©2004 The Institute of Mind and Behavior, Inc.
The Journal of Mind and Behavior
Summer 2004, Volume 25, Number 3
Pages 237–258
ISSN 0271–0137

Extending the Medium Hypothesis: The Dennett–Mangan Controversy and Beyond

Karl F. MacDorman

Osaka University

Mangan's hypothesis, that consciousness is an information-bearing medium, presents an alternative to Dennett's brand of functionalism, and Dennett's counterattacks have yet to address Mangan's main assertion. The medium hypothesis does not entail Cartesian theater assumptions concerning the localization, causal status, and "filling in" of consciousness in the brain. In principle, it is compatible with distributed information transfer between different media, epiphenomenalism, and gaps in visual experience. However, Mangan's strongest empirical argument, based on consciousness' limited "bandwidth," does not necessarily show that transduction between media of different information-bearing capacities occurs between the brain and consciousness. The features of consciousness that he attributes to a lower bandwidth medium can be explained in terms of functional constraints on a single medium. Furthermore, empirical results showing gaps and anomalies in visual experience speak against consciousness being a medium.

Keywords: filling in, functionalism, multiple realizability

Dennett's brand of functionalism boasts a third-person scientific approach to consciousness. Arguments against it typically take a first-person stance (Searle, 1992), which has often left little common ground for progress on developing a scientific understanding of consciousness. One notable exception is Mangan's (1993a) critique of Dennett (1991). It has the virtue of allowing one to examine the cogency of Dennett's views without becoming embroiled in disputes concerning whether it is valid to take a first-person stance. Mangan can join Searle in arguing against functionalism without taking aim at cognitive science or denying that information is an objective feature of the world. Mangan's medium hypothesis can be expressed in third-person terms, namely, that there is a medium operating during a specifiable range of cognitive activi-

Requests for reprints should be sent to Karl F. MacDorman, Ph.D., Department of Adaptive Machine Systems, Graduate School of Engineering, Osaka University, 2-1 Yamada-oka, Suita, Osaka 565–0871 Japan. Email: kfm@ed.ams.eng.osaka-u.ac.jp

ties (e.g., novelty processing, attention, global control) that is distinct from neurons as currently understood and that has a reduced capacity for carrying information (i.e., a narrow bandwidth). However, as a theory of consciousness, the medium hypothesis applies to both first- and third-person standpoints, while encouraging a convergent analysis of them.[1]

Functionalism, as generally construed, assumes that *any* information-bearing medium is conscious as long as it supports certain necessary functional relations (the *multiple realizability* principle). Mangan's critique of this position may be summarized as follows: since we have yet to rule out the possibility of consciousness being a *specific* kind of medium, functionalism cannot stand on strictly logical grounds. If only that specific medium can render what we experience conscious, functionalism is wrong. Furthermore, a scientific approach to an information-bearing medium must be able to investigate *how* a medium bears its information. However, the multiple realizability premise of functionalism could impede the scientific study of consciousness because it discourages medium-based investigations. Hence, a biological approach to studying information-bearing media needs to be applied to consciousness: With respect to biology, what are the similarities and differences between consciousness and other information-bearing media? One difference, according to Mangan, is that a subject's consciousness bears its information as that person's experience.

Regardless of its final standing, the medium hypothesis provides a third-person formulation that can accommodate common intuitions about consciousness. Dennett himself agrees and has complimented Mangan both in print (1993b) and at various public occasions for framing with "enviable clarity and vividness" the very idea he is trying to attack. Mangan (1998) and Dennett (1998) presented their opposing views at the Tucson II consciousness conference, where a lively debate ensued. This paper explores the ramifications of the medium hypothesis and how we might examine it empirically.

The Medium Hypothesis and Competing Accounts of Consciousness

The medium hypothesis, Mangan tells us, is not an all or nothing proposition as its applications range from the solid to the speculative, from third-person arguments to a detailed analysis of phenomenology, from a specific attack on functionalism to a broad program for consciousness research. Since Mangan's

---

[1]As an example of successful interaction between first- and third-person accounts, Mangan (1993c) cites Hering's (1878) opponent process theory derived largely from phenomenology. Herring's theory has helped to direct neurophysiological research, resulting in the discovery of cells in the eye, thalamus, and other visual areas that are inhibited by one primary color and excited by its opponent, or vice versa (DeValois, 1975). Mangan has used a convergent first-person/third-person approach to investigate context representation in consciousness (1991a, 1993b).

arguments (1992, 1993a, 1993b, 1998) are not as well known as Dennett's, they need to be summarized in greater detail below.

Mangan's critique of functionalism. Within our organism we can distinguish a class of systems that function to bear information. Neurons are only one type of information-bearing medium; others include RNA, DNA, blood, and the remarkably complex set of linked but distinct physical media in the ear, from the eardrum to the cochlea. In all cases, a full scientific account of these media must address two different kinds of questions: (1) *What* information does the medium bear? and (2) *How* does it bear that information? Thus, understanding DNA means finding out what genetic information DNA bears (for example, a recessive allele that causes sickle cell anemia but not the common cold) and how DNA bears its information (as nucleotides in base pairs in a double helix).

Aspects of the world that answer questions about what information the medium carries can be multiply instantiated without loss. In other words, the same information one medium bears can be fully transferred to an indefinite number of other media. The human genome project, for example, aimed to determine what information human DNA bears and to store this information in a very different medium, that of a digital computer (Mangan, 1998). Dennett's version of functionalism assumes that *all* questions about consciousness fall into this category.

From the standpoint of the medium hypothesis, aspects of the world that answer questions about how a medium bears its information *cannot* be instantiated in more than one way; that is, they cannot be abstracted from particular facts about a particular medium. One cannot even equate media that have *identical* functional capacities. According to Mangan, ignoring this point is one of the flaws of functionalism. To understand this, we need only consider a parallel example. It may be an accident of evolution that a particular set of four nucleotides constitutes RNA. As a matter of biochemistry, a different set of nucleotides might have worked just as well; however, a medium composed of different nucleotides could not — for the purposes of science — be RNA even if it were functionally identical. The medium would neither be chemically identical to RNA nor in its origin in natural history.

Now, if *consciousness* refers to a distinct information-bearing medium in our organism, if feeling pain is what it means for consciousness to instantiate the information that your arm is broken, then functionalism fails because a pain experience would not necessarily arise when the same information is instantiated by a different medium. Since functionalism tacitly assumes that consciousness is not a medium, and since support for the medium hypothesis undercuts functionalism, the medium hypothesis has the potential for turning functionalism into an empirically testable theory.

Two positive accounts of consciousness as a medium. Dennett (1993b, 1998) has argued that most theorists tacitly take consciousness to be a medium, though few recognize it. Thus, it is not difficult to find points of compatibility between the medium hypothesis and many of the theories that Dennett has criticized. Both Libet (1994) and Lowe (1996) have proposed medium-based accounts, and both claim that their accounts are empirically testable, although neither explicitly refers to consciousness being a medium. However, as we shall see, their accounts hardly exhaust the possible range of medium-based theories.

Libet (1994) posits the existence of a conscious mental field that influences neural events but cannot be reduced to them. It is unlike other fields in that the subject observes it as conscious experience, but others can only detect it through the subject's self-reports.[2] Ontologically speaking, the field is in a distinct phenomenological category, which cannot be reduced to neural processes. Local neuronal areas produce global changes in the conscious mental field; however, to discover the relationship between neural activity and the field, we must compare neural readings with concomitant conscious events. In principle, the predictions of Libet's hypothesis may be empirically testable: the experimenter isolates cortical slabs from the rest of the brain, either surgically or by chemical agents, and then electrically stimulates them. If the isolated slabs produced reportable experiences, his hypothesis would be supported. (However, a functionalist may counter that, if an isolated slab could produce conscious experience, it would be owing to the complex internal functional relations it supported.)

Lowe's (1996) medium-compatible account argues for mental events with causal powers that are independent of the causal powers of neural events.[3] Consciousness is seen as "an *emergent* feature of biological evolution — that is, a novel feature brought into being by biological processes but not itself a biological phenomenon" (p. 78):

> Consider the spider and its web. This is a complex system one feature of which — the web — is wholly produced by elements within the system (the relevant organs of the spider), and yet which, once produced, has independent causal powers. The web does things for the spider which it could not possibly do for itself unaided, and yet it is wholly created by the spider. (p. 80)

---

[2]Although *consciousness* is sometimes seen as something distinctly human and inseparable from language and society, this paper only attempts to examine its phenomenal and cognitive aspects, which are presumably also enjoyed by newborn infants, deaf-mutes, and many species of vertebrates.

[3]Lowe's account centers on what he defines to be the self: "a being which can identify itself as the unique subject of certain thoughts and experiences" (1996, p. 5). Since selves "have" minds without being distinct from minds, many of Lowe's hypotheses concerning the self may extend to consciousness.

We tend to regard the brain's neural activity as being seamlessly interconnected. Psychophysical causation seems implausible partly because it is unclear why causal chains in neural events should have any gaps that mental events must fill in. But we do not need to assume that mental events initiate chains of neural events for mental events to have independent causal powers. Returning to the web analogy, the spider's legs can initiate the spider's movements and its movements have no gaps, even if the web is ignored. This, of course, is because the web is not an initiating cause but an enabling cause that constrains and facilitates the spider's travel. For Lowe, conscious states constrain neural events much the way a web constrains a spider's activity.

Lowe proposes that during deliberative action mental acts may induce "certain patterns of convergence amongst neural events — patterns which, in the absence of such mental causes, would appear to involve miraculous coincidences, in view of the widely distributed and chaotic character of the neural activity prior to convergence" (p. 83). This hypothesis is empirically testable and supported by Libet's (1985) finding that a widely distributed build-up of electrical activity in the cortex ends with a sudden discharge along an efferent pathway causing a bodily movement just after a person's decision to make that movement:

> Thus, the proposal is that the occurrence of the conscious mental event of a decision to move a limb in a specific fashion, while it does not *initiate* any sequence of neural events culminating in such a movement, does serve to coordinate a host of mutually independent neural events so as to induce them to converge upon one specific pattern of efferent activity. (Lowe, 1996, pp. 83–84)

To Lowe, every deliberative movement is caused by a unique conscious decision, but neural events can only be partial causes of deliberative behavior. Thus, Lowe doubts token neural events can be mapped onto "token conscious events in a way which preserves isomorphism between their respective causal liaisons" (p. 84). This implies that consciousness is either epiphenomenal (has no causal influence on the brain and behavior) or, as Lowe suggests, has independent causal powers.

The point here is not to defend Libet or Lowe, but simply to provide some concrete examples of medium-based theories since Mangan does not. The medium hypothesis permits but does not entail Libet's or Lowe's position. The essence of the medium hypothesis is that consciousness is the particular medium that carries its information as a subject's conscious experience; other media cannot realize consciousness. However, the medium hypothesis is agnostic concerning the causal and biological status of consciousness.

Dennett's position: functionalism. Functionalism is perhaps still the dominant theory in the field of philosophy of mind, and some have taken it to be synonymous with cognitive science. Functionalism is the doctrine that all mental

states are constituted by functional relations between sensory input, internal states, and motor output. It is these functional relations that make a mental state the type of state it is, whether it be a toothache, the delight of a kiss, or the sensation of dampness on a wet night. Any specification of mental processes purely in terms of functional relations necessarily admits multiple realizability. That is, *in principle*, it is possible to realize those processes in other media, for example, in the transistors of a silicon semiconductor instead of in the neural pathways of the brain.

To Dennett, the functional role of a mental state must be understood in a teleological sense, namely, in terms of its biological purpose, not in the sense of mathematical function or causal role. One understands the psychological capacities of an organism by examining the functions of its component subsystems and how they work together to produce behavior. An understanding of an organism's biological function comes through a process that resembles reverse engineering. Dennett (1993b) does not consider functionalism to be at odds with behaviorism, but rather he takes the controversial stand that functionalism embodies a recognition of the fact that behavioral capacities place tight constraints on internal processing.

Dennett likens consciousness to fame; it is a kind of cerebral celebrity: "Those contents are conscious that persevere, that monopolize resources long enough to achieve certain typical and 'symptomatic' effects — on memory, on the control of behavior, and so forth" (1993b, p. 929; 2001). Just as being a star is not simply a matter of being on television (for one can go unnoticed on television in some minor or anonymous capacity), being conscious is not simply a matter of being in a medium. Both conscious and celebrity status are to be measured by their effects, not by the kind of medium that instantiates them. Thus, there are no *in principle* limitations on the form an information-bearing medium can take; only practical ones.

Why Dennett's Anti-Cartesian Arguments Fail to Address the Medium Hypothesis

Dennett's response to Mangan has been to reformulate his arguments against the Cartesian theater in terms of the medium hypothesis. The Cartesian theater is akin to a stage on which your brain plays out productions that it has scripted and choreographed for "you," the audience, to experience and act upon. Dennett (1993c, 1996a, p. 72, 1998) recasts the Cartesian theater as the claim that neural signals inexorably flow from the various sensory centers toward a special locale *and medium* in the brain where they are assembled into a multimodal representation of conscious experience. The consciousness medium is likened to a little person (or homunculus) who appreciates conscious experience and decides what to do. But the sheer synaptic intercon-

nectedness of the brain implies that messages cannot converge at a single place.

But even a definitive refutation of the Cartesian theater would not undercut the general claims of the medium hypothesis. The medium hypothesis does not entail the Cartesian theater. The Cartesian theater is only one way to apply the medium hypothesis to consciousness and the brain. It is in fact possible to argue in favor of the medium hypothesis while arguing against the Cartesian theater. By equating the two, Dennett (1993b, 1993c, 1998) incorrectly assumes that the medium hypothesis must be tied to the following assumptions:

Localization: The consciousness medium and the information it bears (its content) must be *spatiotemporally localized* in a specific region of the brain;

Causal Status: The consciousness medium must play a *causal role* in behavior. It cannot be epiphenomenal, but like the neural substrate, must also compute relational mappings; and finally

Filling In: The brain must create a complete, filled in representation of the world (a "finished product") before it can be transduced into the consciousness medium, experienced, and responded to by the subject. (Transduction involves the conversion of energy and information into another form — for example, sense organs transduce photons, sound waves, airborne molecules, and skin pressure into neural signals.)

Yet medium-based conceptions of consciousness are not bound by these assumptions. Even if neuroscientists were to demonstrate the absence of a Cartesian theater in your brain (i.e., that there is no one place where sensorimotor information comes together to form a multimodal representation), this will only refute the most naïve versions of the medium hypothesis.

Distributed transduction: why the localization assumption does not apply. In one analogy, Dennett (1998) likens the brain to the British Empire in microcosm, where British subjects are analogous to neurons. The War of 1812 ends with ambassadors signing a peace treaty in Ghent, but news of this reaches British and American troops too late to prevent the battle of New Orleans. In this example, we cannot make precise determinations about when the British Empire "knew" the war was over. We cannot say, for example, it was when the Prime Minister (or Parliament or King) was informed. By analogy, we cannot make precise determinations about when something "enters" consciousness. The whole idea of objects entering and exiting a certain region is wrong.

However, although we must necessarily conceive of the Prime Minister as occupying a single location in the British Empire, the medium hypothesis does not require us to conceive of a posited conscious medium as occupying a specific location in the brain. There is nothing about the medium hypothesis itself that precludes the possibility of transductive surfaces between the

posited consciousness medium and the neural substrate being distributed throughout wide, perhaps disjoint, areas of the brain. This is in fact Crick and Koch's (1990) hypothesis (also compare Libet, 1994; Lowe, 1996; Todorovic, 1987, p. 550).

Thus, we can recast the example of the British Empire in terms of *distributed transduction*: each person/neuron receives, at its particular time and place, information relating to the War of 1812 being over. The neurons then transfer that information into the consciousness medium. Aspects of experience are influenced by those neurons that are engaged in transducing those aspects. Thus, your experience of the war being over is the total effect over time of all the neurons engaged in transducing information (with the possible addition of whatever causal powers the medium itself might have). In upholding the medium hypothesis, one need not assume that contents must enter consciousness all at once (Mangan, 1993a, following James, 1890).

This, of course, leaves the problem of psychophysical causation: no neuroscientist has yet found evidence for transductive surfaces in the brain where electrochemical energy borne by neurons could pass into another medium. Even if transductive surfaces were to exist, they would be hard to detect. Although the transduction of information typically involves the passage of energy between media, quantum effects show that information has the potential for flowing instantaneously at a distance in the absence of a proximal energy flow. (An example of nonlocal transduction is the Einstein–Podolsky–Rosen effect: the spin of particles emitted by an atom exhibits correlations that persist even when the particles are distantly separated.)[4] Thus, if consciousness is a physical medium and transduction cannot be demonstrated in terms of conservation of energy, this does not necessarily imply the absence of transduction. This line of inquiry, however, remains highly speculative.

Epiphenomenalism: why the causal status assumption does not apply. By equating the medium hypothesis with the Cartesian theater, Dennett assumes that, as an information-bearing medium, consciousness would be required to compute mappings. This allows him to depict the medium hypothesis as homuncular: if there were a Cartesian theater in the brain, we would need yet another theory to explain what the little person inside that theater is doing and how he or she does it (Dennett, 1998). But the medium hypothesis does

[4]For example, when a spherically symmetric atom emits a pair of photons in opposite directions, each photon's rotational polarization is not defined until one of the photons is measured. Particle spins are able to maintain their mutual dependence, and thus provide a means of nonlocal transduction, because they are relatively insensitive to external interactions. Some physicists consider the nonlocality of particle spin correlations to belie the universe's underlying unity. We do not need special mechanisms to explain the nonlocality of global effects but to explain how the holistic connections between particles can decohere to the extent that physics becomes local (see Clarke, 1995).

not *require* consciousness to play a causal role. To bear information, it is not necessary to *process* it.

Not only does the medium hypothesis not entail the Cartesian theater, it seems somewhat at odds with Cartesian dualism. Descartes conceives of the soul (i.e., consciousness) as being a simple substance, both nonphysical and without parts or spatial extension. However, all known biological information-bearing media are complex, physical, and spatially extended. If Mangan's intent is to naturalize consciousness by showing that it can be treated as simply another biological information-bearing media, the medium could not be a simple substance. We cannot ask how a simple substance bears information because we cannot analyze it as there are no parts to analyze. Thus, it would be meaningless to ask medium-specific questions: "*How* does consciousness bear its information?" It is mainly because Descartes placed the soul beyond analysis that anti-homuncular arguments have force against Cartesian dualism.

However, if consciousness were epiphenomenal, that is, if it played no causal role in determining brain activity or behavior, how, in third-person terms, can it bear its information in its own *particular* way? One possible, though obscure explanation is that epiphenomenalism does not prevent consciousness from bearing information or even processing it; epiphenomenalism only bars the results of conscious activity from *returning* to influence the brain and behavior. Although the medium hypothesis is compatible with epiphenomenalism, it does, however, fit much better with a causal view of consciousness. If consciousness were epiphenomenal, it would stand out as a lone anomaly among biological information-bearing media.[5] In addition, a medium-based theory that did not afford consciousness a causal role would remain susceptible to all the arguments typically leveled against epiphenomenalism (cf. replies to Velmans, 1991).

In equating consciousness with "cerebral celebrity," Dennett defines conscious experience in terms of the functional role of its underlying representations. Thus, to Dennett, experience cannot be separated from its cognitive function. This may be why Dennett does not consider versions of the medium hypothesis, such as epiphenomenalism, that permit the decoupling of experience and conscious control.

---

[5]Epiphenomenalism of this kind is far removed from better known supervenience varieties in which the mental and physical are seen as different sides of the same coin. Mangan replaces this dualism with *media pluralism*: there are many information-bearing media, all are physical, and at least one is conscious. Mangan himself has been critical of arguments that deny consciousness a functional role. When Max Velmans (1991) claimed, on the basis of various experiments, that consciousness does not perform many of the cognitive functions for which it has been thought necessary, Mangan (1991b) accused him of falling prey to the *fallacy of functional exclusion*: just because a cognitive function may be performed in the absence of consciousness does not mean that, under normal circumstances, consciousness does not contribute to its performance.

We see this no clearer than in Dennett's (1998) tennis argument where he criticizes the contention, which he attributes to Libet (1993a, 1993b), that it takes between 200 and 500 milliseconds to become conscious of an item. If that were so, Dennett claims, tennis players could not be consciously playing tennis: even 200 ms is much too long for consciousness to aid a person in hitting a ball. Libet's findings, however, only concern stimuli that are just above the threshold of perceptibility. Thus, they provide at best an upper limit on how long it takes to become conscious of something (see Klein, 1995). Typical values are probably much lower, and it is these values that concern us in determining what role consciousness could take in playing tennis.[6]

Why the filling in assumption does not apply. Dennett (1998) identifies the medium hypothesis with the Cartesian theater assumption that the brain must replay in consciousness a filled-in representation of the world. Dennett's (1991, 1992, 1993a) point is that once your brain has drawn a conclusion, it does not need to go back and fill in a representation in some medium presumed to underlie your experience. However, the medium hypothesis does not *logically* entail filling in, although it would seem highly compatible with the filling-in assumption. In principle, a version of the medium hypothesis could broadly embrace the functionalist's account of representation but still require a special medium to render representations conscious. Clearly, such an explanation would lack parsimony (cf. Chalmers's, 1996, discussion of the X-factor). Therefore, evidence against filling in does cast doubt on the medium hypothesis. We will consider in more detail the relationship between the filling-in controversy and the medium hypothesis when we consider empirical issues below.

Deficiencies in Mangan's Bandwidth Argument: The Same Evidence Can Be Explained in Terms of Functional Constraints

Mangan (1998) credits the medium hypothesis with moving functionalism from being a metaphysical position into being, at least in principle, a testable hypothesis, amenable to empirical dispute, rather than simply conceptual dispute. Since we can never fully answer medium-specific questions by reference to a functional equivalent, we can never, *for purposes of science*, equate a

[6]Even if we pass over this point, Dennett's tennis argument only has force against the medium hypothesis if one assumes that everything you do in playing tennis depends on information becoming transduced into consciousness, being *decided upon* there, and then being transduced back onto the neural substrate. But, to adopt the medium hypothesis, there is no need to assume that consciousness plays a causal role in deciding how to hit a tennis ball. Some (if not all) of your well-honed skills, your split second reactions, may come into play outside of consciousness (see Baars, 1988), yet you are fully aware of acting and otherwise being a part of the game.

model of a given medium with that medium.[7] Since information and the media that bears it are distinct for the purposes of science, if the medium hypothesis is logically possible, no strictly logical argument can be given that will conclusively support functionalism. If *consciousness* refers to a medium, functionalism is flawed. If *consciousness* refers to an abstract class of relational mappings, the medium hypothesis is flawed. Other grounds besides strictly logical ones must be brought to bear in ultimately deciding for or against functionalism.

Mangan's (1998) other grounds for the medium hypothesis are (1) its intuitive appeal (you would not have to worry about attributing consciousness to robots *unless* they too had a means of transducing information into the consciousness medium); (2) the fact that it places qualia in a naturalistic framework (qualia is the representation in the consciousness medium);[8] and (3) his bandwidth argument. In searching for empirical support for either functionalism or the medium hypothesis, it is only reasonable to start by scrutinizing Mangan's bandwidth argument, since it rests on observations resulting from scientific inquiries.

Many experiments establish that consciousness has a remarkably limited capacity to bear information at any given moment. The narrowness of consciousness is in marked contrast to the massive capacity of both sensory media and the neural medium of the brain as exhibited by nonconscious parallel processing. All else being equal, different information-bearing media often do have different bandwidths. One empirical hypothesis is that, when part of a system has substantially diminished information-bearing capacity relative to the rest of the system, the observed diminished capacity is owing to the unobserved operation of a different medium. In Mangan's interpretation, the disparity between the narrow bandwidth of consciousness and the wide bandwidth of sensory systems and the neural substrate constitutes indirect evidence that consciousness is a distinct cognitive medium.

Mangan further claims that functionalism cannot explain why conscious processes should occupy precisely those points in a cognitive system where the cognitive system's capacity of articulation seems most limited, that is, where parallel processes give way to serial ones. The structure of consciousness, as exhibited by its *single stream*, *focus/fringe* (Mangan, 1991a, 1993b, 1993c), and

---

[7]One may object that modeling works by the principle of comparative likeness, so that for the purposes of science we can use a model of a phenomenon (say, DNA) as if it were the phenomenon. Nevertheless, scientists can never equate model and phenomenon, for if that were possible they could permanently disregard the empirical world, the very world in which they perform their experiments, and remain forever in the world of their models.

[8]In this way Mangan avoids following Searle (1992) in leaving qualia as a lone anomaly.

chunking limits (Miller, 1956), may be more limited than architectural considerations would lead us to expect (Mangan, 1998), and computer simulations might one day be useful in testing this hypothesis (Mangan, 2001). However, all the phenomena that Mangan explains in terms of a lower-bandwidth medium can also be explained in terms of functional constraints on cognitive processing, as the following alternative explanations show.

The limited capacity of consciousness. The annals of cognitive psychology contain many accounts of the limited capacity of consciousness. Baddeley (1993) has interpreted limited capacity in terms of limitations in working memory. Posner and Rothbart (1991) examined the close link between capacity constraints in consciousness and limitations in both working memory and selective visual attention (also see Cave and Wolfe, 1990; van der Heijden, 1992). However, it is unclear whether consciousness is limited and, thus, is imposing capacity constraints on short-term memory and attention or whether the inherent capacity constraints of these associated phenomena cause consciousness to appear limited.

If we follow Mangan in assuming that consciousness is inherently limited, we need not take recourse in Mangan's hypothesis that consciousness is a lower-bandwidth medium that is distinct from the higher-bandwidth neurons of the brain. If the functionalists are right, it may be that only a limited amount of the brain's information content is capable of meeting the *functional* requirements for being conscious at any one time. The configuration of the brain's computing elements (i.e., neurons), their speed, accuracy, and other performance characteristics impose certain constraints on the kinds of functions that the brain can compute and the speed with which it can compute them. Arguably most processes are not conscious because they fail to meet certain functional requirements. Factors relevant to determining whether a process is conscious — and how conscious it is — may include its accessibility, reportability, and availability for global control, its degree of integration into a global conceptualization and its impact on attention, learning, memory, and behavior (see Chalmers, 1996; Dennett, 1993b).

On this functionalist account, for a process to be conscious, it must at least be in a position to influence a vast number of other brain processes (Baars, 1988). However, the more bits of information a particular bit must coherently be integrated with, the more complex the resulting computation. Thus, computational constraints may create a selective pressure to evolve cognitive systems that integrate separate bits of information as little as is possible without jeopardizing performance (see MacDorman, 1999). In this view, the ideal architecture offloads as many tasks as it can to encapsulated, and presumably nonconscious, modules. As once novel activities become routine, response patterns are automated to function outside of consciousness (see Langer and

Imber, 1979), thus avoiding unnecessary expenditures of finite computational resources.[9]

Mangan would perhaps counter that neurons are adapted precisely for coordinating a massive number of constraints in parallel. Even so, the processes that functionalists identify with consciousness are so computationally demanding, that functional explanations are likely to predict bandwidth limits. The problem of maintaining a consistent and highly articulated representation (analogous to conscious contents) and allowing it to elicit a wide breath of activity is central to such unsolved problems as the frame problem in artificial intelligence.

Focal versus fringe consciousness. A functionalist could explain focal and fringe consciousness in terms of the degree to which conscious contents monopolize the brain's computing resources. More dominant contents are focal; less dominant ones are fringe (Mangan 1993a, 1993b).

The chunking limits of consciousness. The concept of bandwidth does not put an explicit constraint on the number of chunks in consciousness. The capacity of a channel (i.e., the maximum information that it can carry as measured in bits per second, cf. Shannon, 1948) is linearly proportional to its bandwidth (cycles per second). However, the same channel capacity can be used to carry an equivalent amount of information, regardless of whether that information is comprised of many small chunks or just a few large chunks. Fewer chunks do not imply less capacity because the information content of each chunk could have increased. Since large differences in information content hardly influenced the chunking capacity of Miller's (1956) subjects, there may be no direct relation between chunking limits and bandwidth.

Baars (1997, p. 196) has criticized Baddeley (1993) for equating conscious experience with short-term memory since "only the currently rehearsed item is conscious." This implies that consciousness has a lower chunking capacity than short-term memory. Thus, contrary to Mangan, it is not clear that the chunking limits of attention and short-term memory have any bearing on the chunking limits of consciousness. If the information-bearing capacity (i.e., bandwidth) of consciousness were the main factor in determining the capacity of short-term memory, one would expect the capacity of short-term memory to be highly sensitive to the information content of the items in memory. What was so surprising and significant about Miller's (1956) experiments is that they showed precisely the opposite: a sensitivity to the number of chunks and *not* their information content (Baddeley, 1994). It should be evident that

[9]Methods of automating, offloading, and otherwise *dis*integrating cognitive activities from the global conceptualization may all be part of the brain's solution to the frame problem: how it manages to settle on representational forms that avoid the need to reason about stabilities (Janlert, 1996).

Miller's explanation of chunking limits in terms of recoding is broadly compatible with functionalism.

The single stream of consciousness. According to Mangan, although it is not a necessary feature of consciousness that it courses down a single stream, the limited capacity of the consciousness medium enforces a practical sort of unity for reasons of efficiency. However, a functionalist might argue that consciousness enjoys the apparent unity of a single stream because, if there were multiple streams, the separate streams would need to be coordinated, and the computation required for that would serve as a sufficient functional criterion to integrate their conscious contents into a single stream. Another possibility is that your brain indeed manifests multiple streams of consciousness, but you only happen to be the experiencer of one of them.

Finally, it is a mistake for Mangan to attach too much significance to the particular design solutions evolution has found in us, such as a single, limited-capacity stream of consciousness with a focus and fringe. Because evolution's design solutions are often suboptimal, they cannot reveal the fundamental limits of the medium in which they are implemented. Consider, for example, the placement of rods and cones in the human eye. They are wired from the front, and this leaves us with a blind spot where the optic nerve must cross back to the brain. Would it not make more sense for our photoreceptors to be wired from the back (like those of invertebrates), thus preventing the blind spot? A functional limitation can often be attributed to some quirk in a species' evolutionary history. A simplistic response to why we have, for example, one stream of consciousness instead of many may be that evolution has not arrived at the multiple stream solution.

How Might "Filling In" Evidence Constrain Medium-Based Theories of Consciousness?

Conscious experience seems to have remarkable continuity and flow given that there are large gaps in our sensory input (e.g., the blind spot, intervals during eye saccades or blinking). We might assume that the continuity and flow of conscious experience mirror the continuity and flow of an underlying consciousness medium; however, such postulating does not explain why people do not notice gaps in conscious experience. Thus, we might conclude that the brain has already filled in gaps before transducing information from the neural substrate into the consciousness medium.

This position contrasts sharply with Dennett's functionalist account. What happens when you enter a room wallpapered with images of Marilyn Monroe? Dennett (1991, 1993a) argues that once your brain forms the appropriate conclusions, it does not have to fill in all the peripheral images for you to *see* a room wallpapered with Marilyns. You can still have a visual experience, even

if the representations underlying that experience are judgment-like (e.g., "put more Marilyns here"). Thus, there is no filling in of experience in part because there is nothing to fill in — there is no medium of consciousness, to use Mangan's new way of putting things (also see Akins and Winger, 1996; Dennett, 1996b).

On this account, you do not notice gaps in your sensory input because, for you to notice a gap, your brain must be doing something with it to represent it as a gap (see Dennett, 1991, p. 356). Thus, you cannot notice a gap if there are no "epistemologically hungry agencies" waiting for confirmation or disconfirmation from a "blind" area (1991, pp. 355; 1993a, p. 208). Consciousness may be essentially a gappy phenomenon, at least insofar as it makes sense to apply the spatial term gap to functional relations set up among judgment-like representations. Therefore, the reason conscious experience seems to have continuity and flow might be because most gaps in the sensory input never meet the necessary functional criteria to be conscious (e.g., persistence of representation and degree of influence on such cognitive processes as memory, attention, and control).

Churchland and Ramachandran (1993) have criticized Dennett (1991) for claiming that perceptual gaps are ignored and that matching areas of visual experience are not filled in. They point to psychophysical evidence for the filling in of the blind spot (Ramachandran, 1992), artificially induced scotomata (Ramachandran and Gregory, 1991), and scotomata resulting from brain lesions. They also note physiological evidence for cells in the V1 area corresponding to the optic disk whose receptive fields interpolate values from surrounding inputs (the Gattass effect reported in Fiorani, Rosa, Gattass, and Rocha–Miranda, 1992).[10]

Churchland and Ramachandran have taken these results as evidence against Dennett's hypothesis that the representations underlying visual experience are judgment-like (Dennett, 1996b). However, since Dennett's theory does not require that the form of a visual event match its resultant phenomenology, Dennett does not infer that the assertions of subjects concerning visual completion would refute his hypothesis (see Akins and Winger, 1996).

Filling in evidence may, however, be relevant to appraising the relative plausibility of alternative medium-based hypotheses about consciousness, because Mangan's theory *does* claim that the posited medium bears represen-

[10]To cite two further examples, recordings of single cells in the visual cortex (V2) of macaque monkeys have revealed neural correlates of illusory contours (von der Heydt, Peterhans, and Baumgartner, 1984). Extriate cells in V2 and V3 of active monkeys increased their response to a texture pattern with an equiluminant hole to levels matching their response to the pattern without a hole (De Weerd, Gattass, Desimone, and Ungerleider, 1995). The time-course of the increase in the cells' response matched reported perceptual completion in human subjects, lengthening in both instances with the size of the hole.

tations *as experience*. Mangan's view is that nonconscious neural processes "mold" consciousness into a flow of changing experiences much as one might mold a lump of clay.[11] Thus, presumably, phenomenological evidence for or against filling in imposes constraints on medium-based theories. If consciousness were a medium, "filling in" evidence should be able to tell us something about what kind of medium it would be.[12] We will consider two sets of stimuli that seem to confound a straightforward macro-physical interpretation of consciousness as a physical, spatially extended information-bearing medium.

Ramachandran (1992) created an experiment in which the top half of a bar is red, the bottom half is green, and the blind spot is a place along the border between the two colors. However, it is unclear to the subjects where the border is between red and green or, if there is no border, how the colors blend into one another. This ambiguity contrasts with the clear visual experience of a border when the border is in peripheral vision but outside of the blind spot. As Dennett notes,

> If there is *any* sort of filling in worthy of the name, then each sub-area of the bar-as-represented must be filled in either red or green (or "reddish green" as in the Crane and Piantanida [1983] experiment!). Or I suppose the areas could flicker back and forth between red and green, but one way or another, filling in requires explicit representation of the color at each "pixel" within the outline of the bar . . . . But if there isn't filling in, if the brain just concludes that it is a single solid bar with a red top and a green bottom and *does not go into the matter* of where and how the color changes, then there would be no fact of the matter about where "the boundary" was . . . . (1993a, pp. 207–208)

If we examine, say, a strip of film, we would expect there to be a fact of the matter about what its spectral reflectance is at every point along the strip (at least above the microscopic level). By the same token, if consciousness were a medium, one might expect there to be a fact of the matter about what color the bar is for every subarea of our visual experience.

Ramachandran (1992) also created an experiment in which the blind spot is placed over the point in a figure at which a dozen lines intersect (see also Churchland and Ramachandran, 1993). He reports that the 24 opposing spokes appear to complete across the blind spot. However, although the spokes complete, there is no clear point of intersection. MacDorman (2004) demonstrated in an experiment with a pair of identical dials, each consti-

---

[11]For Mangan, not all experiences need be representations, just as a lump of clay need not represent anything. Consciousness is not to be identified with the subject of experience. Rather, its current state is the experience itself, subjectivity being an aspect of that experience.

[12]This exhibits one benefit if consciousness were a medium: in addition to behavioral and neurophysiological evidence, first-person experience would constrain our scientific theories of consciousness. This would help in constructing, comparing, and testing theories: as theories are always underdetermined by the data, this new source of data could make a welcomed contribution to reducing a theory's possible degrees of freedom.

tuted by nine white lines on a black background, that the white circular region that appears when lines intersect in normal peripheral vision is missing from the dial that is centered on the blind spot. Furthermore, out of 27 subjects, more than 85% observed the white circular region disappear in normal peripheral vision within 90 seconds. The *unfilling* of the white circle appears to confirm Dennett's contention that, in addition to the blind spot, filling in does not necessarily occur in ordinary peripheral vision (1991, 1993a).

The twin dials experiment supports the contention that consciousness is not like a macrophysical medium. If consciousness were an information-bearing medium as other known physical media are, would we expect spokes to converge without touching, overlapping, or otherwise intersecting? Should we expect an "area" of visual experience to have no color (or an absence of color) associated with it?[13] Moreover, if one looks beyond visual experience, much of what happens in consciousness appears to be nonspatial or only vaguely spatial (e.g., thinking verbally or adding in one's head). It is unclear how a physical, spatially extended medium would bear such apparently nonspatial information. If one tries to combine the conception of consciousness being a medium with a functionalist explanation of representation and filling in, it becomes less clear what new insights the medium hypothesis would bring to a third-person, scientific understanding of consciousness.

Conclusion

Dennett and Mangan have provided two opposing and potentially testable views on consciousness, both of which are amenable to a third-person formulation and empirical study. To Mangan, our conscious experience is "manifested" by a distinct medium and "shaped" by nonconscious neural processes. To Dennett, our conscious experience is constituted by judgment-like representations, representations implemented by biological brains and, in principle, by artificial brains in robots; the brain does not need to transform judgment-like representations into a visual (or other modal) form, much less transduce them into a privileged medium, for us to enjoy the richness of our experiences (Dennett, 1996b). Is consciousness a medium or content system? While these competing positions hardly exhaust the alternatives, they may well lead to widely different implications for fields as different as robotics and theology. (Could we build a conscious robot? Could consciousness persist after death — and in what form?)

---

[13]From a macrophysical standpoint, it does not even make sense to apply the concepts *area* and *unassigned* to the same visual experience, since area in *physical* space is compatible only with the medium hypothesis while (abstract) color tag assignments are compatible only with functionalism.

Mangan's case against functionalism depends on its incompatibility with the medium hypothesis. His argument is in need of refinement because, on certain interpretations, especially those that eschew multiple realizability, functionalism and the medium hypothesis might be compatible. Chalmers (1996), for example, has argued that universally applicable coherence laws (pp. 242–246) render versions of functionalism compatible with a nonreductive conception of experience (pp. 93–111, and even panpsychism, pp. 293–301). Properties of consciousness may supervene on (i.e., be fully determined by) physical properties in the sense of being "systematically and perfectly correlated" in this universe (p. 36) but not in any possible universe, as logical supervenience would imply. Chalmers upholds that in a universe in which they are uncorrelated, there could be worlds populated by zombies who, apart from being unconscious, are in every way like us (a possibility denied by Dennett).

Nevertheless, Mangan has shown that Dennett's reductive version of functionalism cannot stand only on logical grounds. Further empirical investigations are in order. However, Mangan's bandwidth argument (based on the observation that a change in information-bearing capacity often indicates a change in media) cannot necessarily rule out functionalism because all the features of consciousness that he attributes to a change in bandwidth (e.g., capacity limits, single stream, focus and fringe) can be explained in terms of function alone. As we noted, Miller's results on short-term memory and attention seem more compatible with functionalism than with the medium hypothesis. Mangan assumes that the information capacity of consciousness is proportional to the information capacity of the medium that bears it. But information capacity does not seem to be the crucial factor since the capacity limits of short-term memory and attention are sensitive to the number of chunks in consciousness rather than their information content.

We have found that the medium hypothesis allows for a wide range of theories about consciousness, much wider than Dennett or Mangan seem to acknowledge. It is compatible with both epiphenomenalism and the attribution of independent causal powers to consciousness; it is compatible with local transduction and the distributed transduction of Libet and Lowe.

A major issue for medium-based theories, however, is psychophysical causation: How is information transduced between consciousness and neurons? Neuroscientists do not generally expect to find transductive surfaces in the brain (e.g., points of energy loss or gain), though quantum effects *might* explain a flow of information in the absence of an energy flow. Mangan insists that the medium hypothesis is not undercut by our current inability to observe or otherwise specify the precise physical process that "manifests" the medium. Ignorance of this sort is not uncommon in the early stages of a research program. It took a century to go from Darwin's belief in an unobserved medium that bore hereditary information to Watson and Crick's discovery of

its structure. Mangan reasons that, if Darwin was justified to work on the theory of natural selection without knowledge of its basis in DNA, we too are justified to explore the medium hypothesis.[14] Whether such an investigation will bear fruit remains to be seen.

There are, however, steps we can take now to examine functionalism and the medium hypothesis. The seeming continuity and flow of visual experience stand in marked contrast to the many discontinuities that appear in the visual input. Since Mangan's medium is to bear its information as experience, this suggests that for experience to appear filled in, the medium must also be filled in. In some ways visual anomalies and other evidence against filling in seem to sit better with Dennett's propositional conception of content. Empirical studies of phenomena like filling in may have the potential either to refute the notion of consciousness being a medium or to elucidate its nature.

References

Akins, K.A., and Winger, S. (1996). Ships in the night: Churchland and Ramachandran on Dennett's theory of consciousness. In K. Akins, *Perception* (Vancouver Studies in Cognitive Science, Volume 5, pp. 173–198). New York: Oxford University Press.

Baars, B.J. (1988). *A cognitive theory of consciousness*. Cambridge: Cambridge University Press.

Baars, B.J. (1997). Contrastive phenomenology: A thoroughly empirical approach to consciousness. In N. Block, O. Flanagan, and G. Guzeldere (Eds.), *The nature of consciousness: Philosophical debates* (pp. 187–202). Cambridge, Massachusetts: MIT Press.

Baddeley, A. (1993). Working memory and conscious awareness. In A.F. Collins and M.A. Conway (Eds.), *Theories of memory* (pp. 11–28). Hove, United Kingdom: Erlbaum.

Baddeley, A. (1994). The magic number 7: Still magic after all these years. *Psychological Review, 101*, 353–356.

Blackmore, S.J., Brelstaff, G., Nelson, K., and Troscianko, T. (1995). Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception, 24*, 1075–1081.

Cave, K.R., and Wolfe, J.M. (1990). Modeling the role of parallel processing in visual search. *Cognitive Psychology, 22*, 225–271.

Chalmers, D.J. (1996). *The conscious mind: In search of a fundamental theory*. Oxford: Oxford University Press.

Crick, F., and Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences, 2*, 263–275.

Churchland, P.S., and Ramachandran, V.S. (1993). Filling in: Why Dennett is wrong. In B. Dahlbom (Ed.), *Dennett and his critics: Demystifying mind* (pp. 28–52). Oxford: Blackwell.

Clarke, C.J.S. (1995). The nonlocality of mind. *Journal of Consciousness Studies, 2*, 231–240.

Crane, H.D., and Piantanida, T.P. (1983). On seeing reddish-green and yellowish-blue. *Science, 221*, 1078–1079.

Dennett, D.C. (1991). *Consciousness explained*. Boston: Little, Brown and Co.

Dennett, D.C. (1992). "Filling in" versus finding out: A ubiquitous confusion in cognitive science. In H.L. Pick, Jr., P. van den Broek, and D.C. Knill (Eds.), *Cognition: Conceptual and methodological issues* (pp. 33–49). Washington, DC: American Psychological Association.

[14]Mangan points out that nothing prevents us from using the concept of medium specificity in the absence of a microphysical account nor from obtaining indirect third-person evidence for a medium from behavioral and physiological observations (in a way analogous to how physicists obtained sufficient indirect evidence for cosmic rays on the basis of vapor trails in a cloud chamber).

Dennett, D.C. (1993a). Back from the drawing board. In B. Dahlbom (Ed.), *Dennett and his critics: Demystifying mind* (pp. 203–235). Oxford: Blackwell.
Dennett, D.C. (1993b). The message is: There is no medium. *Philosophy and Phenomenological Research, 53*, 919–931.
Dennett, D.C. (1993c). Caveat emptor. *Consciousness and Cognition, 2*, 48–57.
Dennett, D.C. (1994). The practical requirements for making a conscious robot. *Philosophical Transactions of the Royal Society, A, 349*, 133–146.
Dennett, D.C. (1995). The unimagined preposterousness of zombies. *Journal of Consciousness Studies, 2*, 322–326.
Dennett, D.C. (1996a). *Kinds of minds: Towards an understanding of consciousness*. London: Weidenfeld and Nicolson.
Dennett, D.C. (1996b). Seeing is believing — or is it? In K. Akins, *Perception* (Vancouver Studies in Cognitive Science, Volume 5, pp. 158–172). New York: Oxford University Press.
Dennett, D.C. (1998). The myth of double transduction. In S.R. Hameroff (Ed.), *Toward a science of consciousness II: The second Tucson discussions and debates* (pp. 97–107). Cambridge, Massachusetts: MIT Press.
Dennett, D.C. (2001). Are we explaining consciousness yet? *Cognition, 79*, 221–237.
DeValois, R.L. (1975). Neural coding of color. In E.C. Carterette and M.P. Friedman (Eds.), *Handbook of perception* (Volume 5, pp. 117–166). New York: Academic Press.
De Weerd, P., Gattass, R., Desimone, R., and Ungerleider, L. (1995). Responses of cells in monkey visual cortex during perceptual filling-in of an artificial scotoma. *Nature, 377*, 731–734.
Fiorani, M., Rosa, M.G.P., Gattass, R., and Rocha–Miranda, C.E. (1992). Visual responses outside the "classical" receptive field in primate striate cortex; A possible correlate of perceptual completion. *Proceedings of the National Academy of Science, 89*, 8547–8551.
Hering, E. (1878). *Zur Lehre vom Lichtsinne* [On the Theory of the Light Sense] (second edition). Vienna: Carl Gerold & Sohn.
Hofstadter, D.R., and Dennett, D.C. (1981). *The mind's I: Fantasies and reflections on self and soul*. New York: Basic Books.
James, W. (1890). *The principles of psychology*. New York: H. Holt.
Janlert, L.E. (1996). The frame problem: Freedom or stability? With pictures we can have both. In K.M. Ford and Z.W. Pylyshyn (Eds.), *The robot's dilemma revisited: The frame problem in artificial intelligence* (pp. 35–48). Norwood, New Jersey: Ablex.
Klein, S. (1995). Is quantum mechanics relevant to understanding consciousness?: A review of *Shadows of the Mind* by Roger Penrose. *Psyche: An Interdisciplinary Journal of Research on Consciousness, 2*, 1–6.
Langer, E.J., and Imber, L.G. (1979). When practice makes imperfect: Debilitating effects of overlearning. *Journal of Personality and Social Psychology, 37*, 2014–2024.
Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences, 8*, 529–539.
Libet, B. (1993a). The neural time factor in conscious and unconscious events. *Experimental and theoretical studies of consciousness* (Ciba Foundation Symposia 174, pp. 123–146). Chichester: Wiley.
Libet, B. (1993b). *Neurophysiology of consciousness: Selected papers and new essays*. Boston: Birkhauser.
Libet, B. (1994). A testable field theory of mind–brain interaction. *Journal of Consciousness Studies, 1*, 119–126.
Lowe, E.J. (1996). *Subjects of experience*. Cambridge: Cambridge University Press.
MacDorman, K.F. (1999). Grounding symbols through sensorimotor integration. *Journal of the Robotics Society of Japan, 17*, 20–24.
MacDorman, K.F. (2004). What "unfilling in" says about the nature of representation in the brain. *Second International Symposium on Emergent Mechanisms in the Brain*, March 1–3, Shumagun, Japan, p. 88.
Mangan, B. (1991a). *Meaning and the structure of consciousness*. Unpublished doctoral dissertation. University of California, Berkeley.
Mangan, B. (1991b). Epi-arguments for epiphenomenalism: A peer commentary on Velmans' "Is Human Information Processing Conscious?" *Behavioral and Brain Sciences, 14*, 689.

Mangan, B. (1992). Consciousness unexplained. *The Sciences, 32,* 55–56.
Mangan, B. (1993a). Dennett, consciousness, and the sorrows of functionalism. *Consciousness and Cognition, 2,* 1–17.
Mangan, B. (1993b). Taking phenomenology seriously: The fringe and its implications for cognitive research. *Consciousness and Cognition, 2,* 89–108.
Mangan, B. (1993c). Some philosophical and empirical implications of the fringe. *Consciousness and Cognition, 2,* 142–154.
Mangan, B. (1998). Against functionalism: Consciousness as an information-bearing medium. In S.R. Hameroff (Ed.), *Toward a science of consciousness II: The second Tucson discussions and debates* (pp. 135–141). Cambridge, Massachusetts: MIT Press.
Mangan, B. (2001). Sensation's ghost: The non-sensory fringe of consciousness. *Psyche: An Interdisciplinary Journal of Research on Consciousness, 7*/18). From http://psyche.cs.monash.edu.au/v7/psyche-7-18-mangan.html
Marcel, A.J. (1983). Conscious and unconscious perception: Experiments on visual masking and word recognition. *Cognitive Psychology, 15,* 197–237.
Miller, G. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review, 63,* 81–96.
Posner, M.I., and Rothbart, M.K. (1991). Attentional mechanisms and conscious experience. In A.D. Milner and D. Rugg (Eds.), *The neuropsychology of consciousness* (pp. 91–112). New York: Academic Press.
Putnam, H. (1967a). The mental life of some machines. In H. Castaneda (Ed.), *Intentionality, minds and perception* (pp. 439–460). Detroit: Wayne State University Press.
Putnam, H. (1967b). Psychological predicates. In W.H. Capitan and D.L. Merrill (Eds.), *Art, mind and religion* (pp. 37–48). Pittsburgh: University of Pittsburgh Press.
Ramachandran, V.S. (1992). Blind spots. *Scientific American, 266,* 86–91.
Ramachandran, V.S., and Gregory, R.L. (1991). Perceptual filling in of artificially induced scotomas in human vision. *Nature, 350,* 699–702.
Ryle, G. (1949). *The concept of mind.* London: Hutchinson.
Searle, J.R. (1992). *The rediscovery of the mind.* Cambridge, Massachusetts: MIT Press.
Shannon, C.E. (1948). A mathematical theory of communication. *Bell System Technical Journal, 27,* 379–423, 623–656.
Todorovic, D. (1987). The Craik–O'Brien–Cornsweet effect; New varieties and theoretical implications. *Perception and Psychophysics, 42,* 545–600.
Velmans, M. (1991). Is human information processing conscious? *Behavioral and Brain Sciences, 14,* 651–668.
van der Heijden, A.H.C. (1992). *Selective attention in vision.* London: Routledge.
von der Heydt, R., Peterhans, E., and Baumgartner, G. (1984). Illusory contours and cortical neuron responses. *Science, 224,* 1260–1262.