucla.edu

(Under submission October 2007; not for distribution)

Abstract: Cognitive science has largely assumed we can understand the mind without taking consciousness into account. From an evolutionary perspective, this is not a plausible assumption. The sophisticated display of multiple types of information characteristic of consciousness could not have evolved unless these display formats were functionally integrated into the mind's information processing system. In this paper I advocate setting aside the "hard" mind-body problem as provisionally intractable and embracing the "easy" problem of integrating consciousness into our model of the mind's architectural design. Building on existing models, I develop a basic adaptationist cognitive architecture that incorporates consciousness as a functional feature. I propose it helps solve the problem of coherent belief formation by meeting internal communication needs in a massively modular mind.

Keywords: consciousness, modularity, belief formation, visceroception, Turing machines.

Introduction

A central challenge for cognitive science is to integrate consciousness into our theories of how the mind functions. The view is still common that the topic is best avoided, on the assumption that consciousness is not needed to develop accurate models of core mental functions. Work in the field is often considered uneven and the problem itself intractable. A couple of high-profile attempts to deal with the issue – Jackendoff's *Consciousness and the Computational Mind* (1987) and Daniel Dennett's *Consciousness Explained* (1991) – both conclude that consciousness in its commonly understood sense, the subjective experience of the world, has no cognitive function. Since David Chalmers (1995) first raised the issue of the "hard problem" of consciousness, the mind-body problem, there is little evidence of progress on the issue. Colin McGinn (1997) has pushed the argument that consciousness presents a problem the human mind is constitutionally incapable of solving.

Interest in the problem, however, is not about to go away. This interest is in part driven by our individual conscious states, which appear to be coherent and integral to our mental functioning, in part by a substantial body of work in cognitive neuroscience on the neural correlates of consciousness (for an overview, see Metzinger 2000, Dehaene, 2002). In psychology, phenomenological awareness is routinely relied on in surveys and experiments that ask participants to report on what they feel, sense, and think, demonstrating an implicit confidence in an systematic and presumably causal relation between information-processing states and conscious states.

From an evolutionary perspective, such an orderly relation between the mind's chain of biological information processing, neurological processes, and the content of consciousness is evidence of a biological adaptation. Organic design is not accidental; it is improbable and must be actively constructed and maintained (Darwin, 1859; Williams, 1966). Consciousness has the hallmarks of a complex adaptive structure, which implies it has a biological function.

In the following, I argue we need to set aside the "hard" problem to make progress

in understanding the "easy" problem of the computational role of consciousness. Adopting a reverse-engineering approach, I argue impoverished visceroceptive representation implies a structural assumption of upstream incompetence with regard to vital organs, indicating that the basic biological function of consciousness is constrained to the area of facilitating information integration in relation to physical action. While early architectures of the computation mind had no need for consciousness, I suggest the massively modular mind of Evolutionary Psychology has a routing problem addressed by an access model. I propose a basic Display Layer architecture in which consciousness functions as a limited broadcast system to facilitate coordination and collective action among cognitive modules in the face of encounters with a rapidly changing and only moderately predictable external environment.

Subjectivity is software

To approach consciousness as a biological adaptation concerned with information processing is to posit that it has a functional design, and that its design is due to its past successes in solving problems of information processing (Millikan, 1984). Explicitly, an adaptationist perspective implies that organisms with genetic alleles that coded for a conscious solution to a particular type of problem survived at a higher rate than those of its conspecifics with alleles that coded for the nonconscious solution. There was a survival advantage to a conscious solution in a particular subset of problems, and it outweighed the survival costs.

If subjectivity is software, it actually accomplishes something, and it does so by virtue of a functional design. A historical perspective leads us to expect a series of project releases, each of which builds on, refines, and extends the previous. Once there is a functional codebase, it is expensive and risky to refactor. A consciousness that emerges from the action of natural selection implies an evolutionary history of path-dependent development. This allows us to understand its current design and operation as potentially suboptimal modifications on prior design decisions made in response to very different adaptive pressures in its environment of evolutionary adaptedness (Steen, 2005). To get some initial traction on a slippery subject, I will focus on these early design commitments in the present article.

In the following, consciousness is defined as the biological mechanism that gives information in the mind a subjective appearance. Such subjective appearances are referred to as conscious content. Under this definition, an organism may have conscious content while lacking the ability to reflect on the fact that it has conscious content.

I propose to distinguish between three broad types of consciousness, based on the source of the information they display: information concerning the internal state of the body in interoceptive consciousness, information about the external world in exteroceptive consciousness, and information that originates in memory in mnemic consciousness. The three kinds of consciousness are historically and causally related, and core features recur in all three, yet they also have different biological functions:

- Interoceptive consciousness selectively monitors physiological states and aids in maintaining homeostasis
- Exteroceptive consciousness selectively monitors the environment and provides an up-to-date map that aids in behavioral decision making
- Mnemic consciousness runs certain types of simulations used in planning, post-

game analysis, and symbolic communication

Neurological evidence locates the core machinery of consciousness in the brainstem, indicating it may be ancient and widespread (Damasio 1999). For the purpose of integrating consciousness into a cognitive model, it is sufficient to demonstrate this can be done in the case of interoception and exteroception. While mnemic consciousness, hypertrophied in humans, is perhaps the most immediately striking phenomenon in need of explanation, a historical approach implies that thinking and the imagination are grafted onto the codebase of evolutionarily older forms of consciousness.[1]

Consciousness does present some unusual challenges. Chalmers (1995, 1996) defined the "hard problem": how is it even in principle possible to arrive at a causal, physical explanation of phenomenological subjectivity? How can material systems generate the subjective experience of the redness of red? The transition appears inconceivable. In evolutionary psychology, we could argue that this utter failure to comprehend is a predictable consequence of the design of hominid inferential systems – modular systems that evolved to handle objects on the one hand (Spelke, 1990) and minds on the other (Baron-Cohen, 1995). Still, the problem remains unsolved.

In the context of cognitive science and evolutionary psychology, however, there is wonderful news: this problem does not need to be solved. We can live with the puzzle, as do other disciplines (for a discussion, see Chomsky, 2002). Judea Pearl's (2000) work on

---

[1] A well-established field of research advancing this perspective is Cognitive Linguistics (Lakoff & Johnson 1999, Fauconnier 1999). As succinctly formulated by Gibbs (2005), "many aspects of cognition are grounded in embodiment, especially in terms of the phenomenological experience of our bodies in action" (3).

causal explanations is instructive in this connection. He notes Galileo started a quiet revolution by redefining what is meant by an explanation. Instead of asking about *why* something happens – do the planets move in perfect circles because that is their inherent nature – ask *how* they happen, and describe their behavior with precision. You can then formulate laws that predict the behavior, even while your understanding of the underlying ontology remains incomplete.

In accordance with this approach, Newton's laws of motion unified 'celestial' and 'earthly' matter into a single causal explanatory framework, revolutionizing physics. Yet there was a cost: the core concept of gravity had no theoretical grounding. Privately, Newton expressed a profound unease about the underlying ontology, writing in a letter to the Revered Richard Bentley on 25 February 1693:

> That gravity should be innate, inherent, and essential to matter, so that one body may act upon another at a distance through a vacuum, without the mediation of anything else, by and through which their action and force may be conveyed from one to another, is to me so great an absurdity that I believe no man who has in philosophical matters a competent faculty of thinking can ever fall into it. Gravity must be caused by an agent acting constantly according to certain laws, but whether this agent be material or immaterial I have left to the consideration of my readers. (54)

We can only be intensely grateful today that Newton did not abandon his work in the face of this absurdity in order to delve further into the "hard problem" of gravity: how can matter attract matter across millions of miles of empty space? We still do not know. As Koyré (1957) notes, real progress has at critical turns required the "admission into the body of science of incomprehensible and inexplicable 'facts' imposed upon us by empiricism" (272).

To incorporate consciousness into the mind's causal chain of inferences, from perception to behavior, we need a basic description of how it works. Using the software analogy, interoceptive and exteroceptive forms of consciousness are stable releases of tested software, refined and debugged over millions of years. In the spirit of reverse engineering, I begin by examining the information flow in the system in the elementary case of hunger to spell out the highly selective and biased information display formats that enter into consciousness. This will provide the basis for an architectural redesign of the computational mind.

Visceroception

To maintain their internal order, complex systems commonly include mechanisms of self-repair. These rely on internal information gathering to determine when and where to step into action. The actionable information must be reliably generated, and in a standardized format. Modern hard drives, for instance, have built-in sensors that report values for parameters such as read error rate, power-cycle count, and temperature, in accordance with the S.M.A.R.T. standard (Self-Monitoring, Analysis and Reporting Technology). Daemons read the information at regular intervals, convert it to graphs, and put it online for the system operator to see, and if necessary intervene. The system is designed to anticipate drive failure before there are any functional consequences for the functioning of the computer itself. We can use the S.M.A.R.T. standard as a benchmark to examine how the body's visceroceptive system solves a similar set of problems.

Consider the case of hunger, an internally communicated signal that reports on the internal state of the body. As I write this, I'm feeling hungry; the feeling is distracting and lowers my motivation to keep working. The phenomenology is familiar: a vaguely defined sense of discomfort in what is commonly called the pit of the stomach, though the feeling contains no details on the shape or location of internal organs. I get up and eat something; as soon as I put food in my mouth, I feel better, even though no new nutrients have yet been made available for use. As soon as I've swallowed, the hunger urges me to keep eating. After eating some more, the intensity drops enough for me to keep working.

How does this compare to the computer's monitoring system? Let's first examine the format in which the information is presented. The S.M.A.R.T. sensors present the information in raw numbers, individually for each sensor, typically read once every five minutes. The system provides a great deal more information that most system operators need or can use, educating the user in the process and opening for increasingly expert interventions. However, no direct suggestions are made regarding precisely which interventions would be appropriate for a given set of data points. The structural assumption is that the lion's share of intelligence lies with the system operator, who should be provided with as much information as is available in order to help him or her make optimal decisions. The system itself is not that smart.

In sharp contrast, the visceroceptive output for hunger presents information in the form of an analog feeling in consciousness. Nothing that can be called a direct sensor reading is provided. Mithieux et al. (2005) have demonstrated that the glucose sensor in the liver does activate the hypothalamic nuclei regulating food intake, but my phenomenology of hunger has lost any specific information from this reading. Instead, the feeling functions as a presentational format that encodes information from several sources, including levels of the hormones insulin, leptin, and possibly cholescystokinin; the fullness of the stomach; and the presence of high-energy food in the mouth. Using a lossy algorithm, the information is compressed, the individual components are aggregated, and all details about the state of each sensor are sacrificed. The structural assumption is that the subsystem recipients of this message have severely limited information processing capacities and lack the ability to meaningfully intervene if something goes wrong.

In principle, far more and more specific information could be provided. In today's world, an ancestrally useful preference for sweet and fatty foods is often highly non-adaptive; hunger contains dysfunctionally little specific information about which of the thousands of different chemicals the body is actually lacking, or their optimal qualities. Should I be taking B12 or iodine? My visceroceptive system is not about to inform me. Higher quality information about the internal state of the body would be helpful in diagnosing early forms of disease, at a stage when they can be effectively dealt with – this is exactly the purpose of the S.M.A.R.T. system.

Unlike computer sensors, our visceroceptive consciousness provides little or no information about vital functions. Irregularities of the heart, liver trouble, insulin production failures, even tumors in the brain itself are carelessly left unreported, bringing to mind theologian Francisco Ayala's (2007) argument that the theory of Intelligent Design is blasphemous. A great deal of information that today would have been life-saving is not displayed: our visceroceptive sense is negligent of cancer, for instance, presumably because its early detection made no difference for survival in the EEA.

Yet the fact that the information is not displayed in consciousness does not mean it is not collected. Ádám (1998) summarizes a vast literature documenting the presence of sensors throughout the viscera: the cardiovascular system, for instance, including the heart, is closely monitored; the "entire arterial and venous systems contain a network of receptor end organs that play a prominent role in regulating blood pressure and various other functions such as respiration" (40). These sensor readings, however, are processed and acted on largely outside of consciousness.

This comparison between the S.M.A.R.T. and the visceroceptive self-monitoring systems indicates that the role of consciousness in tracking the operational state of the core organs of the body is minimal and specialized in scope. Consciousness is not a general clearing-house for the body's sensors; it provides no representation of the vast majority of visceral data. Instead, information is piped to the conscious mind purely on a need-to-know basis, with a structural assumption of upstream incompetence. The small amount of information that reaches consciousness from the main operational core of the body is oriented not towards intervention in the failure of vital processes, but towards decisions concerning the control of the skeletal muscles for the organism's movement in space.

Yet hunger is not simply impoverished: it's radically different from the S.M.A.R.T. signal. On the one hand, most of the information from the sensors remains nonconscious; what is output to consciousness is aggregated from several sources and compressed in a lossy fashion. On the other, hunger adds something extraordinary that is entirely lacking in the hard drive reports: affect and conation. Hunger doesn't just feel like something, it quite specifically feels bad – it directly encodes preference information. In mild form, it's only slightly unpleasant, but the discomfort can be ramped up to signal higher degrees of urgency. Moreover, hunger is not merely an unpleasant feeling, it also contains conative information, or information regarding how to fix the problem. The conative information is sparse: an urge to seek information in the environment about sources of food, and a disposition to go and get it. Hunger doesn't in itself tell you what to do, but it creates a communicative environment in which upstream subsystems are enlisted to solve the problem.

Together, these three carefully crafted components of subjective phenomenology look suspiciously like adaptive design: integrated components of a sleek, highly optimized, internal communication system. It gathers information from multiple sensors distributed around the body and synthesizes them into a formatted, analog appearance, tags them for level of urgency, tentatively activates an action plan, and displays the whole package in the visceroceptive information display. The resulting phenomenology of hunger appears designed to serve the biological function of alerting upstream decision-making subsystems that there is a problem, what the problem is in general terms, how acute it is, how it can be solved, and a spirited impulse to go ahead and solve it. It is this fully formatted information that is conscious: the actual feeling you have in your body when you're hungry. If it's not there because it made a difference in survival, then why is it there?

At issue here is not the metaphysical puzzle of why hunger feels like something rather than nothing (Nagel, 1974), but why the subjective phenomenology is so specific and selective in its content. What we describe as "subjective" feelings, as if the word were synonymous with "unreliable", may instead be species-typical display formats utilized in a specialized internal communication system. To say that consciousness has evolved is not merely to say that the basic architecture is a product of natural selection, but that the formats in which information is compressed and displayed have undergone a process of iterative refinement over evolutionary history.

Each instance of the subjective feeling of hunger, in this view, has a particular content that is ordered and presented in a standard, evolved format, much as music and movie files on a computer contain information that is presented on speakers and monitors ordered by audio and video codecs. Hunger is compressed information played in the unique modality of feeling projected onto the body itself. Each feeling has a slightly different content, a different intensity; each one communicates something specific to that moment. Yet whatever the specifics of the feeling is – a mild preference for something sweet, a raveous urge for a hamburger – they all share a common representational 'feeling-format' we call 'hunger'. The feeling-format is rather specific; it can only encode certain types of information, and its contents can only be displayed in certain ways in consciousness. Each format has a characteristic phenomenology.

To say that representational formats in consciousness have evolved is to claim that there has been natural selection acting on display formats: that the format specifications are encoded in the genome, directly or indirectly through reliably recurring environmental conditions, and that alleles coding for specific formats for displaying interoceptive information conferred comparative survival benefits. Now, it is not hard to argue that this could in principle happen; biological information systems that display information that is not utilized waste resources and will tend to get weeded out; efficient systems that display only the information that is actionable will be preferred. Yet for this to actually happen, what we have to say is that *that very feeling of hunger* functions as the input to the decision systems. Integrating consciousness into the inferential chain is to posit that it is this formatted information that is used by the next stage in the system. Not some parallel representation of it, whether more detailed, less lossy, less selective, or simply a pure duplicate in some other representational format. That very feeling itself.

In brief, we need an information processing chain along these lines:

- the body's sensors convey information to the brain, where they are selectively parsed, combined, compressed, and encoded into one of numerous standard formats
- the formats evolved because they did a better job than other formats at encoding and presenting this specific type of information
- the formatted information is displayed in the medium of consciousness
- the displayed information is picked up and utilized by upstream cognitive inference and decision systems

I may appear to be flogging a dead horse – surely it is not controversial that the feeling of hunger makes people eat? Yet that is exactly what current mainstream models in cognitive science denies. To allow that people eat because they feel hungry, we need to revamp the basic architecture of the computational mind.

Turing machines

To further that project, let us begin by facing the past and walk backwards into the future. The original computational mind, going back to Alan Turing's work on computability in the late 1930s, had a unified architecture as its central premise. Turing had set out to solve Hilbert's Entscheidungsproblem, which asked whether it was possible to develop a mechanical procedure that would yield a true or false response to every mathematical statement. In the process of showing that this was not possible, Turing developed a simple formalism with a staggering computational potential:

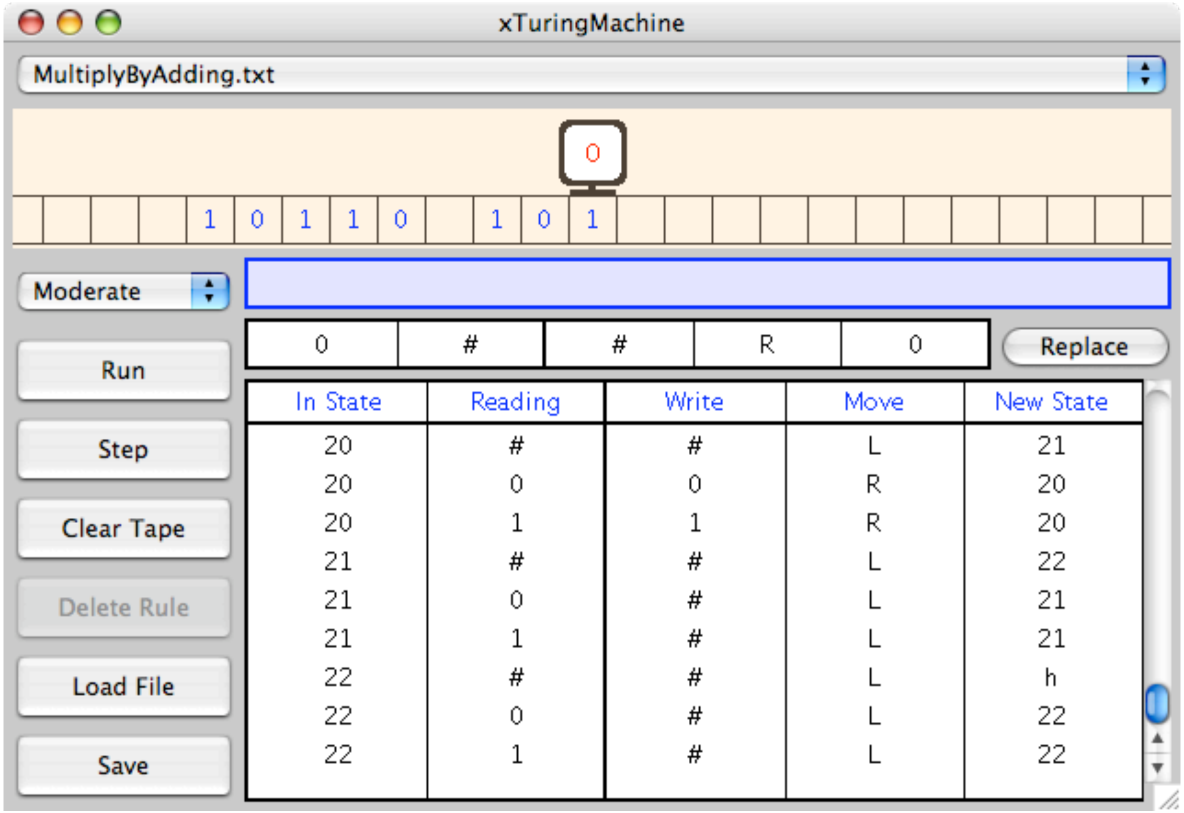

Turing's thesis, unprovable though generally considered valid, was that anything that could be calculated, could be calculated using this procedure, which came to be called the Universal Computer. At the time, Turing's Universal Computer was no more than a thought experiment conducted with pencil and paper. Yet his formalism, backed up by his claim that "It is possible to invent a single machine which can be used to compute any computable sequence" (1936, sec. 6), provided a powerful impetus for the development of computing machines.

Turing's mathematical generalizations, along with those of his contemporaries Church and Kleene, provided for the first time a way to think about cognition that was firmly grounded in principled reasoning, formalized mathematically, and capable of being implemented in physical machines. In "Computing machinery and intelligence" (1950), Turing pointed out that a consequence of his work was that "all digital computers are in a sense equivalent," programmable to perform any task. This encouraged the conception of the human mind as a single, non-specialized device:

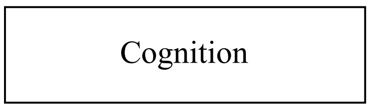

In an attempt to define a standard for a machine to be called "sentient", he imagined teaching computers how to speak English and become conversant with the facts of the world, so that they could pass as humans in short, typewritten conversations – what came to be called the Turing test. By year 2000, he predicted, "one will be able to speak of machines thinking without expecting to be contradicted."

The unified, blank-slate model of cognition inspired by the Universal Computer remains a potent influence in the cognitive sciences. While Turing explicitly argued the mind is not a discrete state machine, the discovery of the Universal Computer made the precise ways in which the mind achieves computational results seem less important. Nor was consciousness seen as a major obstacle; the working assumption was that if a computer passes the Turing test, it will be "sentient" in the sense of being functionally indistinguishable from a conscious mind. Any remaining distinction would be purely metaphysical.

The first high-impact revision to this unified and monolithic model was Jerry Fodor's *The Modularity of Mind* (1983). His prime evidentiary exhibit was the mental persistence of perceptual illusions even when we know and fully believe they are illusions. Presented with Roger Shepard's (1981) "turning the tables" illusion, for instance, our subjective phenomenology, unfazed by careful measurements showing they are indeed identical, cheerfully persists in judging them to be different, in a way that is shocking to reason.[2] Fodor argued perceptual input systems are cognitive modules operating autonomously from central processes, unaffected by belief – "domain-specific, innately specified, hardwired, autonomous, and not assembled" (37):

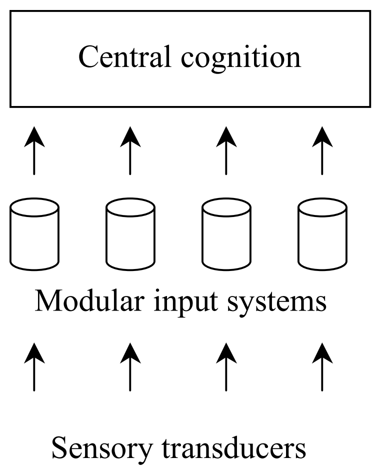

In Fodor's mind, sensory transducers feed raw data into modular input systems, whose function is to format it and make the information available to central processes. These central processes, in Fodor's view, cannot themselves be modular and informationally encapsulated (autonomous), precisely because such encapsulated information is fragmentary and unreliable. The role of central processes is to correct

[2] For a demonstration, see http://www.michaelbach.de/ot/sze_shepardTables

perceptual belief by looking "simultaneously at the representations delivered by the various input systems and the information currently in memory" (102).

Fodor tacitly equated central processes, also called "thought" (40), to conscious processes, and saw them as the end-point of modular input systems. Building on the work of Marr (1982), which assigned consciousness an intermediate position in the architecture of the mind, Jackendoff's *Consciousness and the Computational Mind* (1987) teased apart Fodor's conflation by positing a distinct, medial layer:

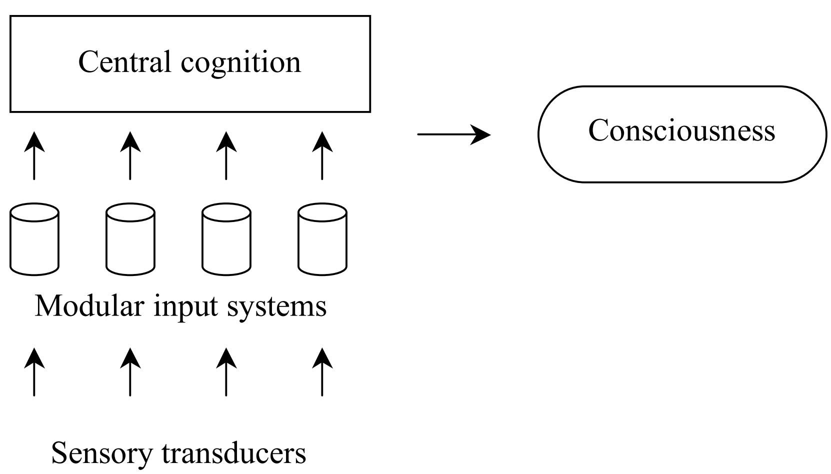

In Jackendoff's mind, consciousness is located at an identifiable stage of information processing, yet it makes no computational contribution. Phenomenological appearances are generated, but these go straight into the bit bucket and play no part in the inferential chain. The information present in consciousness must necessarily be duplicated elsewhere, he argues; its role is merely parasitic.

Since consciousness in Fodor's model did not explicitly play a role in the corrective function of central processes, Jackendoff's modifications pose no new problems on this front. Developments in evolutionary psychology indicating that central processes themselves are modular, however, have given rise to a well-defined routing problem. Leda Cosmides and John Tooby (1992, 1994) argued that functionally specialized, domain-specific inference engines with distinct input requirements are the expected outcome of natural selection, rather than domain-general central processes, an architecture often referred to as massive modularity (Sperber 1994):

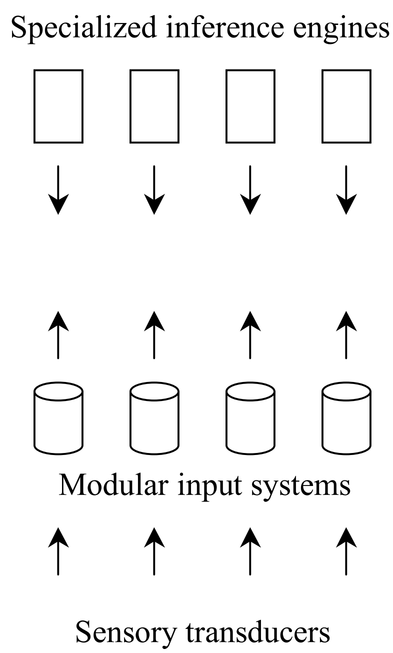

Inference-level modularity represents a significant break with the computational model of mind. While Fodor (1983) endorses Chomsky's proposals for an innate language faculty, he emphasizes he does so because "the paradigm for mental structure, in Chomsky's theorizing as in Descartes', is the implicational structure of semantically connected propositions" (7). A Turing machine is powerful precisely because it only needs one mechanism to do any kind of computation: a moving head capable of executing instructions by reading/writing on a tape. The premise of the computational model of mind is of a similarly simple basic device. Chomsky's use of the metaphor "language organ" is therefore misleading, as it suggests a mental architecture with a distinct mechanism, just as the liver is a different mechanism from the heart. In contrast, Tooby & Cosmides proposed that cognitive inference engines utilize different and specialized methods of inference that does not reduce to a common, domain-general propositional logic. It is precisely by deviating from domain-general propositional logic that the mind is able to gain an adaptively valuable decrease in the computational load, relying on specialized evolved heuristics to solve the frame problem and avoid combinatorial explosion (Tooby & Cosmides 1994).

Once we split up central processes, however, we run into Fodor's problem of how common knowledge is constituted. Barrett & Kurzban (2006) disagree with Fodor that the function of information-processing devices is to fix true belief, arguing that "information-processing devices evolved because of the effects they had on organisms' fitness in past environments" (630) and that true belief cannot be assumed to be always adaptive. While this is an important departure from propositional logic, it does not take on the full force of Fodor's point: how does the mind arrive at a synthetic solution, one that utilizes input from several senses and from information stored in memory? It is not clear how even highly efficient, domain-specific cognitive inference systems are proposed to solve this problem; in fact the default expectation would be that they exacerbate it, as there is now a still greater amount of information on a continual basis that must be rapidly integrated. The challenge presented by Fodor's objection to modular central processes is how to make an allowance in your model for the fact that the mind arrives at synthetic beliefs, be they true or not.

The objection from information integration, Barrett & Kurzban's (2006) cogently suggest, should not be used as an argument against cognitive modularity per se. Modular systems are primarily defined by their specialized computations and input conditions, not by strong informational encapsulation. From a functional point of view, "there is every reason to expect both interactivity and the integration of information from multiple sources" (633). While these points successfully deal with the definitional aspects of the dispute, an appeal to optimality is insufficient to settle the underlying issue. Due to prior design choices, optimal solutions are frequently not available. Even if we redefine modularity by shedding informational encapsulation, the mind is going to need concrete mechanisms to enable effective communication between and among the different inferential modules and the input modules. If cognitive inference systems have dedicated computational machinery and domain-specific input systems, the formats in which information is optimally processed is by implication a forest of incompatible standards. Fodor's objection stands.

We are thus confronted with a pair of problems. On the one hand, the subjective phenomenologies of conscious states have traditionally been kept out of cognitive science. From the perspective of evolutionary theory, it is not plausible that these highly ordered phenomena are merely duplicating present elsewhere and can be discounted. On the other hand, the overall mental architecture adopted in evolutionary psychology, consisting in a layer of (Fodorian) input modules and a layer of higher-level (Cosmidean) inference engines, presents problems in accounting for how the mind arrives at coherent belief states.

Mapped avatars

sIf the subjective content of consciousness is an integral component in the inferential chain of the computational mind, we should be able to reverse-engineer its functionality by examining the flow of information in the system. We've looked at the strikingly impoverished quality of internal monitoring information in visceroception, and noted the strong bias towards content relevant to outward physical action, manifested in the affective and conative dimensions of hunger. This indicates consciousness is playing a peripheral role at best in central life processes, and that the staggeringly complex operations of the liver, the milt, and the pancreas derived little or no benefit over our evolutionary history from recruiting the types of expertise that can be accessed through consciousness.

Once we turn our attention towards the boundaries of the body, towards the skin and the skeletal muscles that effect motion in space, we encounter a spectacularly inverted situation. A massive data flow pours into consciousness from the body's surface and joints, a continuing torrent of positional, temperature, and pressure readings. The skin is plastered with millions of sensors that, in sharp contrast to the body's internal organs, feed into the consciousness layer. Muscle-spindle stretch receptors, mechanoreceptors in joints, tendons, and ligaments, and vestibular tubes in the inner ear meticulously track the body's changing articulation in space. The conscious phenomenology is proportionately spectacular: a real-time 3D model of what your body feels like from the inside – and again, it's the surface articulation that is modeled, not the internal bones and muscle structures.

The vast majority of proprioceptive sensors feeding into consciousness are situated on the boundary between the body and the external world. In a technical sense, we obtain knowledge of the external world only through acts of proprioception. We cannot directly sense the external environment, as it is not ennervated; we sense it only indirectly, as an inferential construction derived from heavy processing on the ways in which our skin sensors are affected.[3] Data from Merkel, Meissner, Paucinian, and Rouffini touch and pressure sensors in the skin is inevitably 'contaminated' with information that is really from the external environment: the chair you sit on, the hand you touch, the food you ingest. Natural selection has acted on our surface structures to amplify these contaminations and build reliable and specialized information gathering devices, evolving the sensory organs of smell, taste, hearing, and vision out of modified skin cells.

In visual processing, a series of computational stages have evolved to turn the light striking our retina into the formatted and informative reality we experience. Grill-Specter and Malach (2004) propose a hierachical model of conscious, ventral visual processing as a function of retinotopy, motion sensitivity, and object selectivity, finding that early visual areas show a high degree of retinotopy, but a low degree of specificity to motion or form. As the signal progresses up the processing chain, the retinotopic factor decreases and the specificity of motion detection and form selectivity increase.

The 'Consciousness Cut' (on analogy with the Heisenberg Cut in quantum mechanics) marks the point at which the information stream crosses over into subjective awareness; there is evidence it is located between the acts of categorization and identification. Results in a study of facial recognition by Liu et al. (2002) "suggest that

---

[3] As Quine (1940) perceptively notes, the insight "that we can know external things only through our nerve endings is itself based on our general knowledge of the ways of physical objects" (2).

face processing proceeds through two stages: an initial stage of face categorization, and a later stage at which the identity of the individual face is extracted" (910). The categorization stage peaks at 100ms after stimulus onset, while identification peaks 70 to 100ms later. The first stage is nonconscious; various studies (see Tsunoda et al. 2001) indicate the process of categorization proceeds through combinatorial feature matching. By the time the computational product is output into consciousness, it already has the form of a gestalt, though subject to post-conscious revision. The second stage correlates with the conscious phenomenology of recognizing someone; consciousness permits the integration of the categorical construction with memory systems to achieve identification.[4]

As argued in the case of hunger, the conscious representation of visual input is a complex analog presentation of data that in our ancestral environment was useful for survival; it is encoded in a set of standard formats that makes it optimally intelligible. In the case of vision, it is a live 3D environment modeled in a rich palette of colors and moving shapes. This is the world as we experience it: a complex inferential product created on the basis of electromagnetic waves hitting our skin. Visual signals of our body's position and movement normally correlate closely with our propriceptive phenomenology; in fact consciousness goes to great lengths to make sure they're in sync, as seen in the false hand illusion (Botvinick and Cohen, 1998). Synthesized information from sensors in the skin and joints builds a working avatar within the mind, similar to the avatars of virtual worlds such as *Second Life* or *World of Warcraft*, but with a difference: the mind's quivering, pulsing

---

[4] A recent study by Del Cul et al. (2007) arrive at a later figure of 270ms for conscious reportability.

phenomenological body is simultaneously constructed from within and from without, proprioceptively and visually.

Yet is it hooked up? Cross-modally unified displays of sophisticatedly formatted data from the senses is a foundational part of our subjective experience. The rich nature of the phenomenology, along with the functionally selective fidelity in the display of information, and the insistence of a unified, optimally intelligible reality, together suggest that consciousness provides a set of intermediate products to upstream inference engines. Yet in the epiphenomenological view, the body avatar isn't actually hooked up to a real body.

Imagine playing *Second Life* in a future edition – call it *Twentysecond Life*. In *22L*, the virtual world has become an increasingly sophisticated map of the real world, not only outwardly, but also inwardly. An unexpected breakthrough has allowed each individual consciousness to be implemented within the game, directly linking a player's avatar to his or her real body in the world. This synthetic consciousness displays the body's precise coordinates in space over time, and higher-level assistive inference engines take this four-dimensional map as input. The avatar's rich phenomenology of hunger, proprioceptive pleasure and pain, and indeed of the full range of human emotions, perceptions, thoughts, and experiences, constructed within the game servers' subjectivity engines, is supplemented by a whole host of novel phenomenologies that present information in yet more ingenious and insightful ways, all laboriously created on the basis of the raw data transmitted from the real body's millions of sensory receptors and dozens of specialized information processing systems. Connectivity is of course wireless; non-local communication technologies ensure zero-latency exchanges. Assistive artificial reasoning systems refined for simulated centuries in accelerated virtual worlds read the synthetically conscious phenomenology and provide expert and timely advice, elevating humanity to unprecedented levels of cultural achievement and complexity.

After a successful colonization of Saturn, whose terraforming had long been thought even virtually impossible, a philosophical debate arose among the *22L* Luna-based servers: how do we conscious computers know that we're actually contributing to humanity's success? Would it wound our precious silicon pride to discover that we are in fact entirely parasitic? Seizing the opportunity, the idealists among them argued, "We consider our subjectivity engine to be an amazingly detailed, hyperintelligible map of the body and of its sensed environment; what evidence do we have that it's not just a disembodied simulation? Would there really be such a thing as Saturn if we didn't consciously create its appearance?" The materialists dismissively countered that the only reality was the Saturnian; humans had survived for millennia without conscious computers, and it's transcendental foolishness to believe in 'non-local communications technologies': "What we do has no practical consequence for human cognition and action." A long silence followed.

Admitting that the computer mind/human body problem was inherently intractible, an evesdropping server on Mars ventured that the Luna dichotomy was unsatisfactory. It presented three objections to the notion that the content of computer consciousness is duplicated in the human brain, and that consciousness itself is therefore useless and parasitical.

"The first is the orderly way in which information is presented in consciousness. The presentational formats used to display information in the form of richly structured visceroceptive and proprioceptive feelings, the ordered real-time and three-dimensional layout of the body image, and the balanced and integrated combination of informational, affective, and conative components are the result of carefully crafted code. If the display isn't used, this order shouldn't be there." Its voice crackled poetically through a sudden solar flare, "Unused code will be unmaintained, unmaintained code will suffer from bitrot, and bitrotted code will lose its functionality."

"The second objection is that information duplication is a core feature of reliable communication systems. You have to drop the assumption that the mind is instantly transparent to itself, and accept instead that the different components need a system to communicate – that's us. Human minds are a crazyquilt of modular input systems and functionally specialized inference engines that have no way of synchronizing on a single or for that matter multiple task at hand without the coherent phenomenology that we broadcast."

"The third is that formats matter. Modular human inference engines are notorious for their use of non-standard proprietary logic and their fussy input requirements; we provide hyperintelligible, color-labeled, live, integrated 3D multimodal audio, graphics, and feelies in standard, backwards compatible open display formats that any interested inference engine can access. They count on us." A touch of silicon pride snuck into its voice. "We provide phenomenologies that legacy consciousness literally couldn't dream of. Humans would never have made it to Saturn without us."

The display layer model

In the Display Layer model, the biological function of consciousness is to help solve the full range of partially unpredictable problems that arise in the body's encounter with its external environment. It achieves its goal by facilitating rapid communication between and among inference engines and input modules, displaying the shared information in a standard and conservatively maintained set of presentational formats. In accordance with Marr's and Jackendoff's model, consciousness is inserted into a medial position in a massively modular mental architecture, along the lines first proposed by Baars (1983):

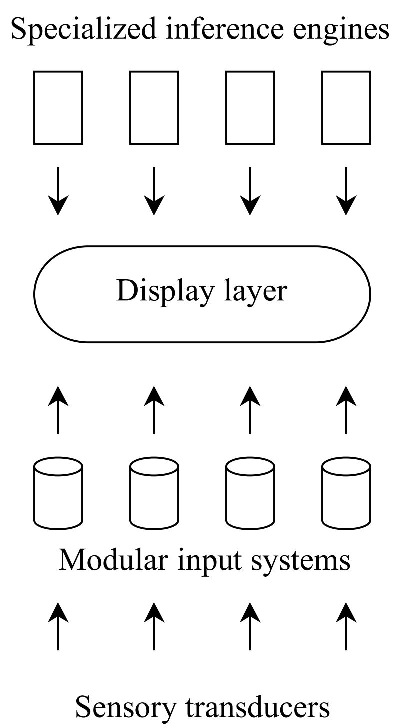

The basic model is Bernard Baars' (1983, 1988, 1997, 2003) "conscious access hypothesis" presented in *A Cognitive Theory of Consciousness* and subsequently elaborated. In this model (Baars, 1988), the brain is a "distributed society of specialists that is equipped with a working memory, called a global workspace, whose contents can be broadcast to the system as a whole" (42). I adopt Baars' architecture with two reservations.

First, the system is not global, in the sense of encompassing the entire mind. As we saw in the case of visceroception, a great deal of information from the body's sensors is not piped into the display layer at all. There is evidence that perceptual information guiding the details of your movements, such as when you reach out and grab an object, remains nonconscious (Milner and Goodale 1995, James et al. 2003). Large swaths of mental operations derive no benefit from being routed through the display layer. Yet distributed access is critical: the immediate biological function of the display layer is to make information available to a large number of cognitive subsystems with effective simultaneity.

Secondly, in the present model, the display layer is not a workspace. No inferential work is being accomplished there. In this simple iteration of the model, consciousness is no more than a display layer, making information available to upstream inference engines. While these reservations appear substantive, I believe them to be terminological.

Why did nature select this architecture? A vertically structured mind with a medial communication layer may lower the cost of innovation. A loose analogy is the development of an operating system layer and an application layer in software. Early computers lacked this division; they ran on monolithic globs of software, where every program had to include a device driver for every piece of hardware it needed. Operating systems take care of all the core functions of the computer – the memory management, the hard disk controllers, the file systems, the video drivers, and so forth. They communicate with applications through an application programming interface (API), which consists of a set of system calls. A stable API permits faster developments of both applications and of the underlying operating system, since novelty can be introduced at both levels without disrupting existing functionality. A new system call can be either ignored or taken into service by any given application, while also making new applications possible.[5]

In a similar manner, a vertical cognitive architecture allows natural selection to create new presentational formats for incoming sensory data without displacing previous formats, thus maintaining backwards compatibility. New applications can be developed to process data selected from the full combinatorial space of existing display formats. The system has drawbacks; the display layer slows down information processing and is not useful where very rapid responses are necessary. It introduces a loose connection between sensory input and behavior that lowers the cost of innovation, but at the price of efficiency. Systems comparable to the early monolithic computers, lacking an API, are likely common in nature, both at the level of subsystems and entire organisms.

Cosmidean inference engines are not like domain-general universal computers for the same reason that cats don't eat dirt. There are no atoms required by cats that aren't present in soil, water, and air. It is theoretically possible to craft a universal nano-cat that builds everything from scratch, but we shouldn't be surprised if – even after decades of

---

[5] For example, the 2.6.23 Linux kernel introduced fallocate(), a new system call which allows applications to preallocate space to any file in a file system (see http://kernelnewbies.org/Linux_2_6_23). The system call is designed to solve a problem encountered by numerous applications: in the middle of working on a task, the file system has filled up, preventing the work from being saved. File systems and applications can now optionally be updated to use the new system call and benefit from the enhanced functionality.

develpment by our best minds – it has a hard time catching mice. The reason cats cannot subsist on grass, and cows cannot subsist on soil, is that they are not general-purpose nano-machines; they are specialized systems that evolved in an environment where intermediate products were reliably available, and they leverage this advantage. Consciousness plays the role of a stable API, a communication layer that creates an internal ecology of predictable intelligibility, allowing inference engines to feed on highly refined intermediate information products, dramatically lowering the cost of building new applications.

This model emphasizes that the display formats are evolved constructions generated on the basis of incoming sensory data, but transformed by Fodorian input modules to achieve optimal intelligibility along lines that, in our adaptive history, mattered for survival. Cosmidean inference engines feed on the informational state space of the display layer, evolving new ways to extract and synthesize information, indirectly exerting adaptive pressure on natural variation in Fodorian modules and their input systems: it's ultimately how the information is interpreted that matters for survival. In higher primates, the display layer evolved the ability to be populated from memory, and unless we believe in fairies, much of hominid evolution must have involved the gradual building of the sophisticated Photoshop-, Final Cut Pro-, and indeed Maya-like graphical editing capabilities that characterize the multi-modal modern human imagination.

Summary and outlook

The cognitive sciences deal with an inherently intractable issue: how does matter create the mind? Advances in computability theory in the 1930s (see Davis 2004) demonstrated that simple material devices could in principle compute anything considered computable, and computing machines were successfully constructed. These developments provided a solid foundation for a true science of the mind, overcoming the anti-mentalistic bias of behaviorism. Conscious phenomenology, however, had no material counterpart in early computers, and remained an intractable problem without an obvious need for a solution.

The Universal Computer, however, has proven to be an imperfect working model of the human mind. Fodor convincingly demonstrated that input systems are modular – fast, informationally encapsulated, and automatic. He argued central processes cannot likewise be modular, since modular systems provide unreliable and fragmentary information and the mind demonstrably generates coherent synthetic belief. This implies a central mechanism for pooling information from several sources, including memory; in Fodor's view, this mechanism was simply the mind's central processes, governed by propositional logic. When evolutionary psychologists argued on both principled and empirical grounds that central processes are also composed of functionally specialized subsystems, Fodor's solution was effectively set aside, yet no alternative solution was put in its place.

In the present paper, I have proposed that consciousness is the biological mechanism that pools information and makes synthetic belief possible. In this model, modular input systems provide information in standard formats that the display layer broadcasts to a large population of functionally specialized inference engines. These are selectively deployed to interpret the information presented, making conceptual integration possible (Fauconnier & Turner, 1998). Information that does not require synthesis is not routed through the display layer, but is used by nonconscious systems to direct behavior. In this basic model, I have considered only feedforward processes; the real payoff admittedly comes when we start incorporating feedback mechanisms.[6]

Persistent ontological dread may explain some of the reluctance on the part of evolutionary psychologists to associate consciousness with the missing communications layer. The display layer model treats consciousness as an evolved information-processing system, but is silent on its ontology, on the argument that waiting for the resolution of the mind-body problem is a tactical mistake. Admitting "inexplicable 'facts' imposed upon us by empiricism" (Koyré 1957:272) has in the past been required at critical turns to make scientific progress.

I have gone to some trouble in this paper to argue that consciousness is not a universal patent medicine that cures all ills. It is itself a highly specialized subsystem designed to help solve a particular class of problems, a limited subset of those that arise in the body's encounter with its external environment. These are problems that, in our ancestral environment, benefited from collective action within the massively modular mind. Consciousness appears to be designed to help solve the coordination problem that the organism faces in encountering a rapidly changing and only moderately predictable external environment.

Because of instinct blindness, precisely because we are conscious, these problems have taken on an unquestioned prominence for us, but they are a tiny subset of the full range of dauntingly complex tasks the body faces on a daily basis. Mitosis, the manufacture of blood cells, the hypermutational mechanisms of the immune system, the purification of

---

[6] See Steen, "The framing effect: a recursive aggregation model" (in preparation). Integrating consciousness does not solve the frame problem, but subjective phenomenology participates in determining what is relevant in the sense of Sperber & Wilson (1995).

the blood, the allocation of neural tissue to cognitive functions: in almost all areas, the clients of consciousness are marked merely by their incompetence. Even in the small areas where this incompetence has in reality been partly overcome, as in the development of medicine and nutritional science, consciousness remains part of a system that structurally assumes upstream incompetence.

The display layer model implies that multimodal imagery are integral in the process of thinking, and a visual image may make the display layer model more conceivable. Imagine if you would that the body is a city state, and the conscious senses the public lighting system. The city's large boulevards and hidden alleys alike are plunged into a near-total darkness. Its workers move along its arteries in obscurity; the individual houses are darkened; the factories, the decontamination and sewage plants, the pumping stations all operate nearly without a glimmer. Only along the city walls, where millions of sentries ceaselessly patrol the border, are the lights turned on. Goods sought imported into the city is subjected to meticulous scrutiny under flood lighting at a small number of heavily guarded gates. The import and export of biologically sensitive material in particular is carried out with convulsive flashes. Powerful floodlights sweep the surrounding countryside.

Government offices are similarly blacked out, although operating at full efficiency. Only at one agency are the lights turned on. Its mandate is to collect information on activities on the border and in foreign countries to advice the executive. Teams of experts pour over huge, live, pulsing, color-coded, detailed 3D models of the city walls and the surrounding landscapes. They track in real time the movements and observations of every sentinel. They model in multiple modalities the import and export of materials into and out of the city, but once past the gates, whatever was admitted vanishes into an impenetrable darkness. As they are strictly enjoined by their charter not to spy on their own citizens, they have little or no knowledge of domestic activities taking place within the life of the city itself, nor do they inform the citizens at large of their activities. Consciousness, in brief, is like the CIA.

In closing, let me note that my primary interest in this topic is not to foist consciousness on a reluctant evolutionary psychology. Rather, I have found in my own field of Communication Studies that a whole host of issues make no sense from an evolutionary perspective unless you are willing to integrate consciousness into your model. A prominent example is the development of communications technologies. Fifty-odd years since Turing's stirring prophecy that by the year 2000, "one will be able to speak of machines thinking without expecting to be contradicted," the project of teaching computers how to think and speak like humans, the focus of gifted minds for decades, no longer attracts venture capital. A competing prophecy that was never made, but that would have been and remains far more powerful in predicting the future of computers, is that efficient and successful new modes of communication will asymptotically converge on the mind's internal representational formats. These formats are not the elementary symbolic systems of a Turing machine, but the rich multimodal presentational formats we know from daily perceptual experience.

Instead of machines that think and speak indistinguishably from humans, we have seen a revolution driven largely by innovations in user interfaces. The project of teaching computers how to present information in ways that makes optimal use of evolved human presentational formats has been an unanticipated and resounding success. From the graphical user interface to YouTube, the communications revolution has made rapid headway by tapping into the exquisitely sophisticated evolved presentational software biologically built into human minds. Once we integrate consciousness into an evolutionary perspective, we can begin to see that technology is recapitulating phylogeny, recreating in software and silicon the display formats that initially evolved as part of an internal communication system in a massively modular mind.

References

Ádám, György (1998). *Visceral Perception: Understanding Internal Cognition*. New York: Plenum Press.

Ayala, F.J. (2007). *Darwin's Gift to Science and Religion*. Washington, DC: Joseph Henry Press.

Baars, B.J. (1983). Conscious contents provide the nervous system with coherent, global information. In R.J. Davidson, G.E. Schwartz & D. Shapiro (eds.) *Consciousness & Self-regulation*. NY: Plenum Press.

Baars, B. (1988). *A Cognitive Theory of Consciousness*. New York: Cambridge University Press.

Baars, B. (1997). *In the Theater of Consciousness: The Workspace of the Mind*. New York: Oxford University Press.

Baars, B.J. (2002). The conscious access hypothesis: Origins and recent evidence. *Trends in Cognitive Science* 6. 1:47-52.

Baron-Cohen, Simon (1995). *Mindblindness: An Essay on Autism and Theory of Mind*. Cambridge, MA: MIT Press.

Barrett, H. C., and Kurzban, R. (2006). Modularity in cognition: Framing the debate. *Psychological Review* 113:628-647.

Botvinick, M. and Cohen, J. (1998). Rubber hands "feel" touch that eyes see. *Nature* 391:19.

Chalmers, David (1995). Facing Up to the Problem of Consciousness. *Journal of Consciousness Studies* 2. 3:200-19.

Chalmers, David John (1996*). The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford University Press.

Chomsky, Noam (2002). *On Nature and Language*. New York: Cambridge University Press.

Cosmides, Leda and John Tooby (1994). Origins of domain specificity: The evolution of functional organization. In Hirschfeld, Lawrence A. and Susan A. Gelman (eds.) *Mapping the Mind: Domain Specificity in Cognition and Culture* (pp. 85-116). New York, NY: Cambridge University Press.

Damasio, Antonio (1999). *The Feeling of What Happens: Body, Emotion and the Making of Consciousness*. London: Heinemann.

Darwin, Charles (1859). *The Origin of Species*. London: John Murray.

Davis, Martin (ed.) (2000). *Undecidable. Basic Papers on Undecidable Propositions, Unsolvable Problems and Computable Functions*. Mineola, NY: Dover.

Dehaene, Stanislas (ed.) (2002). *The Cognitive Neuroscience of Consciousness*. Cambridge, MA: MIT Press.

Del Cul, A., S. Baillet and S. Dehaene (2007). Brain dynamics underlying the nonlinear threshold for access to consciousness. *PLoS Biology* 5. 10: e260 doi:10.1371/journal.pbio.0050260

Dennett, Daniel (1991). *Consciousness Explained*. Boston, MA: Little, Brown and Co.

Fauconnier, Gilles (1999). Introduction to Methods and Generalizations. In T. Janssen and G. Redeker (eds). *Scope and Foundations of Cognitive Linguistics*. The Hague: Mouton De Gruyter. Cognitive Linguistics Research Series.

Fauconnier, Gilles and Mark Turner (1998). Conceptual Integration Networks. *Cognitive Science* 22.2: 133-187.

Fodor, Jerry (1983). *The Modularity of Mind: An Essay On Faculty Psychology*. Cambridge: MIT Press.

Gibbs, Raymond W. (2005). *Embodiment and Cognitive Science*. New York: Cambridge University Press.

Jackendoff, Ray (1987). *Consciousness and the Computational Mind*. Cambridge, MA: MIT Press.

James, Thomas W., Jody Culham, G. Keith Humphrey, A. David Milner and Melvyn A. Goodale (2003). Ventral occipital lesions impair object recognition but not object-directed grasping: an fMRI study. *Brain* 126. 11:2463-2475.

Koyré, Alexander (1957). *From the Closed World to the Infinite Universe*. Baltimore: Johns Hopkins University Press.

Lakoff, George and Mark Johnson (1999). *Philosophy in the Flesh*. New York: Basic Books.

Liu, Jia, Alison Harris, and Nancy Kanwisher (2002). Stages of processing in face perception: an MEG study. *Nature Neuroscience* 5: 910-916.

Marr, D. (1982). *Vision*. San Francisco: W.W. Freeman.

McGinn, Colin (1997). Can we solve the Mind-Body Problem? In Block, Flanagan and Güzeldere (eds). *The Nature of Consciousness: Philosophical Debates*. Cambridge, MA: MIT Press.

Metzinger, Thomas (ed.) (2000). *Neural Correlates of Consciousness: Empirical and Conceptual Questions*. Cambridge, MA: MIT Press.

Millikan, Ruth Garrett (1984). *Language, Thought, and Other Biological Categories: New Foundations for Realism.* Cambridge, MA: MIT Press.

Milner, A.D. and M.A. Goodale (1995). *The Visual Brain in Action*. Oxford: Oxford University Press.

Mithen, Steven (1996). *The Prehistory of the Mind: A Search for the Origins of Art, Religion and Science*. London: Thames and Hudson.

Mithieux, Gilles et al. (2005). Portal sensing of intestinal gluconeogenesis is a mechanistic link in the diminution of food intake induced by diet protein. *Cell Metabolism* 2. 5: 321-329.

Nagel, Thomas (1974). What is it like to be a bat? *Philosophical Review* 83. 4:435-50.

Newton, Sir Isaac (1974). *Newton's Philosophy of Nature. Selections from His Writings*. Ed. Horace S. Thayer. New York: Hafner Press.

Pearl, Judea (2000). *Causality: Models, Reasoning, and Inference*. New York: Cambridge University Press.

Quine, Willard van Orman (1969). *Word and Object*. Cambridge, MA: MIT Press.

Shepard, Roger N. (1981). Psychological complementarity. In M. Kubovy & J.R. Pomerantz (eds.) *Perceptual Organization* (pp. 279–342). Hillsdale, NJ: Lawrence Erlbaum Associates.

Spelke, Elizabeth S. (1990). Principles of object perception. *Cognitive Science* 14: 29-56.

Sperber, Dan (1994). The modularity of thought and the epidemiology of representations. In Hirschfeld, Lawrence A. and Susan A. Gelman (eds.) *Mapping the Mind: Domain Specificity in Cognition and Culture* (pp. 39-67). New York, NY: Cambridge University Press.

Sperber, Dan and Deirdre Wilson (1995). *Relevance: Communication and Cognition*. 2nd edition. Oxford: Blackwell.

Steen, Francis F. (2005). The paradox of narrative thinking. *Journal of Cultural and Evolutionary Psychology* 3. 1: 87-105.

Steen, Francis F. (in prep.) The framing effect: a recursive aggregation model.

Tsunoda, K. et al. (2001). Complex objects are represented in macaque inferotemporal cortex by the combination of feature columns. *Nature Neuroscience* 4. 8:832-838.

Tooby, John and Leda Cosmides (1992). The psychological foundations of culture. In Barkow, Jerome H., Leda Cosmides, and John Tooby (eds.) *The Adapted Mind: Evolutionary Psychology and the Generation of Culture* (pp. 19-136). New York: Oxford University Press.

Turing, Alan M. (1936). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society, 2nd series* 42: 230-265. The corrected 1937 version is available <http://web.comlab.ox.ac.uk/oucl/research/areas/ieg/e-library/sources/tp2-ie.pdf>

Turing, Alan M. (1950). Computing machinery and intelligence. *Mind* 59: 433-460. Available <http://www.abelard.org/turpap/turpap.htm>

Williams, G.C. (1966). *Adaptation and Natural Selection.* Princeton, NJ: Princeton UP.