BEHAVIORAL AND BRAIN SCIENCES (2003) **26,** 375–444
Printed in the United States of America

Gestalt isomorphism and the primacy of subjective conscious experience: A Gestalt Bubble model

Steven Lehar
Peli Lab, The Schepens Eye Research Institute, Boston MA 02114-2500.
slehar@cns.bu.edu **http://cns-alumni.bu.edu/~slehar**

Abstract: A serious crisis is identified in theories of neurocomputation, marked by a persistent disparity between the phenomenological or experiential account of visual perception and the neurophysiological level of description of the visual system. In particular, conventional concepts of neural processing offer no explanation for the holistic global aspects of perception identified by Gestalt theory. The problem is paradigmatic and can be traced to contemporary concepts of the functional role of the neural cell, known as the *Neuron Doctrine.* In the absence of an alternative neurophysiologically plausible model, I propose a *perceptual modeling* approach, to model the percept as experienced subjectively, rather than modeling the objective neurophysiological state of the visual system that supposedly subserves that experience. A Gestalt Bubble model is presented to demonstrate how the elusive Gestalt principles of emergence, reification, and invariance can be expressed in a quantitative model of the subjective experience of visual consciousness. That model in turn reveals a unique computational strategy underlying visual processing, which is unlike any algorithm devised by man, and certainly unlike the atomistic feed-forward model of neurocomputation offered by the Neuron Doctrine paradigm. The perceptual modeling approach reveals the primary function of perception as that of generating a fully spatial virtual-reality replica of the external world in an internal representation. The common objections to this "picture-in-the-head" concept of perceptual representation are shown to be ill founded.

Keywords: brain-anchored; Cartesian theatre; consciousness; emergence; extrinsic constraints; filling-in; Gestalt; homunculus; indirect realism; intrinsic constraints; invariance; isomorphism; multistability; objective phenomenology; perceptual modeling; perspective; phenomenology; psychophysical parallelism; psychophysical postulate; qualia; reification; representationalism; structural coherence

1. Introduction

Contemporary neuroscience finds itself in a state of serious crisis, for the deeper we probe into the workings of the brain, the farther we seem to get from the ultimate goal of providing a neurophysiological account of the mechanism of conscious experience. Nowhere is this impasse more evident than in the study of visual perception, where the apparently clear and promising trail discovered by Hubel and Wiesel (1959) leading up the hierarchy of feature detection from primary to secondary and to higher cortical areas seems to have reached a theoretical dead end. Besides the troublesome issues of the noisy stochastic nature of the neural signal and the very broad tuning of the single cell as a feature detector, the notion of visual processing as a hierarchy of feature detectors seems to suggest some kind of "grandmother cell" model in which the activation of a single cell or a group of cells represents the presence of a particular type of object in the visual field. However, it is not at all clear how such a featural description of the visual scene could even be usefully employed in practical interaction with the world.

Alternative paradigms of neural representation have been proposed, including the suggestion that synchronous oscillations play a role in perceptual representation, although these theories are not yet specified sufficiently to know exactly how they address the issue of perceptual representation. But the most serious indictment of contemporary neurophysiological theories is that they offer no hint of an explanation for the subjective experience of visual consciousness. Visual experience is more than just an abstract recognition of the features present in the visual field – those

STEVEN LEHAR, Ph.D., is an independent researcher at the Schepens Eye Research Institute in Boston, Mass., USA. He is the author of twelve different papers on subjects ranging from new paradigms and forms of neurocomputation, to philosophical papers on epistemology and the structure of conscious experience. A principle focus of Lehar's work is on the implications of Gestalt theory for the nature of perceptual computation and representation in the brain, including the role of feedback in visual processing, and harmonic resonance as an explanation for a number of illusory grouping phenomena. Lehar is also author of *The World In Your Head: A Gestalt View of the Mechanism of Conscious Experience* (2003; Erlbaum), a book that covers most of his theories across a wide range of subjects from vision to cognition to motor control. Lehar is winner of the 1999 Wolfgang Metzger award for significant contribution to Gestalt theory, awarded by the Gestalt Theory and Applications (GTA) society.

features are vividly experienced as solid three-dimensional objects, bounded by colored surfaces, embedded in a spatial void. A number of enigmatic properties of this world of experience were identified decades ago by Gestalt theory, suggestive of a holistic emergent computational strategy whose operational principles remain a mystery.

The problem in modern neuroscience is a paradigmatic one that can be traced to its central concept of neural processing. According to the *Neuron Doctrine,* neurons behave as quasi-independent processors separated by relatively slow chemical synapses, with strictly segregated input and output functions through the dendrites and axon, respectively. It is hard to imagine how such an assembly of independent processors could account for the holistic emergent properties of perception identified by Gestalt theory. In fact, the reason these Gestalt aspects of perception have been largely ignored in recent decades is exactly because they are so difficult to express in terms of the Neuron Doctrine paradigm. More recent proposals that implicate synchronous oscillations as the neurophysiological basis of conscious experience (Crick 1994; Crick & Koch 1990; Eckhorn et al. 1988; Llinas et al. 1994; Singer 1999; Singer & Gray 1995) seem to suggest some kind of holistic global process that appears to be more consistent with Gestalt principles, although it is hard to see how this paradigm, at least as currently conceived, can account for the solid three-dimensional nature of subjective experience. The persistent disparity between the neurophysiological and phenomenal levels of description suggests that either the subjective experience of visual consciousness is somehow illusory, or the state of our understanding of neural representation is far more embryonic than is generally recognized.

Pessoa et al. (1998) made the case for denying the primacy of conscious experience. They argued that although the subjective experience of filling-in phenomena is sometimes accompanied by a neurophysiological correlate, such an isomorphism between experience and neurophysiology is not logically necessary but is merely an empirical issue. For, they claimed, subjective experiences can occur in the absence of a strictly isomorphic correlate. Their view is that although the subjective experience of visual consciousness appears as a "picture" or three-dimensional model of a surrounding world, this does not mean that the information manifest in that experience is necessarily explicitly encoded in the brain. Moreover, that consciousness is an illusion based on a far more compressed or abbreviated representation, in which percepts such as that of a filled-in colored surface can be explained neurophysiologically by "ignoring an absence" rather than by an explicit point-for-point mapping of the perceived surface in the brain.

In fact, nothing could be further from the truth. For to propose that the subjective experience of perception can be more enriched and explicit than the corresponding neurophysiological state, flies in the face of the materialistic basis of modern neuroscience. The modern view is that mind and brain are different aspects of the same physical mechanism. In other words, every perceptual experience, whether a simple percept such as a filled-in surface or a complex percept of a whole scene, has two essential aspects, the subjective experience of the percept and the objective neurophysiological state of the brain that is responsible for that subjective experience. Like the two faces of a coin, these very different entities can be identified as merely different manifestations of the same underlying structure, viewed from the internal first-person perspective as opposed to the external third-person perspective. The dual nature of a percept is analogous to the representation of data in a digital computer, where a pattern of voltages present in a particular memory register can represent some meaningful information, such as a numerical value, a brightness value in an image, or a character of text, when viewed from inside the appropriate software environment, but when viewed in external physical terms those same data take the form of voltages or currents in particular parts of the machine. However, whatever form is selected for encoding data in the computer, the information content of that data cannot possibly be of higher dimensionality than the information explicitly expressed in the physical state of the machine.

The same principle must also hold in perceptual experience, as proposed by Müller (1896) in the *psychophysical postulate.* Müller argued that because the subjective experience of perception is encoded in some neurophysiological state, the information encoded in that conscious experience cannot possibly be any greater than the information encoded in the corresponding neurophysiological state. Although we cannot observe phenomenologically the physical medium by which perceptual information is encoded in the brain, we can observe the information encoded in that medium, expressed in terms of the variables of subjective experience. It follows therefore that it should be possible by direct phenomenological observation to determine the dimensions of conscious experience, and thereby to infer the dimensions of the information encoded neurophysiologically in the brain.

The bottom-up approach that works upward from the properties of the individual neuron and the top-down approach that works downward from the subjective experience of perception are equally valid and complementary approaches to the investigation of the visual mechanism. Eventually, these opposite approaches to the problem must meet somewhere in the middle. To date, however, the gap between them remains as large as it ever was. Both approaches are essential to the investigation of biological vision because each offers its own unique perspective on the problem. The disparity between these two views of the visual representation helps to maintain the focus on exactly those properties that are prominently absent from the conventional neural network view of visual processing.

2. The epistemological divide

There is a central philosophical issue that underlies discussions of phenomenal experience as seen, for example, in the distinction between the Gestaltist and the Gibsonian views of perception. That is, the epistemological question of whether the world we see around us is the real world itself or merely an internal perceptual copy of that world generated by neural processes in our brain. In other words, this is the question of *direct realism* (also known as *naïve realism*) as opposed to *indirect realism* (or *representationalism*). To take a concrete example, consider the vivid spatial experience of this paper that you hold in your hands. The question is whether the rich spatial structure of this experience before you is the physical paper itself, or an internal data structure or pattern of activation within your physical brain. Although this issue is not much discussed in contemporary psychology, it is an old debate that has resurfaced several times in psychology, and the failure to reach consensus on this issue continues to bedevil the debate on the functional role of sensory processing. The reason for the continued confusion is that both direct and indirect realism are frankly incredible, although each is incredible for different reasons.

2.1. Problems with direct realism

The direct realist view is incredible because it suggests that we can have experience of objects out in the world directly, beyond the sensory surface, as if bypassing the chain of sensory processing. For example, if light from this paper is transduced by your retina into a neural signal that is transmitted from your eye to your brain, then the very first aspect of the paper that you can possibly experience is the information at the retinal surface, or the perceptual representation that is downstream of it in your brain. The physical paper itself lies beyond the sensory surface and therefore must be beyond your direct experience. But the perceptual experience of the page stubbornly appears out in the world itself instead of in your brain, in apparent violation of everything we know about the causal chain of vision. Gibson explicitly defended the notion of direct perception and spoke as if perceptual processing occurs somehow out in the world itself rather than as a computation in the brain based on sensory input (Gibson 1972, pp. 217, 239).

Significantly, Gibson refused to discuss sensory processing at all and even denied that the retina records anything like a visual image that is sent to the brain. This leaves the status of the sensory organs in a peculiar kind of limbo, for if the brain does not process sensory input to produce an internal image of the world, what is the purpose of all that computational wetware? Another embarrassment for direct perception is the phenomenon of visual illusions, which are observed out in the world itself; and yet they cannot possibly be in the world for they are the result of perceptual processing that must occur within the brain. With characteristic aplomb, Gibson simply denied that illusions are illusory at all, although it is not clear exactly what he could possibly have meant by that. Modern proponents of Gibson's theories usually take care to disclaim his most radical views (Bruce & Green 1987, pp. 190, 203–204; O'Regan 1992, p. 473; Pessoa et al. 1998), but they present no viable alternative explanation to account for our experience of the world beyond the sensory surface.

The difficulty with the concept of direct perception is most clearly seen when we consider how an artificial vision system could be endowed with such external perception. Although a sensor may record an external quantity in an internal register or variable in a computer, from the internal perspective of the software running on that computer, only the internal value of that variable can be "seen" or can possibly influence the operation of that software. In an exactly analogous manner the pattern of electrochemical activity that corresponds to our conscious experience can take a form that reflects the properties of external objects, but our consciousness is necessarily confined to the experience of those internal effigies of external objects, rather than of the external objects themselves. Unless the principle of direct perception can be demonstrated in a simple artificial sensory system, this explanation remains as mysterious as the property of consciousness it is supposed to explain.

2.2. Problems with indirect realism

The indirect realist view is also incredible, for it suggests that the solid stable structure of the world we perceive to surround us is merely a pattern of energy in the physical brain; that is, the world that appears to be external to our head is actually inside our head. This could only mean that the head we have come to know as our own is not our true physical head but is merely a miniature perceptual copy of our head inside a perceptual copy of the world, all of which is completely contained within our true physical skull. Stated from the internal phenomenal perspective, out beyond the farthest things you can perceive in all directions (i.e., above the dome of the sky and below the earth under your feet, or beyond the walls, floor, and ceiling of the room you perceive around you), beyond those perceived surfaces is the inner surface of your true physical skull encompassing all that you perceive, and beyond that skull is an unimaginably immense external world, of which the world you see around you is merely a miniature virtual-reality replica. The external world and its phenomenal replica cannot be spatially superimposed, for one is inside your physical head and the other is outside. Therefore, the vivid spatial structure of this page that you perceive here in your hands is itself a pattern of activation within your physical brain, and the real paper of which it is a copy is out beyond your direct experience.

I have found a curious dichotomy in the responses of colleagues in discussions on this issue. Many people agree with the statement that everything you perceive is in some sense inside your head, and in fact they often complain that this is so obvious it need hardly be stated. However, when that statement is turned around to say that out beyond everything you perceive is your physical skull, they object most vehemently that that is absurd. And yet the two statements are logically identical, so how can one appear trivially obvious while the other seems patently absurd? The value of this particular mental image is that it helps to smoke out any residual naive realism that may remain hidden in our philosophy. For although this statement can only be true in a topological, rather than a strict topographical, sense, this insight emphasizes the indisputable fact that no aspect of the external world can possibly appear in consciousness except by being represented explicitly in the brain. The existential vertigo occasioned by this mental image is so disorienting that only a handful of researchers have seriously entertained this notion or pursued its implications to its logical conclusion (Harrison 1989; Hoffman 1998; Kant 1781/1991; Koffka 1935; Köhler 1971, p. 125; Lehar 2003b; Russell 1927, pp. 137–143; Smythies 1989; 1994).

Another reason the indirect realist view is incredible is that the observed properties of the world of experience when viewed from the indirect realist perspective are difficult to resolve with contemporary concepts of neurocomputation. For the world we perceive around us appears as a solid spatial structure that maintains its structural integrity as we turn around and move about in the world. Perceived objects within that world maintain their structural integrity and recognized identity as they rotate, translate, and scale by perspective in their motions through the world. These properties of the conscious experience fly in the face of everything we know about neurophysiology, for they suggest some kind of three-dimensional imaging mechanism in the brain, capable of generating three-dimensional volumetric percepts of the degree of detail and complexity observed in the world around us. No plausible mechanism has ever been identified neurophysiologically which exhibits this incredible property. The properties of the phenomenal world are therefore inconsistent with contemporary concepts of neural processing, which is exactly why these properties have been so long ignored.

2.3. Spirituality, supervenience, and other nomological danglers

The perceived incredibility of both direct and indirect realism has led many over the centuries to propose that conscious experience is located neither in the physical brain nor in the external world, but in some separate space that bears no spatial relation to the physical space known to science. These theories fall somewhere between direct and indirect perception because they claim that phenomenal experience is neither in the head, nor out in the world. The original formulation of this thesis was Cartesian dualism, the traditional religious or spiritual view that mind exists in a separate realm that is inaccessible to science. Our inability to detect spiritual entities is not due to any limitations of our detector technology but to the fact that spiritual entities are impossible in principle to detect by physical means. Cartesian dualism is a minority position in contemporary philosophy, at least as a scientific theory of mind, and for very good reason. The chief objection to this kind of dualism is Occam's razor: It is more parsimonious to posit a single universe with one set of physical laws rather than two radically dissimilar parallel universes composed of dissimilar substances and following dissimilar laws, making tenuous contact with each other nowhere else but within a living conscious brain. But if mind and matter come into causal contact, as they clearly do in both sensory and motor function, then surely they must be different parts of one and the same physical universe. There is another, still more serious objection to Cartesian dualism than the issue of parsimony. Since the experiential, or spiritual component of the theory is in principle inaccessible to science, that portion of the theory can be neither confirmed nor refuted. This places the spiritual component of Cartesian dualism beyond the bounds of science and firmly in the realm of religious belief.

A more sophisticated halfway epistemology is seen in the philosophy of critical realism (Broad 1925; Drake et al. 1920; Russell 1921; Sellars 1916). Critical realists avoid religious explanations involving God or spirits, but their concept of conscious experience nevertheless preserves some of the mystery of Cartesian dualism. Critical realists acknowledge that perception is not direct, but instead, is mediated by an intermediate representational entity called *sense-data.* However, critical realists insist that sense-data are

> particular existents of a peculiar kind; they are not physical, . . . and there is no reason to suppose that they are either states of mind or existentially mind-dependent. In having spatial characteristics . . . they resemble physical objects . . . but in their privacy and their dependence on the body . . . of the observer they are more like mental states. (Broad 1925, p. 181)

As with the spirit world of the Cartesian view, sense data and the space in which they are observed are not just difficult to detect, but they are in principle beyond scientific scrutiny.

There is some debate among critical realists over the ontology of conscious experience. In a book on critical realism by a consortium of authors (Drake et al. 1920), Lovejoy, Pratt, and Sellars claimed that the sensa are completely "the character of the mental existent . . . although its existence is not given" (pp. 20–21), while Drake, Rogers, Santayana, and Strong agreed that the data are characteristic of the apprehended object, although "the datum is, *qua* datum, a mere essence, an inputed but not necessarily actual existent. It may or may not have existence" (Drake 1920 in Drake et al. 1920, pp. 20–21, footnote). So the critical realists solved the epistemological problem by defining a unique kind of existent that is experienced, but that does not or may not actually exist. This is a peculiar inversion of the true epistemological situation because, in fact, sense data, or the raw material of conscious experience, are the *only* thing we can know with any real certainty to actually exist. All else, including the entire physical world known to science, is informed conjecture based on that experience.

A more modern reformulation of this muddled epistemology is seen in Davidson's (1970) anomalous monist thesis. Davidson suggested that the mental domain, on the basis of its essential anomalousness and normativity, cannot be the object of serious scientific investigation because the mental is on a wholly different plane from the physical. This argument sounds like the metaphysical dualism of Descartes which disconnects mind from brain entirely, except that Davidson qualified his theory with the monistic proviso that every mental event is connected with specific physical events (in the brain), although there are no laws connecting mental kinds with physical kinds, and this presumably rescues the thesis from metaphysical dualism. Kim (1998) pointed out, however, that this is a negative thesis, for it tells us only how the mental is *not* related to the physical, it says nothing about how they *are* related. As such, this is more an article of faith rather than a real theory of any sort, and in the context of the history of the epistemological debate this can be seen as a last desperate attempt to rescue naïve realism from its own logical contradictions. This kind of physicalism has been appropriately dubbed "token physicalism," for it is indeed a token admission of the undeniable link between mind and the physical brain, without admitting to any of its very significant implications.

To rationalize this view of the mind-brain relation, Davidson (1970) introduced the peculiar notion of *supervenience,* a one-way asymmetrical relation between mind and brain which makes the mind dependent on the brain but forever closes the possibility of phenomenological observation of brain states. As in the case of Cartesian dualism, there are two key objections to this argument. In the first place, the disconnection between the experiential mind and the physical brain is itself merely a hypothesis whose truth remains to be demonstrated. It is at least equally likely *prima facie* that the mind does not supervene on the brain, but rather that the mind is identically equal to the functioning of the physical brain. In fact, this is by far the more parsimonious explanation because it invokes a single explanans, the physical brain, to account for the properties of both mind and brain. After all, physical damage to the brain can result in profound changes in the mind, not just in the information content of the mind or in observed behavior but in the experiential or "what it is like" aspect of conscious experience. The simplest explanation therefore is that consciousness is a physical process taking place in the physical brain, which is why it is altered by physical changes to the physical brain.

But the problem of supervenience is more serious than just the argument of parsimony. If the properties of mind are indeed disconnected from the properties of the physical brain, this would leave the mental domain completely disconnected from the world of reality known to science, as what Feigl (1958) has called a "nomological dangler." If the properties of mind are not determined by the properties of the physical brain, what is it that determines the properties of the mind? For example, phenomenal color experience has been shown to be reducible to the three dimensions of hue, intensity, and saturation. Physical light is not restricted to these three dimensions; the spectrum of a typical sample of colored light contains a separate and distinct magnitude for every spectral frequency of the light, an essentially infinite-dimensional space that is immeasurably greater in information content than the three dimensions of phenomenal color experience. In answer to Koffka's (1935) classical question "Why do things look as they do?", the answer is clearly not "Because they are what they are." That answer is clearly false in the case of color perception, as well as in the cases of visual illusions, dreams, and hallucinations. We now know that the dimensionality of color experience relates directly to the physiology of color vision; it relates to the fact that there are three different cone types in the human retina and it relates to the opponent color process representation in the visual cortex. The dimensions of color experience therefore are not totally disconnected from the properties of the physical brain, as suggested by Davidson (1970), but in fact phenomenal color experience tells us something very specific about the properties of the representation of color in the physical brain. And the same argument holds for spatial vision, for there are a number of prominent distortions of phenomenal space which clearly indicate that phenomenal space is ontologically distinct from the physical space known to science, as will be discussed in section 6.3.

Daniel Dennett (1991) promoted a similar halfway epistemology by drawing a distinction between the neural *vehicles* of mental representation and the phenomenal *contents* of those vehicles. Dennett opened the epistemological crack by claiming that the phenomenal contents do not necessarily bear any similarity whatsoever to the neural vehicles by which they are encoded in the brain. This actually goes beyond Davidson's supervenience because, according to Davidson (1970), mental events that are distinct phenomenally must also be distinct neurophysiologically. This is tantamount to saying that the dimensions of conscious experience cannot be any less than the dimensions of the corresponding neurophysiological state. Dennett effectively removed this limitation by suggesting that even the dimensionality of the phenomenal contents need not match that of the neural vehicles. And into that epistemological crack, Dennett slipped the entire world of conscious experience like a magical disappearing act, where it is experienced but does not actually exist. By the very fact that conscious experience, as conceived by Dennett, is in principle undetectable by scientific means, this concept of consciousness becomes a religious rather than a scientific hypothesis, whose existence can be neither confirmed nor refuted by scientific means. In fact, Dennett even suggested that there is actually no such thing as consciousness per se, and that belief in consciousness is akin to belief in some kind of mythical nonexistent deity (Dennett 1981). This argument of course is only intelligible from a naïve realist perspective, by which the sense-data of conscious experience, so plainly manifest to one and all, are misidentified as the external world itself rather than as something going on in the physical brain.

Another modern theorist, Max Velmans (1990), revived an ancient notion of perception as something projecting out of the head into the world, as proposed by Empedocles and promoted by Malebranche. But Velmans refined this ancient notion with the critical realist proviso that nothing physical actually gets projected from the head; the only thing that is projected is conscious experience, a subjective quality that is undetectable externally by scientific means. But again, as with critical realism, the problem with this notion is that the sense-data that are experienced to exist do not exist in any true physical sense, and therefore the projected entity in Velman's theory is a spiritual entity to be believed in (for those who are so inclined), rather than anything knowable by, or demonstrable to, science. Velmans drew the analogy of a videotape recording that carries the information of a dynamic pictorial scene, expressed in a highly compressed and nonspatial representation, as patterns of magnetic fields on the tape. There is no resemblance or isomorphism between the magnetic tape and the images that it encodes, except for its information content. However, the only reason the videotape even represents a visual scene is because of the existence of a video technology that is capable of reading the magnetic information from the tape and sweeping it out as a spatial image on a video monitor or television screen, where each pixel appears in its proper place in the image. If that equipment did not exist, there would be no images as such on the videotape. But if video technology is to serve as an analogy for spatial representation in the brain, the key question is whether the brain encodes that pictorial information exclusively in abstract compressed form like the magnetic patterns on the tape, or whether the brain reads those compressed signals and projects them as an actual spatial image somewhere in the brain like a television monitor, whenever we have a visuospatial experience. If it is the former, then sense-data are experienced but do not actually exist as a scientific entity, so the spatial image we see is a complete illusion, which, again, is an inversion of the true epistemology. If it is the latter, then there are actual "pictures in the head," a notion that Velmans emphatically rejected.

In fact, the only epistemology that is consistent with the modern materialistic world view is an identity theory (Feigl 1958; Russell 1927) whereby mind is identically equal to physical patterns of energy in the physical brain. To claim otherwise is to relegate the elaborate structure of conscious experience to a mystical state beyond the bounds of science. The dimensions of conscious experience, such as phenomenal color and phenomenal space, are a direct manifestation of certain physical states of our physical brain. The only right answer to Koffka's question (Koffka 1935) is that things appear as they do because that is the way the world is represented in the neurophysiological mechanism of our physical brain. In principle, therefore, the world of conscious experience is accessible to scientific scrutiny after all, both internally through introspection and externally through neurophysiological recording. And introspection is as valid a method of investigation as is neurophysiology, just as in the case of color experience. Of course, the mind can be expected to appear quite different from these two perspectives, just as the data in a computer memory chip appear quite different when examined internally by data access as opposed to externally by electrical probes. But the one quantity that is preserved across the mind/brain barrier is information content, and therefore that quantity can help to identify the neurophysiological mechanism or principle in the brain whose dimensionality, or information content, matches the observed dimensions of conscious experience.

2.4. Selection from incredible alternatives

We are left therefore with three alternatives, each of which appears to be absolutely incredible. Contemporary neuroscience seems to take something of an equivocal position on this issue, recognizing the epistemological limitations of the direct realist view and of the projection hypothesis, yet being unable to account for the incredible properties suggested by the indirect realist view. However, one of these three alternatives simply must be true, to the exclusion of the other two. And the issue is by no means inconsequential, for these opposing views suggest very different ideas of the function of visual processing, or what all that neural wetware is supposed to actually *do.* Therefore, it is of central importance for psychology to address this issue head-on, and to determine which of these competing hypotheses reflects the truth of visual processing. Until this most central issue is resolved definitively, psychology is condemned to remain in what Kuhn (1970) calls a *pre-paradigmatic* state, with different camps arguing at cross-purposes due to a lack of consensus on the foundational assumptions and methodologies of the science. Psychology is, after all, the science of the *psyche,* the subjective side of the mind/brain barrier, and neurophysiology only enters the picture to provide a physical substrate for mind. Therefore, it is of vital importance to reach a consensus on the nature of the *explanandum* of psychology before we can attempt an *explanans.* In particular, we must decide whether the vivid spatial structure of the surrounding world of visual experience is an integral part of the psyche and thus within the explanandum of psychology, or whether it is the external world itself, as it appears to be naively, and thus in the province of physics rather than of psychology.

The problem with the direct realist view is of an epistemological nature, and is therefore a more fundamental objection; for direct realism, as defended by Gibson (1979), is nothing short of magical – that we can see the world out beyond the sensory surface. The projection theory has a similar epistemological problem and is equally magical and mysterious, suggesting as it does that neural processes in our brain are somehow also out in the world. Both of these paradigms have difficulty with the phenomena of dreams and hallucinations (Revonsuo 1995), which present the same kind of phenomenal experience as spatial vision, except independent of the external world in which that perception is supposed to occur in normal vision. It is the implicit or explicit acceptance of this naive concept of perception which has led many to conclude that consciousness is deeply mysterious and forever beyond human comprehension. For example, Searle (1992, p. 96) contended that consciousness is impossible to observe, for when we attempt to observe consciousness we see nothing but whatever it is that we are conscious of; there is no distinction between the observation and the thing observed.

On the other hand, the problem with the indirect realist view is more of a technological or computational limitation, for we cannot imagine how contemporary concepts of neurocomputation, or even of artificial computation for that matter, can account for the properties of perception as observed in visual consciousness. It is clear, however, that the most fundamental principles of neural computation and representation remain to be discovered, and therefore we cannot allow our currently limited notions of neurocomputation to constrain our observations of the nature of visual consciousness. The phenomena of dreams and hallucinations clearly demonstrate that the brain is capable of generating vivid spatial percepts of a surrounding world independent of that external world, and that capacity must be a property of the physical mechanism of the brain. Normal conscious perception can therefore be characterized as a guided hallucination (Revonsuo 1995), which is as much a matter of active construction as it is of passive detection. If we accept the truth of indirect realism, this immediately disposes of at least one mysterious or miraculous component of consciousness, which is its unobservability. Consciousness is indeed observable, contrary to Searle's contention, because the objects of experience are first and foremost the product or "output" of consciousness, and only in secondary fashion are they also representative of objects in the external world. Searle's (1992) difficulty in observing consciousness is analogous to saying that you cannot see the moving patterns of glowing phosphor on your television screen, all you see is the ball game that is showing on that screen. The indirect realist view of television is that what you are seeing is first and foremost glowing phosphor patterns on a glass screen, and only in secondary fashion are those moving images also representative of the remote ball game.

The choice therefore is between accepting a magical mysterious account of perception and consciousness that seems impossible in principle to implement in any artificial vision system, or facing the seemingly incredible truth that the world we perceive around us is indeed an internal data structure within our physical brain (Lehar 2003b). The principal focus of neurophysiology should now be to identify the operational principles behind the three-dimensional volumetric imaging mechanism in the brain, the mechanism responsible for generating the solid stable world of visual experience that we observe to surround us in conscious experience.

3. Problems in modeling perception

The computational modeling of perceptual processes is a formidable undertaking. But the problem is exacerbated by the fact that a neural network model of perception attempts to model two entities simultaneously: the subjective experience of perception and the neurophysiological mechanism by which that experience is generated in the brain. The chief problem with this approach is that our knowledge of neurophysiological principles is known to be incomplete. We do not understand the computational functionality of even the simplest neural systems. For example, the lowly house fly, with its tiny pinpoint of a brain, seems to thumb its nose at our lofty algorithms and complex computational models as it dodges effortlessly between the tangled branches of a shrub in dappled sunlight, compensating for gusty cross-winds to avoid colliding with the branches. This remarkable performance by this lowly creature far exceeds the performance of our most powerful computer algorithms and our most sophisticated neural network models of human perception. In fact, the "dirty little secret" of neuroscience, as Searle (1997, p. 198) called it, is that we have no idea what the right level of analysis of the brain should be because there is no universally accepted theory of how the brain actually codes perceptual or experiential information. The epistemological question highlights this uncertainty, for it shows that there is not much consensus on whether the world of conscious experience is even explicitly represented in the brain at all, the majority view being, apparently, that it is not. Palmer (1999) went even further, saying that "to this writer's knowledge, no one has ever suggested any theory that the scientific community regards as giving even a remotely plausible causal account of how experience arises from neural events." Without this key piece of knowledge, how can we even begin to model the computational processes of perception in neurophysiological terms?

One approach is to begin with the neurophysiology of the brain and attempt to discover what it is computing at the local level of the individual neuron, the elemental building block of the nervous system. The fruit of this branch of investigation is neural network theory. But it is unclear whether neural network theory offers an adequate characterization of the actual processing going on in the brain, or whether it is asking too much of simple integrate-and-fire elements, no matter how cleverly connected in patterns of synaptic connections, to provide anything like an adequate account of the observed properties of conscious experience. Churchland (1984) argued in the affirmative, that we do have enough knowledge of the principles of neurocomputation to begin to propose realistic models of perceptual processing. Palmer (1992) and Opie (1999) presented dynamic neural network models of Gestalt phenomena, such as the perceptual grouping of triangles, showing how the dynamics of perceptual phenomena can be modeled by a dynamic neural network model. But those models are proposed in the abstract, presenting general principles rather than complete and detailed models of specific perceptual phenomena expressed as sense-data. For example, Palmer (1992) discussed the perceptual experience of an equilateral triangle, perceived as an arrow pointing in one of three directions. Palmer modeled this perceptual phenomenon as a competition between three dynamic neural network nodes in a mutually inhibitory relationship, resulting in a "winner-take-all" behavior. Although this model is compelling as a demonstration of Gestalt principles in a neural network model, Palmer left out the most difficult part of the problem, which is not just the competition between three alternative percepts but the perceptual representation of the percept itself. The perceptual experience of a triangle cannot be reduced to just three phenomenal values but is observed as a fully reified triangular structure that spans a specific portion of perceived space. This sense-data component of the phenomenal experience is very much more difficult to account for in neural network terms.

In recent decades a number of attempts have been made to quantify the sense-data of visual consciousness in computational models (see Lesher 1995, for a review). Zucker et al. (1988) presented a model of curve completion that accounts for the emergent nature of perceptual processing by incorporating a feedback loop in which local feature detectors tuned to detect oriented edges feed up to global curvature detector cells, and those cells in turn feed back down to the local edge level to fill in missing pieces of the global curve. A similar bottom-up/top-down feedback is given in Grossberg and Mingolla's (1985) visual model to account for boundary completion in illusory figures like the Kanizsa square by generating an explicit line of neural activation along the illusory contour. An extension of that model (Grossberg & Todoroviçz 1988) accounted for the filling-in of the surface brightness percept in the Kanizsa figure, with an explicit diffusion of neural activation within the region of the illusory surface. These models have had a significant impact on the discussion of the nature of visual illusions because they highlight the fact that illusory features, like the illusory surface of a Kanizsa figure, are observed as extended image-like data structures, and therefore a complete model of the phenomenon must also produce a fully reified image-like spatial structure as its output. In fact, Grossberg's concept of visual reification in his Boundary Contour System (Grossberg & Mingolla 1985) and Feature Contour System (Grossberg & Todoroviç 1988) were the original inspiration behind the perceptual modeling proposed in the present hypothesis.

Although these models finally offer a reasonable account of perceptual experience (in two dimensions), they also demonstrate the profound limitations of a neural network architecture for perceptual representation because neural network theory is no different in principle than a template theory (Lehar 2003a), a concept whose limitations are well known. Grossberg and Mingolla (1985) account for collinear illusory contour completion by way of specialized elongated receptive fields, tuned to detect and enhance collinearity. This concept works well enough for simple collinear boundary completion (as long as it remains restricted to two dimensions), but any attempt to extend this model to higher order perceptual processing runs headlong into a combinatorial explosion in required receptive fields (Lehar 2003a). For example, perceptual completion is observed not only for collinear alignments but it can also define illusory vertices composed of two, three, or more edges that meet at a vertex (Lehar 2003a). Grossberg himself proposed an extension to his model equipped with "corner detector" receptive fields (Grossberg & Mingolla 1985), although this line of thought was subsequently quietly abandoned because, just as with the cells that perform collinear completion, the corner detectors would have to be provided at every location and every orientation across the visual field. To extend the model to account for T, V, Y, and X intersections, specialized receptive fields would have to be provided for each of those features at every location and at every orientation across the visual field. This combinatorial explosion in the required number of specialized receptive fields does not bode well for neural network theory as a general principle of neurocomputation.

The most serious limitation of Grossberg's approach to perception is that, curiously, Grossberg and his colleagues did not extend their logic to the issue of three-dimensional spatial perception. In going from two dimensions to three, Grossberg no longer advocated explicit spatial filling-in, but instead represented the depth dimension by binocular disparity, using left and right eye image pairs (Grossberg 1987; 1990; 1994). Although a stereo pair does encode depth information, it does not do so in a volumetric manner because it can only encode one depth or disparity value for every (x,y) point on the image. This makes it impossible for Grossberg's model to represent transparency with multiple depth values at a single (x,y) location, or to represent the experience of empty space between the observer and a visible object. Moreover, it precludes the kind of volumetric filling-in required to account, for example, for the three-dimensional version of the Ehrenstein illusion constructed of a set of rods arranged radially around a circular void (Ware & Kennedy 1978). The filling-in processes in this illusion take place through the depth dimension, which produces an illusory percept of a glowing disk, hanging in space, as a volumetric spatial structure. If Grossberg's argument for explicit filling-in of the two-dimensional illusions is at all valid, then that argument should apply equally to volumetric filling-in also.

The reason Grossberg declined to extend his model into the third dimension is neurophysiologically motivated. For although Grossberg's model is a de facto perceptual model, it is actually presented as a neural network model; that is, the computational units of the model represent actual neurons in the brain rather than perceptual entities. And this highlights the problem of perceptual modeling in neural network terms, for whenever there is a conflict between the perceptual phenomenon and our current understanding of neurophysiological principles, there is then a conflict between the neural and the perceptual models of the phenomenon. In this case the percept is clearly volumetric, but the corresponding cortical neurophysiology is assumed to be two-dimensional. Another reason Grossberg was reluctant to extend his model into the third dimension is that, even for simple collinear completion, such an approach would require a volumetric block of neural elements each equipped with elongated receptive fields; and those fields must be replicated at every orientation in three dimensions and at every volumetric location across the entire volume of the perceptual representation – a notion that seems too implausible to contemplate, let alone the idea of T, V, Y, and X intersections defined in three dimensions. But until a mapping has been established between the conscious experience and the corresponding neurophysiological state, there is no way to verify whether the model has correctly replicated the psychophysical data. Because these models straddle the mind/brain barrier, they run headlong into the issue that Chalmers (1995) dubbed the "hard problem" of consciousness. Simply stated, even if we were to discover the exact neurophysiological correlates of conscious experience, there would always remain a final explanatory gap between the physiological and the phenomenal levels of description. For example, if the activation of a particular cell in the brain were found to correlate with the experience of red at some point in the visual field, there would remain a vivid subjective quality, or *quale,* to the experience of red that is not in any way identical to any externally observable physical variable such as the electrical activity of a cell. In other words, there is a subjective experiential component of perception that can never be captured in a model expressed in objective neurophysiological terms.

Even more problematic for neural models of perception is the question of whether perceptual information is expressed neurophysiologically in *explicit* or *implicit* form. For example, Dennett (1992) argued that the perceptual experience of a filled-in colored surface is encoded in more abstracted form in the brain, in the manner of an edge image that records only the transitions along image edges. Support for this concept is seen in the retinal ganglion cells that respond only along spatial or temporal discontinuities in the retinal image and produce no response within regions of uniform color or brightness. This concept also appears to make sense from an information-theoretic standpoint, for uniform regions of color represent redundant information that can be compressed to a single value, as is the practice in image compression algorithms. These kinds of theoretical difficulties have led many neuroscientists to simply ignore the conscious experience and to focus instead on the hard evidence of the neurophysiological properties of the brain.

4. A perceptual modeling approach

The quantification of conscious experience is not quite as hopeless as it might seem. Nagel (1974) suggested that we set aside temporarily the relation between mind and brain and devise a new method of *objective phenomenology* – in other words, quantify the structural features of the subjective experience in objective terms without committing to any particular neurophysiological theory of perceptual representation. For example, if we quantify the experience of vision as a three-dimensional data structure, like a model of volumes and surfaces in a surrounding space to a certain perceptual resolution, this description could be meaningful even to a congenitally blind person or to an alien creature who had never personally experienced the phenomenon of human vision. Although this description could never capture everything of that experience, such as the qualia of color experience, it would at least capture the structural characteristics of that subjective experience in an objective form that would be comprehensible to beings incapable of having those experiences.

Chalmers (1995) extended this line of reasoning with the observation that the subjective experience and its corresponding neurophysiological state carry the same information content. On that ground, Chalmers proposed a principle of *structural coherence* between the structure of phenomenal experience and the structure of objectively reportable awareness, to reflect the central fact that consciousness and physiology do not float free of one another but cohere in an intimate way. In essence this is a restatement of the Gestalt principle of *isomorphism,* of which more in section 5. The connecting link between mind and brain therefore is *information* in information-theoretic terms (Shannon 1948) because the concept of information is defined at a sufficiently high level of abstraction to be independent of any particular physical realization, and yet it is specified sufficiently to be measurable in any physical system given that the coding scheme is known. A similar argument was made by Clark (1993, p. 50). Chalmers moderated his claim of the principle of structural coherence by stating that it is a hypothesis that is "extremely speculative." However, the principle is actually solidly grounded epistemologically because the alternative is untenable. If we accept the fact that the physical states of the brain correlate directly with conscious experience, then the claim that conscious experience contains more explicit information than does the physiological state on which it was based amounts to a kind of dualism that would necessarily involve some kind of nonphysical "mind stuff" to encode the excess information observed in experience that is not encoded by the physical state. Some theorists have even proposed a kind of hidden dimension of physical reality to house the unaccounted information in conscious experience (Harrison 1989; Smythies 1994).

The philosophical problems inherent in neural network models of perceptual experience can be avoided by proposing a *perceptual modeling* approach (Lehar 2003b), which models the conscious experience directly in the subjective variables of perceived color, shape, and motion, as opposed to *neural modeling*, where the conscious experience is modeled in the neurophysiological variables of neural activations or spiking frequencies, or the like. The variables encoded in the perceptual model therefore correspond to what philosophers call the sense-data or primitives of raw conscious experience, except that these variables are not supposed to *be* the sense-data themselves, they merely represent the value or magnitude of the sense-data they are defined to represent. In essence this amounts to modeling the information content of subjective experience, which is the quantity that is common between mind and brain, thus allowing an objectively quantified description of a subjective experience. In fact, this approach is exactly the concept behind the description of phenomenal color space in the dimensions of hue, intensity, and saturation, as seen in the CIE (Commission Internationale L'Eclairage) chromaticity space. The geometrical dimensions of that space have been tailored to match the properties of the subjective experience of color as measured psychophysically, expressed in terms that are agnostic to any particular neurophysiological theory of color representation.

Clark (1993) presented a systematic description of other sensory qualities in quantitative terms, based on this same concept of "objective phenomenology." The thorny issue of the hard problem of consciousness is thus neatly sidestepped because the perceptual model remains safely on the *subjective* side of the mind/brain barrier, and therefore the variables expressed in the model refer explicitly to subjective qualia rather than to neurophysiological states of the brain. The problems of explicit versus implicit representation are also neatly circumvented because those issues pertain to the relation between mind and brain and so do not apply to a model that does not straddle the mind/brain barrier. For example, the subjective experience of a Necker cube is of a solid three-dimensional structure, and for that reason the perceptual model of that experience should also be an explicit three-dimensional structure. The spontaneous reversals of the Necker cube, on the other hand, are experienced as a dynamic process, and on that ground should be represented in the perceptual model as a dynamic process – that is, as a literal reversal of the solid three-dimensional structure. The issues of whether a perceived structure can be encoded neurophysiologically as a process or whether a perceived process can be encoded as a structure are therefore irrelevant to the perceptual model, which by definition models a perceived structure as a structure, and a perceived process as a process.

This is of course only an interim solution, for eventually the neurophysiological basis of conscious experience must also be identified; nevertheless, the perceptual model does offer objective information about the informational content encoded in the physical mechanism of the brain. This is a necessary prerequisite to a search for the neurophysiological basis of conscious experience, for we must clearly circumscribe that which we are to explain before we can attempt an explanation of it. This approach has served psychology well in the past, particularly in the field of color perception where the quantification of the dimensions of color experience led directly to great advances in our understanding of the neurophysiology of color vision. The failure to quantify the dimensions of spatial experience has been responsible for decades of futile debate about its neurophysiological correlates. I will show that application of this perceptual modeling approach to the realm of spatial vision opens a wide chasm between phenomenology and contemporary concepts of neurocomputation and thereby offers a valuable check on theories of perception based principally on neurophysiological concepts.

5. The Gestalt principle of isomorphism

The Gestalt principle of isomorphism represents a subtle but significant extension to Müller's psychophysical postulate and to Chalmers's principle of structural coherence. In the case of structured experience, equal dimensionality between the subjective experience and its neurophysiological correlate implies similarity of structure or form. For example, the percept of a filled-in colored surface, whether real or illusory, encodes a separate and distinct experience of color at every distinct spatial location within that surface to a particular resolution. Each point of that surface is not experienced in isolation but in its proper spatial relation to every other point in the perceived surface. In other words, the experience is extended in at least two dimensions, and therefore the neurophysiological correlate of that experience must also encode at least two dimensions of perceptual information. The mapping of phenomenal color space was established by the method of *multidimensional scaling* (Coren et al. 1994, p. 57) in which color values are ordered in psychophysical studies on the basis of their perceived similarity, to determine which colors are judged to be nearest to each other or which colors are judged to be between which other colors in phenomenal color space. A similar procedure could just as well be applied to spatial perception to determine the mapping of phenomenal space. If two points in a perceived surface are judged psychophysically to be nearer to each other when they are actually nearer and farther when they are actually farther, and if other spatial relations such as betweenness are also preserved phenomenally, then direct evidence is thereby provided that phenomenal space is mapped in a spatial representation that preserves those spatial relations in the stimulus. The outcome of this proposed experiment is so obvious it need hardly be performed. And yet its implication – that our phenomenal representation of space is spatially mapped – is not often considered in contemporary theories of spatial representation.

5.1. Structural versus functional isomorphism

The isomorphism required by Gestalt theory is not a strict *structural* isomorphism, a literal isomorphism in the physical structure of the representation, but rather, it is merely a *functional* isomorphism, a behavior of the system *as if* it were physically isomorphic (Köhler 1969, p. 92). This is because the exact geometrical configuration of perceptual storage in the brain cannot be observed phenomenologically any more than the configuration of silicon chips on a memory card can be determined by software examination of the data stored within those chips. Nevertheless, the mapping between the stored perceptual image and the corresponding spatial percept must be preserved, as in the case of the digital image, so that every stored color value is meaningfully related to its rightful place in the spatial percept.

The distinction between structural and functional isomorphism can be clarified with a specific example. Consider the spatial percept of a block resting on a surface, depicted schematically in Figure 1A. The information content of this perceptual experience can be captured in a painted cardboard model built explicitly like Figure 1A, with explicit volumes, bounded by colored surfaces, embedded in a spatial void. Because perceptual resolution is finite, the model should also be considered only to a finite resolution; that is, the infinite subdivision of the continuous space of the actual model world is not considered to be part of the model, which can only validly represent subdivision of space to the resolution limit of perception. The same perceptual information can also be captured in quantized or digital form in a volumetric or voxel (volume-pixel) image in which each voxel represents a finite volume of the corresponding perceptual experience, as long as the resolution of this representation matches the spatial resolution of the percept itself; in other words, the size of the voxels should match the smallest perceivable feature in the corresponding spatial percept. Both the painted cardboard model and its quantized voxel equivalent are structurally or topographically isomorphic with the corresponding percept; they have the same information content as the spatial percept that they represent.

Consider now the flattened representation depicted in Figure 1B, which is identical to the model in Figure 1A except that the depth dimension is compressed relative to the other two dimensions, like a bas-relief. If the defined scale of the model (the length in the representation relative to the length that it represents) is also correspondingly compressed, as suggested by the compressed gridlines in the figure, then this model is also isomorphic with the perceptual experience of Figure 1A. In other words the flattening of the depth dimension is not really registered in the model because the perceived cube spans the same number of gridlines in Figure 1B (in all three dimensions) as it does in Figure 1A, and therefore this flattened model encodes a non-flattened perceptual experience. Though this model is now no longer structurally isomorphic with the original perceptual experience, it does remain topologically isomorphic, preserving neighborhood relations, as well as betweenness, and so forth. In a mathematical system with infinite resolution, this model would encode the same information as the one in Figure 1A. However in a real physical representation there is always some limit to the resolution of the system, or how much information can be stored in each unit distance in the model itself. In a representational system with finite resolution, therefore, the depth information in Figure 1B would necessarily be encoded at a lower resolution than that in the other two dimensions. If our own perceptual apparatus employed this kind of representation, this flattening would not be experienced directly; the only manifestation of the flattening of the representation would be a reduction in the resolution of perceived depth relative to the other two dimensions, making it more difficult to distinguish differences of perceived depth than differences of perceived height and width.

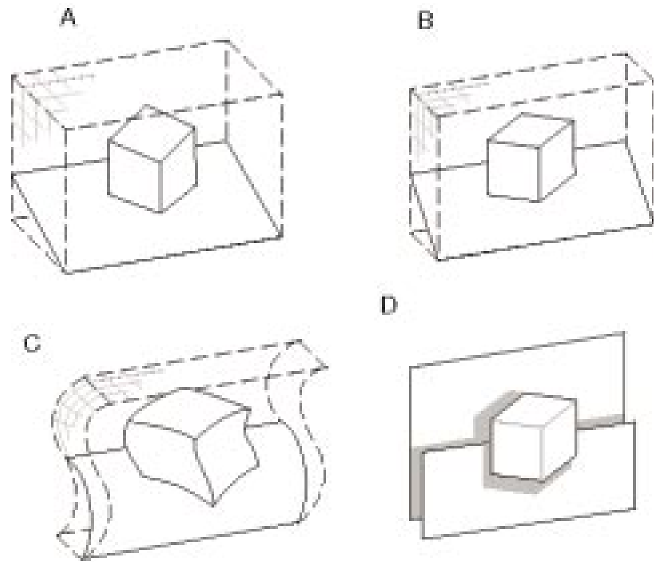

Figure 1. **A.** A volumetric spatial model, for example built of painted cardboard surfaces, is structurally isomorphic with a perceptual experience of a block resting on a surface if it has the same information content. **B.** If the model is compressed in one dimension relative to the other two, the model can still be isomorphic with the original percept if the representational scale of the model (indicated by the shaded gridlines) is also correspondingly compressed, although this is no longer a structural isomorphism but merely a topological isomorphism. **C.** The model can even be warped like the gyri and sulci of the cortical surface and remain isomorphic with the original percept. **D.** But a model composed of a small number of discrete depth planes is not isomorphic with the original percept because it no longer encodes the same information content.

Consider now the warped model depicted in Figure 1C, which is like the flattened model of Figure 1B with a wavy distortion applied, as if warped like the gyri and sulci of the cortical surface. This warped representation is also isomorphic with the perceptual experience it represents for it encodes the same information content as the flattened space in Figure 1B, although again this is a topological rather than a topographical isomorphism. The warping of this space would not be apparent to the percipient because the very definition of straightness is warped along with the space itself, as suggested by the warped gridlines in the figure. In contrast, consider the flattened representation depicted in Figure 1D, where the perceptual representation has been segmented into discrete depth planes that distinguish only foreground from background objects. This model is no longer isomorphic with the perceptual experience it supposedly represents because, unlike this model, the perceptual experience manifests a specific and distinct depth value for every point in each of the surfaces of the percept. Furthermore, the perceptual experience manifests an experience of empty space surrounding the perceived objects, every point of which is experienced simultaneously and in parallel as a volumetric continuum of a certain spatial resolution, whereas the model depicted in Figure 1D encodes only a small number of discrete depth planes. This kind of model therefore is inadequate as a perceptual model of the information content of conscious experience because the dimensions of its representation are less than the dimensions of the experience it attempts to model.

A functional isomorphism must also preserve the functional transformations observed in perception, and the exact requirements for a functional isomorphism depend on the functionality in question. For example, when a colored surface is perceived to translate coherently across perceived space, the corresponding color values in the perceptual representation of that surface must also translate coherently through the perceptual map. If that memory is discontinuous, like a digital image distributed across separate memory chips on a printed circuit board, then the perceptual representation of that moving surface must jump seamlessly across those discontinuities in order to account for the subjective experience of a continuous translation across the visual field. In other words, a functional isomorphism requires a functional connectivity in the representation, as if a structurally isomorphic memory were warped, distorted, or fragmented, but at the same time, the functional connectivity between its component parts were preserved. Consider a representational mechanism, such as that shown in Figure 1A, equipped with additional computational hardware capable of performing spatial transformations on the volumetric image in the representation. The representational mechanism might be equipped with functions that could rotate, translate, and scale the spatial pattern in the representation on demand. This representation would thereby be invariant to rotation, translation, and scale, because the spatial pattern of the block itself would be encoded independent of its rotation, translation, and scale. The fact that an object in perception maintains its structural integrity and recognized identity despite rotation, translation, and scaling by perspective is clear evidence for this kind of invariance in human perception and recognition. If the warped model shown in Figure 1C were equipped with these same transformational functions, the warped representation would also be functionally isomorphic with the non-warped representation as long as those transformations were performed correctly with respect to the warped geometry of that space.

A functional isomorphism is even possible for a representation that is fragmented into separate pieces, if those pieces are wired together in such a way that they continue to perform the spatial transformations exactly as in the corresponding undistorted mechanism. A functional isomorphism can even survive in a volumetric representation whose individual elements or voxels are scrambled randomly across space, if the functional connections between those elements are preserved through the scrambling. The result is a representation that is neither topographically nor topologically isomorphic with the perceptual experience it represents. However, it does remains a volumetric representation, with an explicit encoding of each point in the represented space to a particular spatial resolution, and it remains functionally isomorphic with the spatial experience that it represents, capable of performing coherent rotation, translation, and scaling transformations of the perceptual structures expressed in the representation.

An explicit volumetric spatial representation capable of spatial transformation functions, as described above, is more efficiently implemented in either a topographically isomorphic form or a topologically isomorphic form, which require shorter and more orderly connections between adjacent elements in the representation. However, the argument for structural or topological isomorphism is an argument of representational efficiency and simplicity, rather than of logical necessity. On the other hand, a *functional* isomorphism *is* strictly required in order to account for the properties of the perceptual world as observed subjectively. The volumetric structure of visual consciousness and perceptual invariance to rotation, translation, and scale offer direct and concrete evidence for an explicit volumetric spatial representation in the brain, which is at least functionally isomorphic with the corresponding spatial experience.

A neurophysiological model of perceptual processing and representation should concern itself with the actual mechanism in the brain. In the case of a distorted representation (as in Fig. 1C), the warping of that perceptual map would be a significant feature of the model. A perceptual model, on the other hand, is concerned with the structure of the percept itself, independent of any warping of the representational manifold. Even for a representation that is functionally but not structurally isomorphic, a description of the functional transformations performed in that representation is most simply expressed in a structurally isomorphic form, just as a panning or scrolling function in image data is most simply expressed as a spatial shifting of image data even when that shifting is actually performed in hardware in a non-isomorphic memory array. For that reason, the functional operation of a warped mechanism like Figure 1C is most simply described as the operation of the functionally equivalent undistorted mechanism in Figure 1A. In the present discussion, therefore, our concern will be chiefly with the *functional architecture* of perception, a description of the spatial transformations observed in perception, whatever form those transformations might take in the physical brain. And those transformations are most simply described as if taking place in a physically isomorphic space.

In the discussion that follows, the terminology "spatial representation," "data expressed in spatial form," "literal volumetric replica of the world inside your head," "three-dimensional pattern of opaque-state units," "explicit three-dimensional replica of the surface," and "volumetric spatial medium," will refer not to a topographically isomorphic model of space, as suggested in Figure 1A, but to a functionally isomorphic model of space like the warped model in Figure 1C, in which the explicit volumetric representation is possibly warped and distorted but still encodes an explicit value for every volumetric point in perceived space as well as the neighborhood relations between those values. This is in contrast to the more commonly assumed flattened or abstracted cortical representation depicted in Figure 1D, where the volumetric mapping is no longer preserved.

5.2. Second-order, complementary, and other paramorphisms

The issue of isomorphism is so profoundly problematic for theories of perceptual representation that theorists have gone to no end of trouble in an effort to dispel the issue and to argue that isomorphism is not actually necessary. A careful examination of these proposals, however, reveals the naïve realist assumptions on which they are founded.

Shepard and Chipman (1970) argued that when we perceive a square, for example, there is no need for an internal perceptual replica of that square in the brain of the percipient. They argued that we learn the appropriate use of words such as "square" from a verbal community that has access only to the public object and not to any such private image. If there is some internal event that corresponds to our experience of a square, whether it is the activation of a cell or cell assembly in the brain, our ability to form an association between this event and the word "square" requires only that this event have a regular relation to the external object of causality, not of structural isomorphism. To insist, additionally, that these neurons must be spatially arranged in precisely the form of a square does not in the least help to explain how they come to trigger the naming response "square," at least according to Shepard and Chipman.

As can be discerned from their brief introductory paragraph summarized above, Shepard and Chipman neatly turned the tables on the debate by characterizing the perception of a square as the issue of learning the naming response "square," which is an issue of recognition rather than of perception. To be sure, recognition is an important aspect of perception, and the problem of learning a naming response is a formidable one that deserves further investigation. But the recognition response is by no means the same thing as the perceptual experience of the square as a continuous filled-in, square-shaped region of sense-data experienced in the visual field. How can so intelligent and educated researchers come to make such a profound error in identification of the issue at hand? The answer is clear from their assertion that a verbal community has access only to the public object and not to any private image. This naïve realist assumption is passed off casually as a statement of fact, but in fact it reveals an implicit commitment to the notion that the three-dimensional volumetric objects that we observe to occupy the space of our perceptual field are the actual objects themselves, and that therefore they need not be replicated or re-represented again in the brain. The fact that this assumption has gone unchallenged, and even largely unnoticed by the community at large, demonstrates how deeply the assumptions of naïve realism have become entrenched in contemporary thought.

Shepard (1981) made another attempt to dispel the issue of isomorphism by arguing for psychophysical *complementarity* rather than isomorphism. Appropriately enough, Shepard cited that grand master of naïve realism, B. F. Skinner, who argued that even if we were to discover a part of the brain in which the physical pattern of neural activity had the very same shape as the corresponding external object – say, a square – we would not in this way have made any progress toward explaining how the subject is able to recognize that object as a square, or to learn to associate to it a unique verbal response "square." So again the issue of perception is confounded with the issue of recognition response. Skinner's statement is true enough, as far as it goes. But what Shepard and Skinner failed to acknowledge is that it would be very much harder to learn to recognize a square if you could not "see" it, that is, if you did not have direct access to an internal representation of the square as a square-shaped sense-datum to associate with the appropriate recognition response. To claim that we can experience the square without such an internal replica is just plain magic. Furthermore, until we do discover a part of the brain in which the physical pattern of neural activity (or some other physically measurable quantity) has the very same shape as the corresponding external object, the phenomenal aspect of that volumetric spatial structure remains as a nomological dangler, something that is experienced as a spatial picture, something that is clearly distinct from the actual square in the real world (especially when that square is illusory), but something that does not actually exist in any space known to science. Like the Behaviorists before him, Shepard attempted to discount the entire edifice of conscious experience as if it simply did not exist as a scientific entity.

There is a further difficulty with the notion of psychophysical complementarity. Shepard (1981) argued that the relation of the mental representation to the external object it represents might be one of complementarity, rather than one of similarity or resemblance. Just as a lock has a hidden structure that is to some extent complementary to the visible contour of the key that fits it, the internal structure uniquely activated by a given object must have a structure that somehow meshes with the pattern manifested by its object; in other words, the "shape" of the representation is complementary to, rather than isomorphic with, the object that it represents. But, again, this notion of perceptual representation is only coherent from a naïve realist perspective. If we interpret this argument from an indirect perceptual view, it would have to be that the square shape we experience in immediate consciousness is complementary to the external square, which is beyond our direct experience. In other words, the real "square" in the external world is not actually square as we observe it to be, but rather it would have to be somehow complementary to the square shape we observe in conscious experience, an idea that is obviously absurd.

In yet another, somewhat different, defense of naïve realism, Shepard (1981, p. 292) argued that the relation between the external object and its internal representation might be a kind of *paramorphism* rather than isomorphism, as seen for example in the Fourier transform of an image, which encodes all of the information in a spatial image but in a very abstract nonspatial form. Again, this argument is founded on the naïve assumption that the world we see around us is the world itself, and that therefore the paramorphic representation of that world is not identified as the image of the world we see around us but as our verbal or conceptual recognition of that world. If the perceptual brain did indeed employ a Fourier representation instead of a spatial one, then the world we see around us would necessarily appear in the form of a Fourier transform rather than as a spatial structure, which, again, is obviously absurd. The fact that the world around us appears as a volumetric spatial structure is direct and concrete evidence for a spatial representation in the brain. What is most interesting about this issue is that Shepard clearly did not fully comprehend the position that he challenged, and therefore his criticisms of isomorphism inevitably missed the mark.

Steven Palmer (1999) on the other hand struck at the very heart of the issue of isomorphism. Palmer drew a distinction between two different aspects of conscious experience, the intrinsic qualities of experiences themselves versus the *relational structure* that holds among those experiences. The intrinsic qualities, such as the color qualia in the experience of color, are in principle impossible to communicate from one mind to another, and therefore they are inaccessible to science (except through phenomenology), a restriction that Palmer calls the *subjectivity barrier.* All that can be communicated about conscious experience is the relational structure that holds among those experiences. In the case of color experience, for example, subjects say that orange is more similar to red than it is to green or blue, and that aqua is experienced as intermediate between green and blue, and so forth. It was exactly these relational facts of color experience that were used to define the color solid in the CIE chromaticity diagram. A relational structure like the color solid encodes a great deal of information implicitly about the relations between its variables in a manner that is practically impossible to express as explicit relations because the number of binary, trinary, and other relations between colors implicitly expressed in the color solid is so astronomical as to defy any kind of exhaustive listing or discrete associative links. And yet all of those relations are evidently available to the psychophysical subject when making phenomenal color judgments. This strongly suggests that the variables of phenomenal color experience are encoded in the brain as a relational structure whose information content is identical to that of the color solid, rather than as a list of the astronomical number of relations between individual colors that are expressed implicitly within the color solid.

The subjectivity barrier is often cited as an insurmountable obstacle to meaningful phenomenological examination of brain states. But Palmer observed that the isomorphism constraint goes both ways. Not only is it impossible to express intrinsic color experience in objective external terms, but even if there were some way to quantify the intrinsic qualities of experience, it would then be impossible to infer the structure of the brain from that intrinsic information. The relational structure, on the other hand, does offer direct evidence for the dimensions of color experience as expressed in the physical brain because relational information is the only information that can cross the subjectivity barrier in either direction. Palmer's analysis of isomorphism has profound implications not only for color perception but also for the perception of space – although curiously Palmer avoided discussing the issue of spatial perception, presumably because including such a controversial thesis might imperil the chance of having his paper published. But a spatial percept, like that of a square, is clearly a relational structure in the sense that every point of the percept is presented simultaneously in proper spatial relation to every other point in the square. In other words, our experience of a square is of a spatial structure, and therefore the information encoded in spatial perception is an explicit spatial one, whether expressed in topographical or only topological isomorphic form.

6. The dimensions of conscious experience

The phenomenal world is composed of solid volumes, bounded by colored surfaces, embedded in a spatial void. Every point on every visible surface is perceived at an explicit spatial location in three dimensions (Clark 1993; Lehar 2003b), and all of the visible points on a perceived object, such as a cube or a sphere or this page, are perceived simultaneously in the form of continuous surfaces in depth. The perception of multiple transparent surfaces, as well as the experience of empty space between the observer and a visible surface, reveals that multiple depth values can be perceived at any spatial location. I propose to model the information in perception as a computational transformation from a two-dimensional colored image (or two images in the binocular case) to a three-dimensional volumetric data structure in which every point can encode either the experience of transparency or the experience of a perceived color at that location. The appearance of a color value at some point in this representational manifold corresponds *by definition* to the subjective experience of that color at the corresponding point in phenomenal space. If we can describe the generation of this volumetric data structure from the two-dimensional retinal image as a computational transformation, we will have quantified the information processing that is apparent in perception as a necessary prerequisite to the search for a neurophysiological mechanism that can perform that same transformation.

6.1. The Cartesian theatre and the homunculus problem

This "picture-in-the-head" or "Cartesian theatre" concept of visual representation has been criticized on the grounds that there would have to be a miniature observer to view this miniature internal scene, resulting in an infinite regress of observers within observers (Dennett 1991; 1992; O'Regan 1992; Pessoa et al. 1998). In fact, there is no need for an internal observer of the scene because the internal representation is simply a data structure like any other data in a computer, except that these data are expressed in spatial form (Earle 1998; Lehar 2003b; Singh & Hoffman 1998). If a picture in the head required a homunculus to view it, then the same argument would hold for any other form of information in the brain, which would also require a homunculus to read or interpret that information. But, in fact, any information encoded in the brain needs only to be available to other internal processes rather than to a miniature copy of the whole brain. The fact that the brain does go to the trouble of constructing a full spatial analog of the external environment merely suggests that it has ways to make use of these spatial data. For example, field theories of navigation have been proposed (Gibson & Crooks 1938; Koffka 1935, pp. 42–46) in which perceived objects in the perceived environment exert spatial fieldlike forces of attraction and repulsion, drawing the body toward attractive percepts and repelling it from aversive percepts, as a spatial computation taking place in a spatial medium.

If the idea of an explicit spatial representation in the brain seems to "fly in the face of what we know about the neural substrates of space perception" (Pessoa et al. 1998, Authors' Response sect. R3.2, p. 789), it is our theories of spatial representation that are in urgent need of revision, for to deny the spatial nature of the perceptual representation in the brain is to deny the spatial nature so clearly evident in the world we perceive around us. To paraphrase Descartes, it is not only the existence of myself that is verified by the fact that I think, but when I experience the vivid spatial presence of objects in the phenomenal world, those objects are certain to exist, at least in the form of a subjective experience, with the properties I experience them to have: location, spatial extension, color, and shape. I think them, therefore they exist (Price 1932, p. 3). All that remains uncertain is whether those percepts exist also as objective external objects as well as internal perceptual ones, and whether their perceived properties correspond to objective properties. But their existence and fully spatial nature in my internal perceptual world are beyond question if I experience them so, even if only as a hallucination.

6.2. Bounded nature of the perceptual world

The idea of perception as a literal volumetric replica of the world inside your head immediately raises the question of boundedness: How can an explicit spatial representation encode the infinity of external space in a finite volumetric system? The solution to this problem can be found by inspection, for phenomenological examination reveals that perceived space is not infinite but is bounded (Lehar 2003b). This can be seen most clearly in the night sky, where the distant stars produce a domelike percept that presents the stars at equal distance from the observer, and that distance is perceived to be less than infinite. The lower half of perceptual space is usually filled with a percept of the ground underfoot, but it too becomes hemispherical when viewed from far enough above the surface, as from an airplane or a hot air balloon. Thus the dome of the sky above and the bowl of the earth below define a finite approximately spherical space (Heelan 1983) that encodes distances out to infinity within a representational structure that is both finite and bounded. Although the properties of perceived space are approximately Euclidean near the body, there are peculiar global distortions evident in perceived space that provide clear evidence of the phenomenal world being an internal rather than an external entity.

Figure 2. The perceptual representation of a man walking down a long straight road. The sides of the road are perceived to be parallel and equidistant throughout their length, and at the same time they are perceived to converge to a point both up ahead and behind, and that point is perceived at a distance that is less than infinite. This peculiar violation of Euclidean geometry is perhaps the best evidence for the internal nature of the perceived world, for it shows evidence, out in the world around us, of the perspective projection due to the optics of the eye.

6.3. The phenomenon of perspective

Consider the phenomenon of perspective, as seen for example when standing on a long straight road that stretches to the horizon in a straight line in opposite directions. The sides of the road appear to converge to a point both up ahead and back behind, but, while converging, they are also perceived to pass to either side of the percipient, and, at the same time, the road is perceived to be straight and parallel throughout its entire length. This property of perceived space is so familiar in everyday experience as to seem totally unremarkable. And yet this most prominent violation of Euclidean geometry offers clear evidence for the non-Euclidean nature of perceived space, for the two sides of the road must in some sense be perceived as being bowed, and yet they are also perceived as being straight. This can only mean that the space within which we perceive the road to be embedded must itself be curved. In fact, the observed warping of perceived space is exactly the property that allows the finite representational space to encode an infinite external space. This property is achieved by using a variable representational scale, that is, the ratio of the physical distance in the perceptual representation relative to the distance in external space that it represents. This scale is observed to vary as a function of distance from the center of our perceived world, such a way that objects close to the body are encoded at a larger representational scale than objects in the distance, and beyond a certain limiting distance the representational scale, at least in the depth dimension, falls to zero – that is, objects beyond a certain distance lose all perceptual depth. This is seen, for example, when the sun and moon and distant mountains appear as if cut out of paper and pasted against the dome of the sky.

The distortion of perceived space is suggested in Figure 2, which depicts the perceptual representation for a man walking down a road. The phenomenon of perspective is by definition a transformation defined from a three-dimensional world through a focal point to a two-dimensional surface. The appearance of perspective on the retinal surface therefore is no mystery and is similar in principle to the image formed by the lens in a camera. What is remarkable in perception is that perspective is not observed on a two-dimensional surface but is somehow embedded in the three-dimensional space of our perceptual world. Nowhere in the objective world of external reality is there anything that is remotely similar to the phenomenon of perspective as we experience it phenomenologically, where a perspective foreshortening is observed not on a two-dimensional image but in three dimensions on a solid volumetric object. The appearance of perspective in the three-dimensional world we perceive around us is perhaps the strongest evidence for the internal nature of the world of experience, for it shows that the world that appears to be the source of the light that enters our eye must actually be downstream of the retina, for it exhibits the traces of perspective distortion imposed by the lens of the eye, although in a completely different form.

This view of perspective offers an explanation for another otherwise paradoxical but familiar property of perceived space whereby more distant objects are perceived to be both smaller and, at the same time, undiminished in size. This corresponds to the difference in subjects' reports depending on whether they are given *objective* instruction or *projective* instruction (Coren et al. 1994, p. 500) in how to report their observations, for both types of information are available perceptually. This duality in size perception is often described as a cognitive compensation for the foreshortening of perspective, as if the perceptual representation of more distant objects is indeed smaller but is somehow labeled with the correct size as some kind of symbolic tag representing objective size attached to each object in perception. However, this kind of explanation is misleading, for the objective measure of size is not a discrete quantity attached to individual objects but is more of a continuum, or gradient of difference between objective and projective size, that varies monotonically as a function of distance from the percipient. In other words, this phenomenon is best described as a warping of the space itself within which the objects are represented, so that objects that are warped coherently along with the space in which they are embedded appear undistorted perceptually. The mathematical form of this warping will be discussed in more detail in section 8.7 below.

6.4. The embodied percipient

This model of spatial representation emphasizes another aspect of perception that is often ignored in models of vision: Our percept of the world includes a percept of our own body within that world, and our body is located at a very special location at the center of that world, and it remains at the center of perceived space even as we move about in the external world. Perception is embodied by its very nature, for the percept of our body is the only thing that gives an objective measure of scale in the world, and a view of the world around us is useless if it is not explicitly related to our body in that world. The little man at the center of the spherical world of perception therefore is not a miniature observer of the internal scene but is itself a spatial percept, constructed of the same perceptual material as the rest of the spatial scene, for that scene would be incomplete without a replica of the percipient's own body in his perceived world. Gibson (1979) was right, therefore, in his emphasis on the interaction of the active organism with its environment. Gibson's only error was the epistemological one of failing to recognize that the organism and its environment that are active in perception, are themselves internal perceptual replicas of their external counterparts. It was this epistemological confusion that led to the bizarre aspects of Gibson's otherwise valuable theoretical contributions.

6.5. The ultimate question of consciousness

Indirect realism offers direct evidence for a spatial representation in the brain, but there remains one final question regarding the ultimate nature of consciousness. Even if there is a spatial representation in the brain, why should it be conscious of itself? Why should it not behave much like a machine that performs its function using either a spatial or a symbolic principle of computation but, presumably, performs its function without any conscious experience of what it is doing? Why should human consciousness be any different?

But there is a large unstated assumption implied in the very framing of this consciousness question. The assumption is that a machine could not possibly be conscious. This assumption is generally taken for granted because the alternative, that everything in the universe must have some primitive level of consciousness, seems so absurd from the outset that, like solipsism, we tend to discount it even if we cannot disprove it on logical grounds. But can we really be sure that this alternative is so absurd? Obviously, like solipsism, the possibility of *panpsychism* or, more likely, *panexperientialism* (Chalmers 1995; Rosenberg 2003) is a question that might never be provable one way or the other. Nevertheless, it is of vital importance that we get this question right, because if we come down on the wrong side of this paradigmatic fence, that will necessarily throw all the rest of our philosophy completely out of kilter.

If we accept the materialist view that mind is a physical process taking place in the physical mechanism of the brain, and since we know that mind is conscious, then we already have direct and incontrovertible evidence that a physical process taking place in a physical mechanism can under certain conditions be conscious. Now, it is true that the brain is a very special kind of mechanism. But what makes the brain so special is not its substance, for it is made of the ordinary substance of matter and energy. What sets the brain apart from normal matter is its complex organization. The most likely explanation, therefore, is that what makes our consciousness special is not its substance but its complex organization. The fundamental "stuff" of which our consciousness is composed – the basic qualia of color and spatial extension – are apparently common with the qualia of children, as far back as I can remember; although I also remember a less complex organization of my experiences as a child. It is also likely, on logical grounds, that animals have some kind of conscious qualia because the information encoded in their perceptual state cannot be experienced without some kind of quale, or carrier to express that information in the form of experience. If the experience of mind is identified as the functioning of the physical brain, then the functioning animal brain must also involve an experience of mind. Whether the subjective qualia of different species, or even of different individuals of our own species, are necessarily the same as ours experientially is a question that is difficult, even impossible in principle, to answer definitively. But the simplest, most parsimonious explanation is that our own conscious qualia evolved from those of our animal ancestors, and differ from those earlier forms more in their level of complex organization than in their fundamental nature.

The natural reluctance we all feel to extending consciousness to our animal ancestors, and even more so to plants or to inanimate matter, is a stubborn legacy of our anthropocentric past. But the history of scientific discovery has been characterized by a regular progression of *anthrodecentralization,* demoting humans from the central position in the universe under the personal supervision of God, to lost creatures on the surface of a tiny blip of matter orbiting a very unremarkable star among countless billions of stars in an unremarkable galaxy amongst countless billions of other galaxies as far as the telescopic eye can see. Modern biology has now discovered that there is no vital force in living things, but only a complex organization of the ordinary matter of the universe, following the ordinary laws of that universe. There is no reason on earth why consciousness should not also be considered to be a manifestation of the ordinary matter of the universe following the ordinary laws of that universe, although expressed in a complex organization in the case of the human brain. A claim to the contrary would necessarily fall under the category of an extraordinary claim, which, as Carl Sagan pointed out, would require extraordinary evidence for it to be accepted by reasonable men.

When we examine the chain of biocomplexity from the simplest pure chemical to the most complex human brain, there is a continuous progression from single atoms, to simple compound molecules, to complex organic molecules, to proteins and DNA, to viruses, to simple single-celled organisms, and all the way up the evolutionary chain to the brain of man. If we are to claim that consciousness is uniquely human, or unique to animals above a certain complexity, then there would necessarily be some kind of abrupt transition along that progression where that consciousness comes suddenly into existence, and that abrupt transition would occur both for the individual during gestation and for the species during evolution. The claim that consciousness is unique to humans, or to animals, or to living creatures, is bedeviled by the fact that there are always transitionary forms to be found that are intermediate between humans and animals, between animals and plants, and between living and nonliving creatures such as hypercomplex molecules like viruses; and there is also a continuous progression during gestation from fertilized egg to full human body. If we posit that consciousness appears abruptly at any one of these transitions where the only observed difference is a slight increase in complexity of organization, then we again lapse into nomological danglers and vital force because the postulated conscious quality that supposedly appears abruptly at that point is undetectable to science, and therefore it is a quality of the supervenient spirit world rather than anything knowable by, or demonstrable to, science. Surely the time has come to finally accept the full implications of Darwin's theory of evolution and acknowledge the fact that our nature and our consciousness are not of a separate spiritual realm but are composed of the very same material substance and energy of which the rest of the universe is composed.

The inescapable conclusion is that all matter and energy have some kind of primal protoconsciousness, what Chalmers (1995) calls "panexperientialism" to distinguish it from panpsychism, the view that everything is conscious in any human kind of sense. The more plausible panexperientialism posits merely that there exists a very simple protoconsciousness in inanimate matter that is a fundamental property of that matter. For inanimate matter, this protoconsciousness is something so simple and primitive that we would hardly recognize it as consciousness at all. And yet when this protoconsciousness is organized in the right manner in a human brain, it gives rise to the wonderful splendor of human consciousness. We are not external observers of the physical universe, rather we ourselves are part of that universe and our experience is a tiny fragment of the experience of the larger universe around us, although expressed in a very much more complex form in the human brain. This way of describing consciousness is the only true monism that really equates mind with the functioning of physical matter, without recourse to nomological danglers and spiritual mumbo jumbo.

This identity relation between mind and matter casts a new light on Searle's (1997) assertion that "a computer is not even a computer to a computer." What would the consciousness of a computer be like, if a computer did have consciousness? Consider the hypothesis that consciousness is a manifestation of forces and energy, or energetic wrinkles in space-time, or what Rosenberg (2003) called manifestations of causality in the physical world. The consciousness of a computer would thereby correspond to the patterns of energy in its chips and wires. In a digital computer that consciousness would be a very binary affair, and it is also in the very nature of digital computation that complex calculations are divided into a number of very simple steps, each of which can be computed independent of the problem as a whole. The consciousness of a computer would thereby be a very fragmented kind of thing, with each flip-flop or logic gate experiencing only the energy state in its local inputs and outputs because those are the only forces that influence the local logic gate. There is a very different kind of energy structure in an analog spatial system like a soap bubble, whose entire surface is under tension against the outward pressure of the captured air. A push on any point of the bubble has an immediate influence on the bubble as a whole, on the entire gestalt, whose causal structure works in an emergent manner to try to restore the spherical shape. If a soap bubble has any form of primal consciousness, that protoconsciousness would be of an elastic spherical form under stress, as a unitary gestalt.

It is curious that in his Chinese room analogy, Searle (1980) assumed a fragmented, rule-based mechanism as his model of conscious experience, because in his analogy the Chinese translation is performed step-by-step, very much like the computation in a digital computer. No wonder there is no emergent global consciousness from such a fragmented computational analogy. But does a globally integrated analog structure like a soap bubble have a corresponding global consciousness independent of the individual consciousnesses of its constituent parts? And does that larger consciousness include the consciousnesses of its individual parts? The answers to these questions can be found by inspection of our own consciousness.

By the fact that we ourselves have global consciousness, we can infer that larger global phenomena in the brain do give rise to global emergent consciousness that takes the form we observe in the perceived world around us. And that global consciousness does not appear to include a consciousness of its individual elements, for we are completely unaware of the component electrons, molecules, and neurons of our own physical brain that must be responsible for that global percept. Our personal conscious experience is therefore confined to an awareness of the spatial structures of the patterns of energy in our brain, although presumably there would also be many more independent and disconnected consciousnesses in the energy structures of our physical body of which we are not directly aware, and most likely there are also multiple independent conscious entities within our own brain, which make up the "unconscious mind," and of which our central narrative consciousness is not directly aware.

Consider the experience of swallowing food. I am conscious of the inside of my mouth as a vivid three-dimensional structure "colored" by sensations of taste and texture, warmth and cold. But this spatial consciousness terminates abruptly at the threshold of my throat, beyond which my spatial consciousness of the food is abruptly cut off. The rest of my alimentary canal performs wavelike motions of peristaltic contraction, very much like the kind of manipulation that occurs consciously in my mouth, but these motions are all beyond my own personal conscious awareness. Is my alimentary canal conscious of itself, or does it perform its function totally in the absence of conscious experience? It seems that conscious experience has a direct functional role, because my consciousness of my own mouth helps me to chew the food and direct it intelligently down my throat without choking. If the food I swallowed was a hot and spicy vindaloo curry, I know in an indirect way that my stomach is feeling the burning pain because I can feel it churning and grinding in protest, although I cannot feel its pain directly, only remotely, like a loud argument heard through the wall in an adjacent motel room. And the next morning, as the vindaloo curry passes another abrupt threshold portal, I become suddenly aware of the pain again as part of my own personal experience. The simplest explanation therefore is that my alimentary canal has a similar conscious experience; it feels the waves of peristaltic contraction, which are its own conscious wavelike thoughts, just as I feel the inside of my mouth, although, unlike my central narrative consciousness, presumably the alimentary consciousness is not burdened by memories or aspirations or any real self-consciousness except of itself as a spatial structure and of the vital imperative to propel arriving food farther down the pipeline. This hypothesis is supported by the fact that there exist rare individuals who have conscious control over their own bowel functions, such that they can consciously control their own alimentary peristaltic contractions just as we can control the contractions of our mouths. And that extra level of control is accompanied by a conscious experience of their alimentary canal in places where we have none. In fact, it is not clear if it is even intelligible to have control of a body part without having some kind of consciousness of that control.

There is not, therefore, a single "bridge locus" that is the only place in the brain where consciousness occurs, but rather there is one global representational mechanism which has verbal and cognitive access to the components of ordinary consciousness including memories and aspirations, and then there are countless additional independent conscious energy structures disconnected from our global or narrative consciousness of which we remain personally unaware. Each of those islands of consciousness has an isolated experience of its own energy structure.

If consciousness is indeed identical to energy structure, then the spherical bubble can be conscious of its own spherical form, although it has neither memory nor aspirations, nor any kind of understanding except an understanding of its own spherical energy structure. How then does human consciousness come to be aware not only of its own structure but also that of the external environment around it beyond the bounds of the physical brain? It does so by constructing a much more complex and elaborate bubble structure in the human brain, composed of patterns of electrochemical energy that take the form of a replica of the external world, complete with a replica of our own body at the center of that representational space. So, in answer to Searle's contention, the computer too could acquire a consciousness of itself, if it were loaded with a representation of itself. The pattern of that representation in the computer would thereby appear as a computer to the computer. Of course, like us, the computer would not notice that what it was seeing was not really an image of its *real* self, as viewed from the outside, but merely a miniature representation of itself that would be entirely contained within itself, because its computational consciousness could not extend beyond the confines of its computational brain. The computer would not know that everything of which it was aware was actually surrounded by the larger physical computer, which in turn was composed of entirely different and independent sets of conscious energy structures in the physical structures of its frame and screws and power supply.

If this notion of panexperientialism, or protoconsciousness of inanimate matter, sounds bizarre and far-fetched, we should bear in mind that whatever the ultimate solution to the mind-brain quandary, it is sure to do considerable violence to our normal, everyday, commonsense notions of reality. When it comes to these fundamental issues of existence, our intuitive instincts are almost certain to fail us, and therefore every alternative should be given serious consideration, however implausible it might at first seem intuitively. For, as intuitively incredible as the notion of panexperientialism might seem, the alternatives are all fraught with even more profound philosophical paradoxes and contradictions. But whatever our theoretical inclinations on the ultimate question of consciousness, it is important to point out that this is a separate and independent issue from the question of whether the internal representation of the brain is spatial or symbolic. Whichever way the answer to the ultimate question goes, whether consciousness is uniquely human or is shared with the living and nonliving worlds, unless we wish to believe in some magical nomological dangler that extends mind halfway into the spirit world, we must face the observational fact that there is a spatial representation in the brain.

7. The Gestalt properties of perception

One of the most formidable obstacles facing computational models of the perceptual process is that perception exhibits certain global Gestalt properties such as emergence, reification, multistability, and invariance that are difficult to account for either neurophysiologically or even in computational terms such as computer algorithms. The ubiquity of these properties in all aspects of perception, as well as their preattentive nature, suggests that Gestalt phenomena are fundamental to the nature of the perceptual mechanism. I propose that no useful progress can possibly be made in our understanding of neural processing until the computational principles behind Gestalt theory have been identified.

7.1. Emergence

Figure 3 shows a picture that is familiar in vision circles, for it reveals the principle of emergence in a most compelling form. The picture appears initially as a random pattern of irregular shapes, but a remarkable transformation is observed in this percept as soon as one recognizes the subject of the picture as a dalmation in the patchy sunlight under overhanging trees. What is remarkable about this percept is that the dog is perceived so vividly despite the fact that much of its perimeter is missing. Furthermore, visual edges, which form a part of the perimeter of the dog, are locally indistinguishable from other less significant edges. Therefore, any local portion of this image does not contain the information necessary to distinguish significant from insignificant edges.

Although Gestalt theory did not offer any specific computational mechanism to explain emergence in visual perception, Koffka (1935) suggested a physical analogy of the soap bubble to demonstrate the operational principle behind emergence. The spherical shape of a soap bubble is not encoded in the form of a spherical template or abstract mathematical code, but rather, that form emerges from the parallel action of innumerable local forces of surface ten-

Figure 3. The dog picture is familiar in vision circles for it demonstrates the principle of emergence in perception. The local regions of this image do not contain sufficient information to distinguish significant form contours from insignificant noisy edges. As soon as the picture is recognized as that of a dog in the dappled sunshine under trees, the contours of the dog pop out perceptually, filling in visual edges in regions where no edges are present in the input.

sion acting in unison. The characteristic feature of emergence is that the final global form is not computed in a single pass but continuously, like a relaxation to equilibrium in a dynamic system model. In other words, the forces acting on the system induce a change in the system configuration, and that change in turn modifies the forces acting on the system. The system configuration and the forces that drive it are changing continuously in time until equilibrium is attained, at which point the system remains in a state of dynamic equilibrium. Even at this point, its static state belies a dynamic balance of forces ready to spring back into motion as soon as the balance is upset.

Emergence is actually the issue that inspired Davidson's (1970) theory of anomalous monism. Davidson argued (p. 247) that mental events resist capture in the nomological net of physical theory, for mentalistic propositions do not display the law-like character of physical ones. Davidson asserted (p. 248) that "there are no strict deterministic laws on the basis of which mental events can be predicted and explained," and this is the principle of the anomalism of the mental. But Wolfgang Köhler (1924) showed that in fact there is no magic in emergence; rather, emergence is a common property of certain kinds of physical systems, such as the soap bubble taking on its spherical shape, or water seeking its own level in a vessel, or global weather patterns defying lawful prediction based on their present state. To insist that mind supervenes on the brain in some mysterious way is like saying that the soap bubble supervenes on soapy water, or that the water level supervenes on the body of water in a vessel, or that global weather patterns supervene on the earth's physical atmosphere. But this is no different than saying that these are emergent processes that are already the simplest model of themselves. Emergence in perception does not imply that the mind supervenes on the brain, but rather it indicates that the neurophysiological processes involved in perception exhibit the kind of holistic emergence seen in the soap bubble, where a multitude of tiny forces act together simultaneously to produce a final perceptual state by way of a process that cannot be reduced to simple laws.

7.2. Reification

The Kanizsa figure (Kanizsa 1979), shown in Figure 4A, is one of the most familiar illusions introduced by Gestalt theory. In this figure the triangular configuration is not only recognized as being present in the image, but that triangle is filled-in perceptually, producing visual edges in places where no edges are present in the input; and those edges in turn are observed to bound a uniform triangular region that is brighter than the white background of the figure. Idesawa (1991) and Tse (1999a; 1999b) extended this concept with a set of even more sophisticated illusions, including those shown in Figures 4B–D, in which the illusory percept takes the form of a three-dimensional volume. These figures demonstrate that the visual system performs a perceptual *reification,* a filling-in of a more complete and explicit perceptual entity based on a less complete visual input. Reification is a general principle of perceptual processing, of which boundary completion and surface filling-in are more specific computational components. The identification of this generative aspect of perception is one of the most significant contributions of Gestalt theory.

7.3. Multistability

A familiar example of multistability in perception is seen in the Necker cube, shown in Figure 5A. Prolonged viewing of this stimulus results in spontaneous reversals, in which the entire percept is observed to invert in depth. Figure 5B shows how large regions of the percept invert coherently in bistable fashion. Even more compelling examples of multistability are seen in surrealistic paintings by Salvator Dali

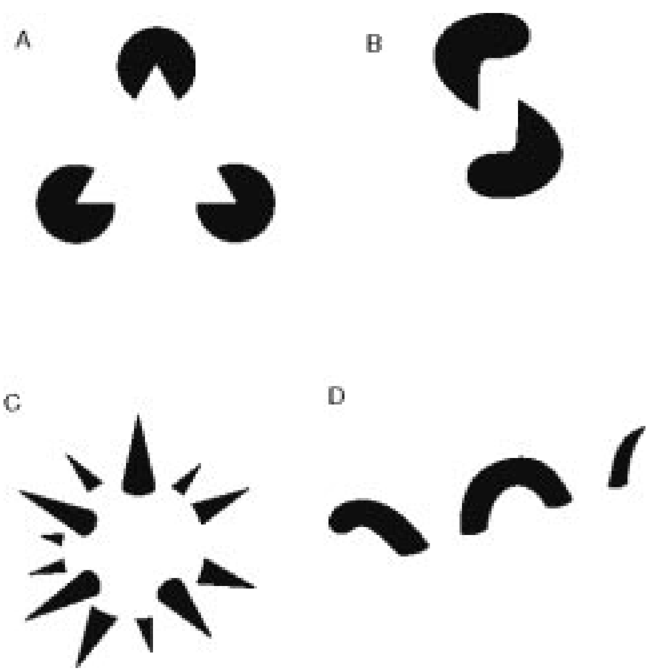

Figure 4. **A.** The Kanizsa triangle. **B.** Tse's volumetric worm. **C.** Idesawa's spiky sphere. **D.** Tse's "sea monster."

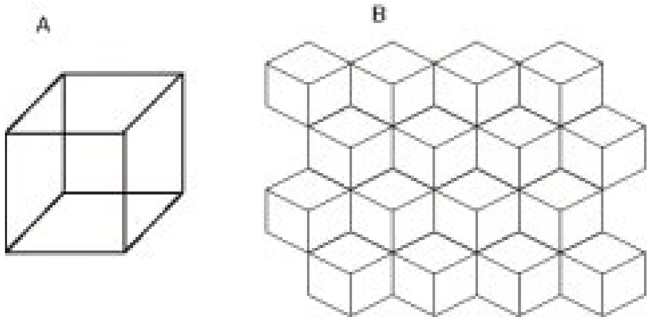

Figure 5. **A.** The Necker cube demonstrates multistability in perception. **B.** This figure shows how large regions of the percept flip coherently between perceptual states.

and etchings by Escher, in which large and complex regions of an image are seen to invert perceptually, losing all resemblance to their former appearance (Attneave 1971). The significance for theories of visual processing is that perception cannot be considered as simply a feed-forward processing performed on the visual input to produce a perceptual output, as it is most often characterized in computational models of vision, but rather perception must involve some kind of dynamic process whose stable states represent the final percept.

7.4. Invariance

A central focus of Gestalt theory is the issue of invariance – how an object, like a square or a triangle, can be recognized regardless of its rotation, translation, or scale, or whatever its contrast polarity against the background, or whether it is depicted in solid or outline form, or whether it is defined in terms of texture, motion, or binocular disparity. This invariance is not restricted to the two-dimensional plane but is also observed through rotation in depth, and even in invariance to perspective transformation. For example, the rectangular shape of a tabletop is recognized even when its retinal projection is in the form of a trapezoid due to perspective, and yet when we view the tabletop from any particular perspective we can still identify the exact contours in the visual field that correspond to the boundaries of the perceived table, to the highest resolution of the visual system. The ease with which these invariances are handled in biological vision suggests that invariance is fundamental to the visual representation.

Our failure to find a neurophysiological explanation for Gestalt phenomena does not suggest that no such explanation exists, only that we must be looking for it in the wrong places. The enigmatic nature of Gestalt phenomena highlights the importance of the search for a computational mechanism that exhibits these same properties. In the next section, I present a model that demonstrates how these Gestalt principles can be expressed in a computational model that is isomorphic with the subjective experience of vision.

8. The computational mechanism of perception

The basic function of visual perception can be described as the transformation from a two-dimensional retinal image, or a pair of images in the binocular case, to a solid three-dimensional percept. Figure 6A depicts a two-dimensional stimulus that produces a three-dimensional percept of a solid cube complete in three dimensions. For simplicity, a simple line drawing is depicted in the figure, but the argument applies more appropriately to a view of a real cube observed in the world. Every point on every visible surface of the percept is experienced at a specific location in depth, and each of those surfaces is experienced as a planar continuum, with a specific three-dimensional slope in depth. The information in this perceptual experience can therefore be expressed as a three-dimensional model, as suggested in Figure 6B, constructed on the basis of the input image in Figure 6A.

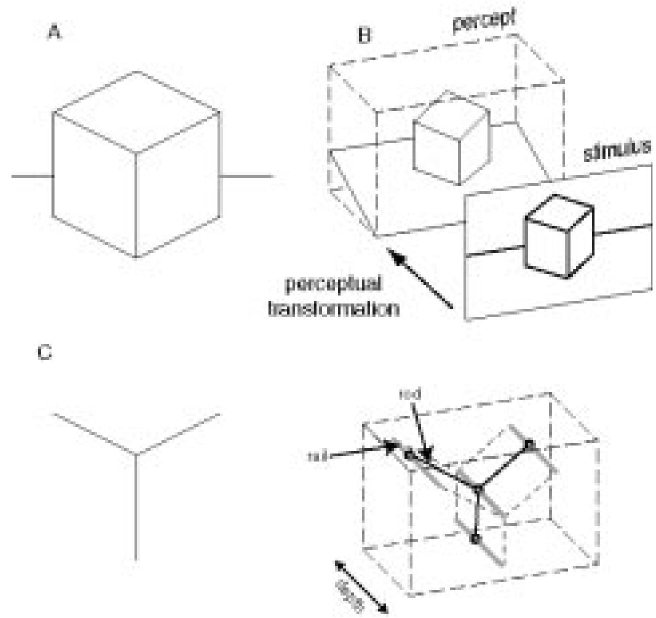

Figure 6. **A.** A line drawing stimulates **B,** a volumetric spatial percept with an explicit depth value at every point on every visible surface, and an amodal percept of hidden rear surfaces. **C.** The central Y vertex from panel A, which tends to be perceived as a corner in depth. **D.** A dynamic rod-and-rail model of the emergence of the depth percept in panel C by relaxation of local constraints.

The transformation from a two-dimensional image space to a three-dimensional perceptual space is known as the *inverse optics* problem because the intent is to reverse the optical projection in the eye, in which three-dimensional information from the world is collapsed into a two-dimensional image. However, the inverse optics problem is underconstrained, for there are an infinite number of possible three-dimensional configurations that can give rise to the same two-dimensional projection. How does the visual system select from this infinite range of possible percepts to produce the single perceptual interpretation observed phenomenally? The answer to this question is of central significance to understanding the principles behind perception, for it reveals a computational strategy quite unlike anything devised by man, and certainly unlike the algorithmic decision sequences embodied in the paradigm of digital computation. The transformation observed in visual perception gives us the clearest insight into the nature of this unique computational strategy. I propose that the principles of emergence, reification, and multistability are intimately involved in this reconstruction, and that in fact these Gestalt properties are exactly the properties needed for the visual system to address the fundamental ambiguities inherent in reflected light imagery.

The principle behind the perceptual transformation can be expressed in general terms as follows. For any given visual input there is an infinite range of possible configurations of objects in the external world which could have given rise to that same stimulus. The configuration of the stimulus constrains the range of possible perceptual interpretations to those that line up with the stimulus in the two dimensions of the retinal image. Although each individual interpretation within that range is equally likely with respect to the stimulus, some of those perceptual alternatives are intrinsically more likely than others, in the sense that they are more typical of objects commonly found in the world. I propose that the perceptual representation has the property that the more likely structural configurations are also more stable in the perceptual representation, and therefore the procedure used by the visual system is to essentially construct or reify all possible interpretations of a visual stimulus in parallel, as constrained by the configuration of the input, and then to select from that range of possible percepts the most stable perceptual configuration by a process of emergence. In other words, perception can be viewed as the computation of the intersection of two sets of constraints, which might be called *extrinsic* and *intrinsic* constraints. The extrinsic constraints are those determined by the visual stimulus, the intrinsic constraints are determined by the structural stability of the percept.

Arnheim (1969) presented an insightful analysis of this concept, which can be reformulated as follows. Consider (for simplicity) just the central Y vertex of Figure 6A depicted in Figure 6C. Arnheim proposed that the extrinsic constraints of inverse optics can be expressed for this stimulus using a rod-and-rail analogy as shown in Figure 6D. The three rods, representing the three edges in the visual input, are constrained in two dimensions to the configuration seen in the input, but are free to slide in depth along the four rails. The rods must be elastic between their endpoints, so that they can expand and contract in length. By sliding along the rails, the rods can take on any of the infinite three-dimensional configurations corresponding to the two-dimensional input of Figure 6C. For example, the final percept could theoretically range from a percept of a convex vertex protruding from the depth of the page to a concave vertex intruding into the depth of the page, with a continuum of intermediate perceptual states between these limits.

There are other possibilities beyond these, such as percepts where each of the three rods is at a different depth and therefore they do not meet in the middle of the stimulus. However, these alternative perceptual states are not all equally likely to be experienced. Hochberg and Brooks (1960) showed that the final percept is the one that exhibits the greatest simplicity, or *prägnanz.* In the case of the vertex of Figure 6C the percept tends to appear as three rods whose ends coincide in depth at the center, and meet at a mutual right angle, defining either a concave or convex corner. This reduces the infinite range of possible configurations to two discrete perceptual states. This constraint can be expressed emergently in the rod-and-rail model by joining the three rods flexibly at the central vertex, and installing spring forces that tend to hold the three rods at mutual right angles at the vertex. With this mechanism in place to define the intrinsic or structural constraints, the rod-and-rail model becomes a dynamic system that slides in depth along the rails, and this system is bistable between a concave and a convex right-angled percept, as observed phenomenally in Figure 6C. Although this model reveals the dynamic interaction between intrinsic and extrinsic constraints, this particular analogy is hardwired to modeling the percept of the triangular vertex of Figure 6C. I will now develop a more general model that operates on this same dynamic principle, but is designed to handle arbitrary input patterns.

8.1. A Gestalt Bubble model

For the perceptual representation, I propose (Lehar 2003b) a volumetric block or matrix of dynamic computational elements, as suggested in Figure 7A, each of which can exist in one of two states, transparent or opaque, with opaque-state units being active at all points in the volume of perceptual space where a colored surface is experienced. In other words, upon viewing a stimulus like that in Figure 6A, the perceptual representation of this stimulus is modeled as a three-dimensional pattern of opaque-state units embedded in the volume of the perceptual matrix in exactly the configuration observed in the subjective perceptual experience when viewing Figure 6A – with opaque-state elements at all points in the volumetric space that are within a perceived surface in three dimensions, as suggested in Figure 6B. All other elements in the block are in the transparent state to represent the experience of the spatial void within which perceived objects are perceived to be embedded. More generally, opaque-state elements should also encode the subjective dimensions of color (hue, intensity, and saturation), and intermediate states between transparent and opaque would be required to account for the perception of semitransparent surfaces, although for now the discussion will be limited to two states and the monochromatic case. The transformation of perception can now be defined as the turning on of the appropriate pattern of elements in this volumetric representation in response to the visual input, in order to replicate the three-dimensional configuration of surfaces experienced in the subjective percept.

8.2. Surface percept interpolation

The perceived surfaces due to a stimulus like Figure 6A appear to span the structure of the percept defined by the edges in the stimulus, somewhat like a milky bubble surface clinging to a cubical wire frame. Although the featureless portions of the stimulus between the visual edges offer no explicit visual information, a continuous surface is perceived within those regions, as well as across the white background behind the block figure, with a specific depth and surface orientation value encoded explicitly at each point in the percept. This three-dimensional surface interpolation function can be expressed in the perceptual model by assigning to every element in the opaque state a surface orientation value in three dimensions, and by defining a dynamic interaction between opaque-state units to fill in the region between them with a continuous surface percept. In order to express this process as an emergent one, the dynamics of this surface interpolation function must be defined in terms of local fieldlike forces analogous to the lo-

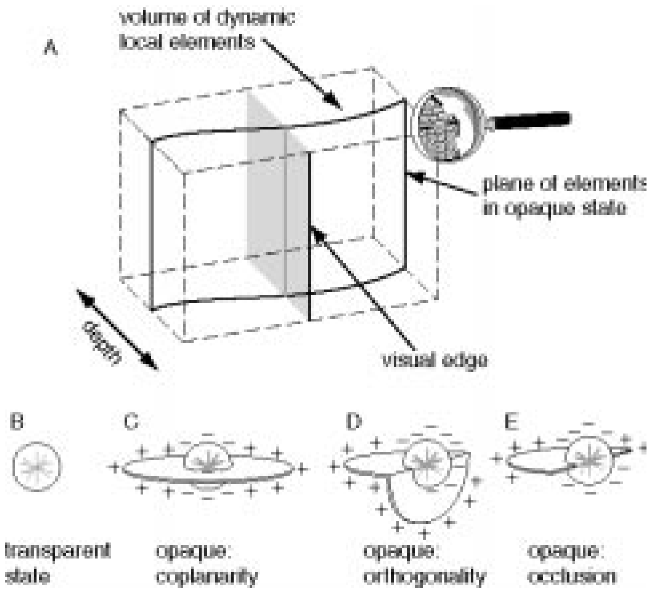

Figure 7. **A.** The Gestalt Bubble model consisting of a block of dynamic local elements which can be in one of several states. **B.** The transparent state, no neighborhood interactions. **C.** The opaque coplanarity state which tends to complete smooth surfaces. **D.** The opaque orthogonality state which tends to complete perceptual corners. **E.** The opaque occlusion state which tends to complete surface edges.

cal forces of surface tension active at any point in a soap bubble. Figure 7B represents a point in the perceptual matrix in the transparent state, representing a percept of empty space at that location. Elements in this state do not project field-like forces to adjacent elements. Figure 7C depicts an opaque-state unit representing a local portion of a perceived surface at a specific three-dimensional location and with a specific surface orientation. The planar field of this element, depicted somewhat like a planetary ring in Figure 7C, represents both the perceived surface represented by this element, as well as a fieldlike influence propagated by that element to adjacent units. This planar field fades smoothly with distance from the center with a Gaussian function. The effect of this field is to recruit adjacent elements within that field of influence to take on a similar state – that is, to induce transparent state units to switch to the opaque state, and opaque-state units to rotate toward a similar surface orientation value. The final state and orientation taken on by any element is computed as a spatial average or weighted sum of the states of neighboring units as communicated through their planar fields of influence, with the greatest influence from nearby opaque elements in the matrix. The influence is reciprocal between neighboring elements, thereby defining a circular relation as suggested by the principle of emergence. To prevent runaway positive feedback and uncontrolled propagation of surface signal, an inhibitory dynamic is also incorporated. This dynamic suppresses surface formation out of the plane of the emergent surface by endowing the local field of each unit with an inhibitory field that suppresses the opaque state in neighboring elements in all directions outside of the plane of its local field. The mathematical specification of the local field of influence between opaque-state units is outlined in greater detail in the Appendix. However, the intent of the model is expressed more naturally in the global properties as described here, so the details of the local field influences are presented as only one possible implementation of the concept, provided in order to ground this somewhat nebulous idea in more concrete terms.

The global properties of the system should be such that if the elements in the matrix were initially assigned randomly to either the transparent or opaque state, with random surface orientations for opaque-state units, the mutual fieldlike influences would tend to amplify any group of opaque-state elements whose planar fields happened to be aligned in an approximate plane; and as that plane of active units feeds back on its own activation, the orientations of its elements would conform ever closer to that of the plane, and elements outside of the plane would be suppressed to the transparent state. This would result in the emergence of a single plane of opaque-state units as a dynamic global pattern of activation embedded in the volume of the matrix, and that surface would be able to flex and stretch much like a bubble surface. Although, unlike a real bubble, this surface is defined not as a physical membrane but as a dynamic sheet of active elements embedded in the matrix. This volumetric surface interpolation function will now serve as the backdrop for an emergent reconstruction of the spatial percept around a three-dimensional skeleton or framework constructed on the basis of the visual edges in the scene.

8.3. Local effects of a visual edge

A visual edge can be perceived as an object in its own right, like a thin rod or wire surrounded by empty space. More often, however, an edge is seen as a discontinuity in a surface, either as a corner or a fold, or perhaps as an occlusion edge like the outer perimeter of a flat figure viewed against a more distant background. The interaction between a visual edge and a perceived surface can therefore be modeled as follows: The two-dimensional edge from the retinal stimulus projects a different kind of field of influence into the depth dimension of the volumetric matrix, as suggested by the gray shading in Figure 7A, to represent the three-dimensional locus of all possible edges that project to the two-dimensional edge in the image. In other words, this field expresses the inverse optics probability field or extrinsic constraint due to a single visual edge. Wherever this field intersects opaque-state elements in the volume of the matrix, it changes the shape of their local fields of influence from a coplanar interaction to an orthogonal, or corner, interaction as suggested by the local force field in Figure 7D. The corner of this field should align parallel to the visual edge but otherwise remain unconstrained in orientation except by interactions with adjacent opaque units. Visual edges can also denote occlusion, and so opaque-state elements can also exist in an occlusion state, with a coplanarity interaction in one direction only, as suggested by the occlusion field in Figure 7E. Therefore, in the presence of a single visual edge, a local element in the opaque state should have an equal probability of changing into the orthogonality or occlusion state, with the orthogonal or occlusion edge aligned parallel to the inducing visual edge. Elements in the orthogonal state tend to promote orthogonality in adjacent elements along the perceived corner, whereas elements in the occlusion state promote occlusion along that edge. In other words, an edge will tend to be perceived as a corner or occlusion percept along its entire length, although the whole edge may change state back and

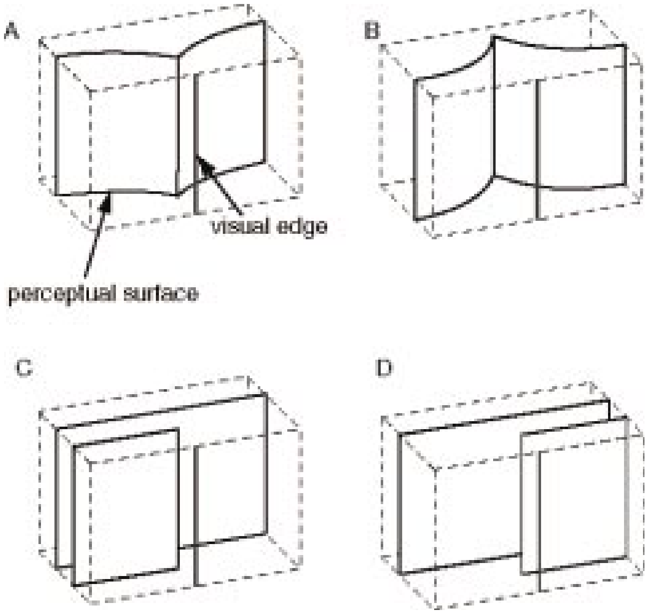

Figure 8. Several possible stable states of the Gestalt Bubble model in response to a single visual edge.

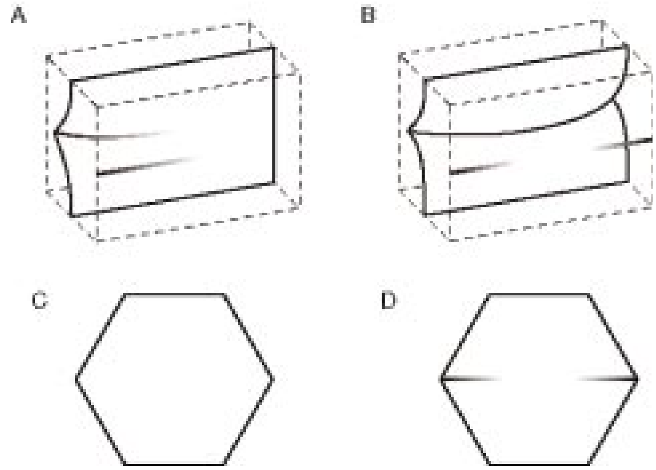

Figure 9. **A.** Boundary completion in the Gestalt bubble model: A single line ending creates a crease in the perceptual surface. **B.** Two line endings generate a crease joining them. **C.** A regular hexagon figure transforms into **D,** a percept of a folded surface in depth, with the addition of suggestive lines, with the assistance of a global gestalt that is consistent with that perceptual interpretation.

forth as a unit in a multistable manner. (The Appendix presents a more detailed mathematical description of how these orthogonality and occlusion fields might be defined.) The presence of the visual edge in Figure 7A tends to crease or break the perceived surface into one of the several possible configurations shown in Figures 8A–D. The final configuration selected by the system would depend not only on the local image region depicted in Figure 8, but also on forces from adjacent regions of the image, in order to fuse the orthogonal or occlusion state elements seamlessly into nearby coplanar surface percepts.

8.4. Global effects of configurations of edges

Visual illusions like the Kanizsa figure shown in Figure 4A suggest that edges in a stimulus that are in a collinear configuration tend to link up in perceptual space to define a larger global edge connecting the local edges. This kind of collinear boundary completion is expressed in this model as a physical process analogous to the propagation of a crack or fold in a physical medium. A visual edge that fades gradually produces a crease in the perceptual medium which tends to propagate outward beyond the edge as suggested in Figure 9A. If two such edges are found in a collinear configuration, the perceptual surface will tend to crease or fold between them as suggested in Figure 9B. This tendency is accentuated if additional evidence from adjacent regions supports this configuration. This can be seen in Figure 9D where fading horizontal lines are seen to link up across the figure to create a percept of a folded surface in depth, which would otherwise appear as a regular hexagon as in Figure 9C.

Gestalt theory emphasizes the significance of closure as a prominent factor in perceptual segmentation because an enclosed contour is seen to promote a figure/ground segregation (Koffka 1935, p. 178). For example, an outline square tends to be seen as a square surface in front of a background surface that is complete and continuous behind the square, as suggested in the perceptual model depicted in Figure 10A. The problem is that closure is a *gestaltqualität,* a quality defined by a global configuration that is difficult to specify in terms of any local featural requirements, especially in the case of irregular or fragmented contours as seen in Figure 10B. In this model, an enclosed contour breaks away a piece of the perceptual surface, completing the background amodally behind the occluding foreground figure. In the presence of irregular or fragmented edges the influences of the individual edge fragments act collectively to break the perceptual surface along that contour as suggested in Figure 10C, like the breaking of a physical surface that is weakened along an irregular line of cracks or holes. The final scission of figure from ground is therefore driven not so much by the exact path of the individual irregular edges as by the global configuration of the emergent gestalt.

8.5. Vertices and intersections

In the case of vertices or intersections between visual edges, the different edges interact with one another, favoring the percept of a single vertex at that point. For example, the three edges defining the three-way Y vertex shown in Figure 6C promote the percept of a single three-dimensional corner whose depth profile depends on whether the corner is perceived as convex or concave. In the case of Figure 6A, the cubical percept constrains the central Y vertex as a convex rather than a concave trihedral percept. I propose that this dynamic behavior can be implemented using the same kinds of local field forces described in the Appendix to promote mutually orthogonal completion in three dimensions, wherever visual edges meet at an angle in two dimensions. Figure 11A depicts the three-dimensional influence of the two-dimensional Y vertex when projected on the front face of the volumetric matrix. Each plane of this three-planed structure promotes the emergence of a corner or occlusion percept at some depth within that plane. But the effects due to these individual edges are not independent. Consider first the vertical edge projecting from

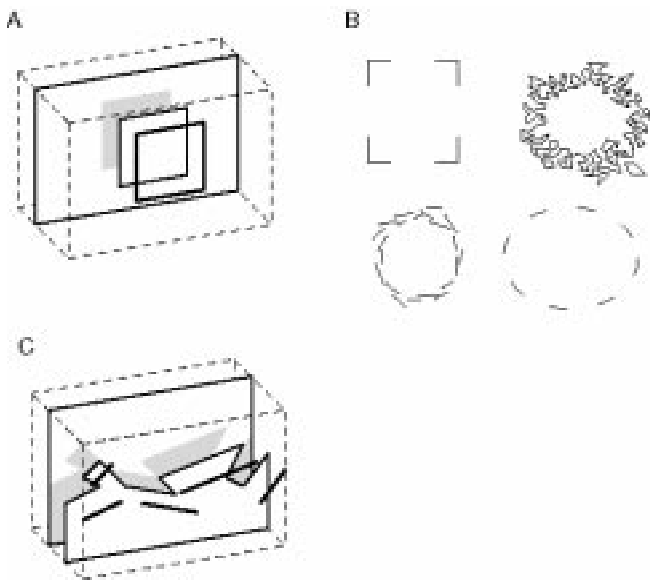

Figure 10. **A.** The perception of closure and figure/ground segregation are explained in the Gestalt Bubble model exactly as perceived, in this case as a foreground square in front of a background surface that completes behind the square. **B.** Even irregular and fragmented surfaces produce a figure/ground segregation. **C.** The perceived boundary of the fragmented figure follows the global emergent gestalt rather than the exact path of individual edges.

the bottom of the vertex. By itself, this edge might produce a folded percept as suggested in Figure 11B, which could occur through a range of depths and a variety of orientations in depth, as well as in concave or convex form. But the two angled planes of this percept each intersect the other

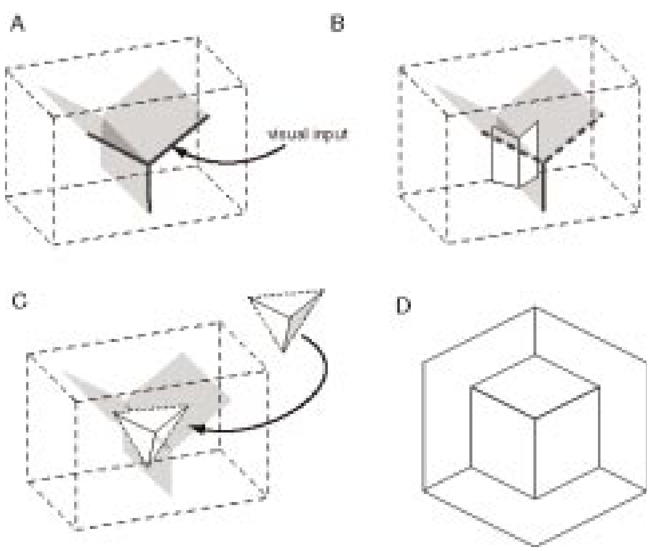

Figure 11. **A.** The three-dimensional field of influence due to a two-dimensional Y vertex projected into the depth dimension of the volumetric matrix. **B.** Each field of influence, for example the one due to the vertical edge, stimulates a folded surface percept. The folded surface intersects the other fields of influence due to the other two edges, thereby tending to produce a percept of a single corner percept. **C.** One of many possible emergent surface percepts in response to that stimulus, in the form of a convex trihedral surface percept. **D.** The percept can also be of a concave trihedral corner, as seen sometimes at the center in this bistable figure.

two fields of influence due to the other two edges of the stimulus, as suggested in Figure 11B, thus favoring the emergence of those edges' perceptual folds at that same depth, and resulting in a single trihedral percept at some depth in the volumetric matrix, as suggested in Figure 11C. Any dimension of this percept that is not explicitly specified or constrained by the visual input remains unconstrained. In other words, the trihedral percept is embedded in the volumetric matrix in such a way that its three component corner percepts are free to slide inward or outward in depth, to rotate through a small range of angles, and to flip in a bistable manner between a convex and a concave trihedral configuration. The model now expresses the multistability of the rod-and-rail analogy shown in Figure 6D but in a more generalized form that is no longer hardwired to the Y-vertex input shown in Figure 6C; it can accommodate any arbitrary configuration of lines in the input image. A local visual feature like an isolated Y vertex generally exhibits a larger number of stable states, whereas in the context of adjacent features the number of stable solutions is often diminished. This explains why the cubical percept of Figure 6A is stable, but its central Y vertex alone, as shown in Figure 6C, is bistable. The fundamental multistability of Figure 6A can be revealed by the addition of a different spatial context, as depicted in Figure 11D.

8.6. Perspective cues

Perspective cues offer another example of a computation that is inordinately complicated in most models. In a fully reified spatial model, however, perspective can be computed relatively easily with only a small change in the geometry of the model. Figure 12A shows a trapezoid stimulus, which has a tendency to be perceived in depth, with the shorter top side being perceived as of the same length as the longer base but apparently diminished by perspective. Arnheim (1969) suggested a simple distortion to the volumetric model to account for this phenomenon, which can be reformulated as follows. The height and width of the volumetric matrix are diminished as a function of depth, as suggested in Figure 12B, transforming the block shape into a truncated pyramid that tapers in depth. The vertical and horizontal dimensions represented by that space, however, are not diminished; in other words, the larger front face and the smaller rear face of the volumetric structure represent equal areas in perceived space by unequal areas in representational space, as suggested by the converging gridlines in the figure. All of the spatial interactions described above (e.g., the collinear propagation of corner and occlusion percepts) would be similarly distorted in this space. Even the angular measure of orthogonality is distorted somewhat by this transformation. The perceived cube depicted in the solid volume of Figure 12B is metrically shrunk in height and width as a function of depth, but because this shrinking is in the same proportion as the shrinking of the space itself, the depicted irregular cube represents a percept of a regular cube with equal sides and orthogonal faces.

The propagation of the field of influence in depth due to a two-dimensional visual input, however, does not shrink with depth. A projection of the trapezoid of Figure 12A would occur in this model as depicted in Figure 12C, projecting the trapezoidal form backward in parallel, independent of the convergence of the space around it. The shaded surfaces in Figure 12C therefore represent the loci of all

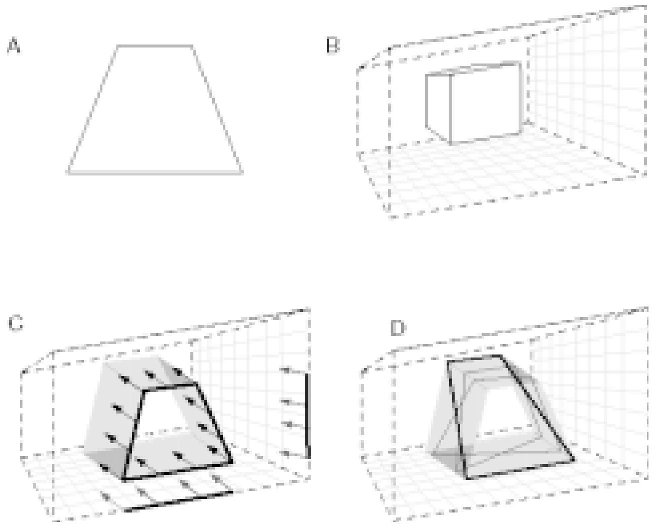

Figure 12. **A.** A trapezoidal stimulus that tends to be perceived as a rectangle viewed in perspective. **B.** The perspective modified spatial representation whose dimensions are shrunken in height and breadth as a function of depth. **C.** The parallel projection of a field of influence into depth of the two-dimensional trapezoidal stimulus. **D.** Several possible perceptual interpretations of the trapezoidal stimulus, one of which (depicted in black outline) represents a regular rectangle viewed in perspective, because the convergence of its sides exactly matches the convergence of the space itself.

possible spatial interpretations of the two-dimensional trapezoid stimulus of Figure 12A, or the extrinsic constraints for the spatial percept due to this stimulus. One possible perceptual interpretation is of a trapezoid parallel to the plane of the page, which can be perceived to be either nearer or farther in depth, but because the scale size shrinks as a function of depth, the percept will be experienced as larger in absolute size (as measured against the shrunken spatial scale) when perceived as farther away and as smaller in absolute size (as measured against the expanded scale) when perceived to be closer in depth. This corresponds to the phenomenon known as Emmert's Law (Coren et al. 1994), whereby a retinal afterimage appears larger when viewed against a distant background than when viewed against a nearer background.

There are also an infinite number of alternative perceptual interpretations of the trapezoidal stimulus, some of which are depicted by the dark shaded lines in Figure 12D. Most of these alternative percepts are geometrically irregular, representing figures with unequal sides and odd angles. But of all the possibilities, there is one special case, depicted by the bold outline in Figure 12D, in which the convergence of the sides of the perceived form happens to coincide exactly with the convergence of the space itself. In other words, this particular percept represents a regular rectangle viewed in perspective, with parallel sides and right-angled corners, whose nearer (bottom) and farther (top) horizontal edges are the same length in the distorted perceptual space. Although this rectangular percept represents the most stable interpretation, other possible interpretations might be suggested by different contexts. The most significant feature of this concept of perceptual processing is that the result of the computation is expressed not in the form of abstract variables encoding the depth and slope of the perceived rectangle, but in the form of an explicit three-dimensional replica of the surface as it is perceived to exist in the world.

8.7. Bounding the representation

An explicit volumetric representation of perceived space as proposed here must necessarily be bounded in some way in order to allow a finite representational space to map to the infinity of external space, as suggested in Figure 2 (in sect. 6.3). The nonlinear compression of the depth dimension observed in phenomenal space can be modeled mathematically with a vergence measure, which maps the infinity of Euclidean distance into a finite bounded range, as suggested in Figure 13A. This produces a representation reminiscent of museum dioramas, like the one depicted in Figure 13B, where objects in the foreground are represented in full depth and the depth dimension is increasingly compressed with distance from the viewer, eventually collapsing into a flat plane corresponding to the background. This vergence measure is presented here merely as a nonlinear compression of depth in a monocular spatial representation, as opposed to a real vergence value measured in a binocular system, although this system could of course serve both purposes in biological vision. Assuming unit separation between the eyes in a binocular system, this compression is defined by the equation

$$\nu = 2 \operatorname{atan}(1/2r),$$

where ν is the vergence measure of depth and r is the Euclidean range, or distance in depth. Because vergence is large at short range and smaller at long range, it is actually the "π-complement" vergence measure ρ that is used in the representation, where $\rho = (\pi - \nu)$ and ρ ranges from 0 at $r = 0$ to π at $r =$ infinity.

What does this kind of compression mean in an isomor-

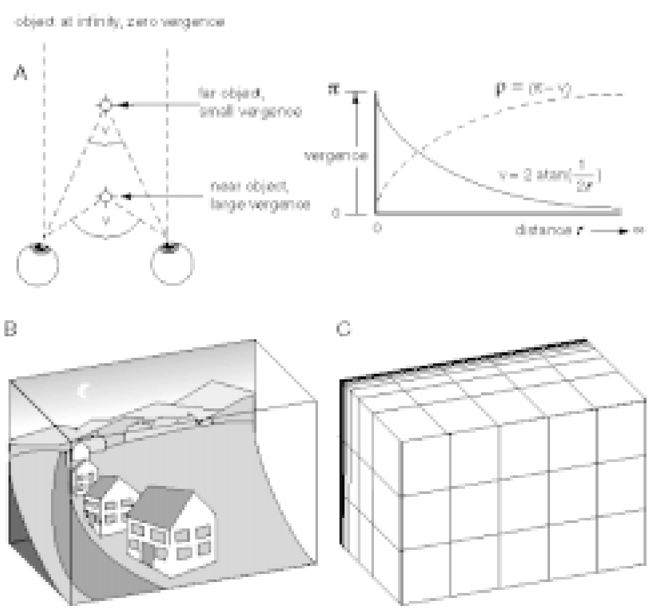

Figure 13. **A.** A vergence representation maps infinite distance into a finite range. **B.** This produces a mapping reminiscent of a museum diorama. **C.** The compressed reference grid in this compressed space defines intervals that are perceived to be of uniform size.

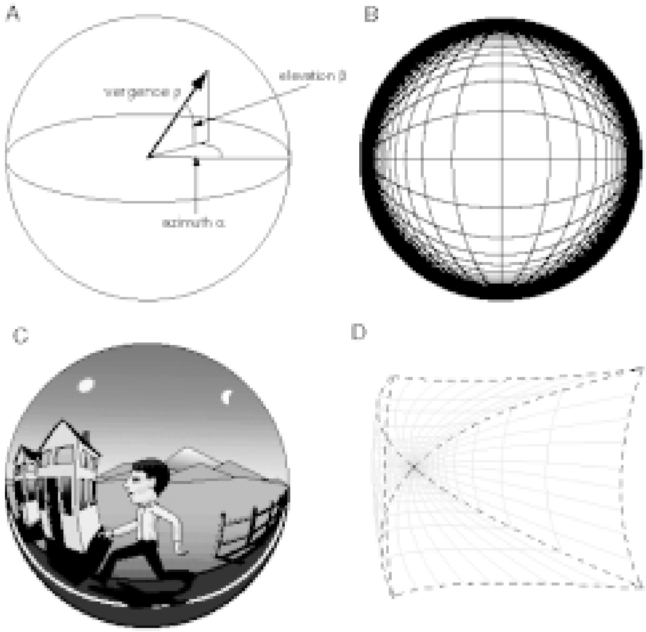

Figure 14. **A.** An azimuth/elevation/vergence representation maps the infinity of three-dimensional Euclidean space into a finite perceptual space. **B.** The deformation of the infinite Cartesian grid caused by the perspective transformation of the azimuth/elevation/vergence representation. **C.** A view of a man walking down a road represented in the perspective distorted space. **D.** A section of the spherical space depicted in the same format as the perspective space shown in Figure 12.

phic representation? If the perceptual frame of reference is compressed along with the objects in that space, then the compression need not be perceptually apparent. Figure 13C depicts this kind of compressed reference grid. The unequal intervals between adjacent depth gridlines define intervals that are perceived to be of equal length, so the flattened cubes defined by the distorted grid would appear perceptually as regular cubes, of equal height, breadth, and depth. This compression of the reference grid to match the compression of space would, in a mathematical system with infinite resolution, completely conceal the compression from the percipient. In a real physical implementation there are two effects of this compression that would remain apparent perceptually, due to the fact that the spatial matrix itself would have to have a finite perceptual resolution. The resolution of depth within this space is reduced as a function of depth, and beyond a certain limiting depth, all objects are perceived to be flattened into two dimensions, with zero extent in depth. This phenomenon is observed perceptually, where the sun, moon, and distant mountains appear as if they are pasted against the flat dome of the sky.

The other two dimensions of space can also be bounded by converting the x and y of Euclidean space into azimuth and elevation angles, α and β, producing an angle/angle/vergence representation, as shown in Figure 14A. Mathematically, this transformation converts the point $P(\alpha,\beta,r)$ in polar coordinates to point $Q(\alpha,\beta,\rho)$ in this bounded spherical representation. In other words, azimuth and elevation angles are preserved by this transformation, and the radial distance in depth r is compressed to the vergence representation ρ as described above. This spherical coordinate system has the ecological advantage that the space near the body is represented at the highest spatial resolution, whereas the less important more distant parts of space are represented at lower resolution. All depths beyond a certain radial distance are mapped to the surface of the representation which corresponds to perceptual infinity.

The mathematical form of this distortion is depicted in Figure 14B, where the distorted grid depicts the perceptual representation of an infinite Cartesian grid with horizontal and vertical gridlines spaced at equal intervals. This geometrical transformation from the infinite Cartesian grid actually represents a unique kind of perspective transformation on the Cartesian grid, with the transformed space appearing as a perspective view of a Cartesian grid when viewed from inside, with all parallel lines converging to a point in opposite directions. The significance of this observation is that by mapping space into a perspective-distorted grid, the distortion of perspective is removed, in the same way that plotting log data on a log plot removes the logarithmic component of the data. Figure 14C shows how this space would represent the perceptual experience of a man walking down a road. If the distorted reference grid of Figure 14B is used to measure lines and distances in Figure 14C, the bowed line of the road on which the man is walking is aligned with the bowed reference grid and therefore is perceived to be straight. The distortion of straight lines into curves in the perceptual representation is not immediately apparent to the percipient because they are perceived to be straight. However, in a global sense there are peculiar distortions apparent to the percipient which are caused by this deformation of Euclidean space, for although the sides of the road are perceived to be parallel, they are also perceived to meet at a point on the horizon.

The fact that two lines can be perceived to be both straight and parallel and yet to converge to a point both in front of and behind the percipient indicates that our internal representation itself must be curved. The proposed representation of space has exactly this property. Parallel lines do not extend to infinity but meet at a point beyond which they are no longer represented. Likewise, the vertical walls of the houses in Figure 14C bow outward, away from the observer, but in doing so they follow the curvature of the reference lines in the grid of Figure 14B and are therefore perceived as being both straight and vertical. Because curved lines in this spherical representation in fact represent straight lines in external space, all of the spatial interactions discussed in the previous section, including the coplanar interactions and the collinear creasing of perceived surfaces, must follow the grain or curvature of collinearity defined within this distorted coordinate system. The distance scale encoded in the grid of Figure 14B replaces the regularly spaced Cartesian grid by a nonlinear collapsing grid whose intervals are spaced ever closer as they approach perceptual infinity but nevertheless represent equal intervals in external space. This nonlinear collapsing scale thereby provides an objective measure of distance in the perspective-distorted perceptual world. The houses in Figure 14C, for example, would be perceived to be approximately the same size and depth, although the more-distant house is experienced at a lower perceptual resolution.

Figure 14D depicts how a slice of Euclidean space of fixed height and width would appear in the perceptual sphere, extending to perceptual infinity in one direction, like a slice cut from the spherical representation of Figure 14C. This slice is similar to the truncated pyramid shape shown in Figure 12B, with the difference that the horizontal and vertical scales of representational space diminish in a nonlinear fashion as a function of distance in depth. In other words, the sides of the pyramid in Figure 14B converge in curves rather than in straight lines, and the pyramid is no longer truncated, but extends in depth all the way to the vanishing point at representational infinity. An input image is projected into this spherical space using the same principles as before.

8.8. Brain anchoring

One of the most disturbing properties of the phenomenal world for models of the perceptual mechanism involves the subjective impression that the phenomenal world rotates relative to our perceived head as our head turns relative to the world, and that objects in perception are observed to translate and rotate while maintaining their perceived structural integrity and recognized identity. This suggests that the internal representation of external objects and surfaces is not anchored to the tissue of the brain, as suggested by current concepts of neural representation, but that perceptual structures are free to rotate and translate coherently relative to the neural substrate, as suggested in Köhler's field theory (Köhler & Held 1947). This issue of brain anchoring is so troublesome that it is often cited as a counterargument for an isomorphic representation because it is so difficult to conceive of the solid spatial percept of the surrounding world having to be reconstructed anew in all its rich spatial detail with every turn of the head (Gibson 1979; O'Regan 1992). However, an argument can be made for the adaptive value of a neural representation of the external world that could break free of the tissue of the sensory or cortical surface in order to lock on to the more meaningful coordinates of the external world, if only a plausible mechanism could be conceived to achieve this useful property.

Even in the absence of a neural model with the required properties, the invariance property can be encoded in a perceptual model. In the case of rotation invariance, this property can be quantified by proposing that the spatial structure of a perceived object and its orientation are encoded as separable variables. This would allow the structural representation to be updated progressively from successive views of an object that is rotating through a range of orientations. The rotation invariance property does not mean, however, that the encoded form has no defined orientation, but rather that the perceived form is presented to consciousness at the orientation and rate of rotation that the external object is currently perceived to possess. In other words, when viewing a rotating object, like a person doing a cartwheel or a skater spinning about her vertical axis, every part of that visual stimulus is used to update the corresponding part of the internal percept even as that percept rotates within the perceptual manifold to remain in synchrony with the rotation of the external object. The perceptual model need not explain how this invariance is achieved neurophysiologically, it must merely express the invariance property computationally, regardless of the "neural plausibility" or computational efficiency of that calculation, for the perceptual model is more a quantitative description of the phenomenon than a theory of neurocomputation.

The property of translation invariance can be similarly quantified in the representation by proposing that the structural representation can be calculated from a stimulus that is translating across the sensory surface, to update a perceptual effigy that translates with respect to the representational manifold, while maintaining its structural integrity. This accounts for the structural constancy of the perceived world as it scrolls past a percipient walking through a scene, with each element of that scene following the proper curved perspective lines as depicted in Figure 2, expanding outward from a point up ahead, and collapsing back to a point behind, as would be seen in a cartoon movie rendition of Figure 2.

The fundamental invariance of such a representation offers an explanation for another property of visual perception, the way that individual impressions left by each visual saccade are observed to appear phenomenally at the appropriate location within the global framework of visual space depending on the direction of gaze. This property can be quantified in the perceptual model as follows. The two-dimensional image from the spherical surface of the retina is copied onto a spherical surface in front of the eyeball of the perceptual effigy, from whence the image is projected radially outward in an expanding cone into the depth dimension of the internal perceptual world, as suggested in Figure 15, as an inverse analog of the cone of light received from the world by the eye. Eye, head, and body orientation relative to the external world are taken into account in order to direct the visual projection of the retinal image into the appropriate sector of perceived space, as determined from proprioceptive and kinesthetic sensations, in order to update the image of the body configuration relative to external space. The percept of the surrounding environment therefore serves as a kind of three-dimensional frame

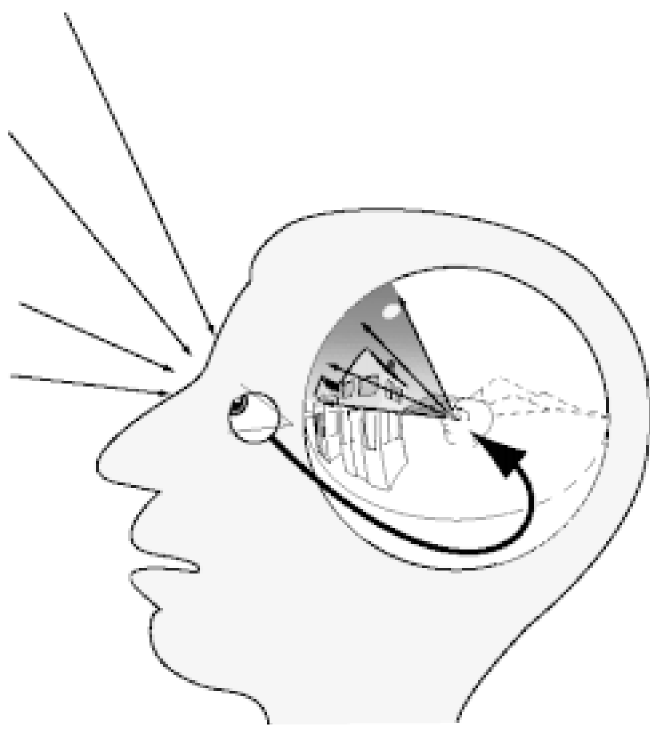

Figure 15. The image from the retina is projected into the perceptual sphere from the center outward in the direction of gaze, as an inverse analog of the cone of light that enters the eye in the external world, taking into account eye, head, and body orientation in order to update the appropriate portion of perceptual space.

buffer expressed in global coordinates. This frame buffer accumulates the information gathered in successive visual saccades and maintains an image of that external environment in the proper orientation relative to a spatial model of the body, compensating for body rotations or translations through the world. Portions of the environment that have not been recently updated gradually fade from perceptual memory, which is why it is easy to bump one's head after bending for some time under an overhanging shelf, or why it is possible to advance only a few steps safely after closing one's eyes while walking.

9. Discussion

The picture of visual processing revealed by the phenomenological approach is radically different from the picture revealed by neurophysiological studies. In fact, the computational transformations observed phenomenologically are implausible in terms of contemporary concepts of neurocomputation and even in terms of computer algorithms. However, the history of psychology is replete with examples of plausibility arguments based on the limited technology of the time, arguments that were later invalidated by the emergence of new technologies. The outstanding achievements of modern technology, especially in the field of information processing systems, might seem to justify our confidence in judging the plausibility of proposed processing algorithms. And yet, despite the remarkable capabilities of modern computers, certain classes of problems appear to be fundamentally beyond the capacity of the digital computer. In fact, the very problems that are most difficult for computers to address, such as extraction of spatial structure from a visual scene, especially in the presence of attached shadows, cast shadows, specular reflections, occlusions, and perspective distortions, as well as the problems of navigation in a natural environment, are problems that are routinely handled by biological vision systems, even those of simpler animals. On the other hand, the kinds of problems that are easily solved by computers, such as perfect recall of vast quantities of meaningless data, perfect memory over indefinite periods, detection of the tiniest variation in otherwise identical data, or exact repeatability of even the most complex computations, are the kinds of problems that are inordinately difficult for biological intelligence, even that of the most complex of animals. It is therefore safe to assume that the computational principles of biological vision are fundamentally different from those of digital computation, and plausibility arguments predicated on contemporary concepts of what is computable are not applicable to biological vision. If we allow that our contemporary concepts of neurocomputation are so embryonic that they should not restrict our observations of the phenomenal properties of perception, the evidence for a Gestalt Bubble model of perceptual processing becomes overwhelming.

The phenomena of hallucinations and dreams demonstrate that the mind is capable of generating complete spatial percepts of the world, including a percept of the body and the space around it (Revonsuo 1995). It is unlikely that this remarkable capacity is used only to create such illusory percepts. More likely, dreams and hallucinations reveal the capabilities of an imaging system that is normally driven by the sensory input, generating perceptual constructs that are coupled to external reality.

Studies of mental imagery (Kosslyn 1980; 1994) have characterized the properties of this imaging capacity and have confirmed the three-dimensional nature of the encoding and processing of mental imagery. Pinker (1980) showed that the scanning time between objects in a remembered three-dimensional scene increases linearly with increasing distance between objects in three dimensions. Shepard and Metzler (1971) showed that the time for rotation of mental images is proportional to the angle through which they are rotated. Kosslyn (1975) showed that it takes time to expand the size of mental images, and that smaller mental images are more difficult to scrutinize. As unexpected as these findings may seem for theorists of neural representation, they are perfectly consistent with the subjective experience of mental imagery. On the basis of these findings, Pinker (1988) derived a volumetric spatial medium to account for the observed properties of mental image manipulation which is very similar to the model proposed here, with a volumetric azimuth/elevation coordinate system that is addressable both in subjective viewer-centered and objective viewer-independent coordinates, and with a compressive depth scale.

The phenomenon of hemi-neglect (Heilman & Watson 1977; Heilman et al. 1985; Kolb & Whishaw 1996; McFie & Zangwill 1960) reveals the effects of damage to spatial representation, as the capacity to represent spatial percepts in one half of phenomenal space is destroyed. Hemi-neglect patients are not simply blind to objects to one side, but are blind to the very existence of a space in that direction as a potential holder of objects. For example, they typically eat food only from the right half of the plate, and express surprise at the unexpected appearance of more food when the plate is rotated 180 degrees. This condition persists even when patients are cognitively aware of their deficit (Sacks 1985). Bisiach and Luzatti (1978) and Bisiach et al. (1981) showed how this condition can also impair mental imaging ability. They described a neglect patient who, when instructed to recall a familiar scene viewed from a certain direction, could recall only objects from the right half of his remembered space. When instructed to mentally turn around and face in the opposite direction, the patient could then recall only objects from the other side of the scene, objects that now fell in the right half of his mental image space. The condition of hemi-neglect therefore suggests damage to the left half of a three-dimensional imaging mechanism that is used both for perception and for the generation of mental imagery. Note that hemi-neglect also includes a neglect of the left side of the body, which is consistent with the fact that the body percept is included as an integral part of the perceptual representation.

The condition of hemi-neglect initially caused a great stir in psychological circles because it appeared to be concrete evidence for an explicit spatial representation in the brain (Bisiach & Luzzatti 1978; Bisiach et al. 1981; Denny-Brown & Chambers 1958; de Renzi 1982). It is curious that half of phenomenal space should have to disappear for psychologists to take account of its existence in the first place. But after the initial excitement, the naïve realists quickly marshalled their defenses with an array of arguments that many believe disposed of the troublesome issue of hemi-neglect. Some argued that hemi-neglect is not a failure of spatial representation but rather an imbalance of attention, or "orienting response"; that is, half of phenomenal space does not actually disappear, but the neglect patient is merely inclined to ignore its presence (Heilman & Watson 1977; Heilman et al. 1985; Kinsbourne 1987; 1993). Even if this argument were valid, it would not account for the presence in visual consciousness of the spatial structure of the phenomenal world whenever it is *not* being ignored or neglected; it would merely offer a convenient escape clause to make neglect syndrome seem no *more* mysterious than normal spatial perception. Others argue that the phenomenon of hemi-neglect fractionates to a number of distinct patterns of impairment (Vallar 1998, p. 88). For example, many neglect patients can describe the global gestalt of a figure, but when copying its local features, they leave out those on the left side (Marshall & Halligan 1995). Present accounts of the multiple forms of neglect refer to several spatial maps and their interaction (e.g., Ladavas et al. 1997). This highlights a conflict between the phenomenal and neurophysiological evidence, the former presenting a unified spatial structure in visual experience and the latter suggesting discrete mechanisms in different cortical areas. To the naive realist this suggests that the spatial percept must be somehow illusory, which thereby supposedly relieves neuroscience from any obligation to account for its manifest properties.

What is curious about the debate over neglect is the passion that it engenders. The evidence presented by each side never seems to convince the opposition, because the debate is not really about neglect but about its implications for perceptual representation, and that issue is not so much a matter of experimental evidence but of the interpretation of that evidence, or the foundational assumptions with which one comes to the debate in the first place. Whatever the physiological reality behind the phenomenon of hemi-neglect, the Gestalt Bubble model offers at least a concrete description of this otherwise paradoxical phenomenon.

The idea that this spatial imaging system employs an explicit volumetric spatial representation is suggested by the fact that disparity tuned cells have been found in the cortex (Barlow et al. 1967), as predicted by the Projection Field Theory of binocular vision (Boring 1933; Charnwood 1951; Julesz 1971; Kaufman 1974; Marr & Poggio 1976), which is itself a volumetric model. Psychophysical evidence for a volumetric representation comes from the fact that perceived objects in depth exhibit attraction and repulsion in depth (Mitchison 1993; Westheimer & Levi 1987), in a manner that is suggestive of a short-range attraction and longer range repulsion in depth, analogous to the center-surround processing in the retina. Brookes and Stevens (1989) discussed the analogy between brightness and depth perception; they showed that a number of brightness illusions attributed to such center-surround processing have corresponding illusions in depth. Similarly, Anstis and Howard (1978) demonstrated a Craik-O'Brien-Cornsweet illusion in depth (cf. Cornsweet 1970) by cutting the near surface of a block of wood with a depth profile matching the brightness cusp of the brightness illusion, resulting in an illusory percept of a difference in depth of the surfaces on either side of the cusp. As in the brightness illusion, the depth difference at the cusp appears to propagate a perceptual influence out to the ends of the block, suggesting a spatial diffusion of depth percept between depth edges.

The many manifestations of constancy in perception have always posed a serious challenge for theories of perception because they reveal that the percept exhibits properties of the distal object rather than of the proximal stimulus, or pattern of stimulation on the sensory surface. The Gestalt Bubble model explains this by the fact that the information encoded in the internal perceptual representation itself reflects the properties of the distal object rather than the proximal stimulus. Size constancy is explained by the fact that objects perceived to be more distant are represented closer to the outer surface of the perceptual sphere, where the collapsing reference grid corrects for the shrinkage of the retinal image due to perspective. An object perceived to be receding in depth, therefore, is expected perceptually to shrink in retinal size along with the shrinking of the grid in depth, and, conversely, shrinking objects tend to be perceived as receding. Rock and Brosgole (1964) showed that perceptual grouping by proximity is determined not by proximity in the two-dimensional retinal projection of the figure, but rather by the three-dimensional perceptual interpretation. A similar finding was shown by Green and Odum (1986). Shape constancy is exemplified by the fact that a rectangle seen in perspective is not perceived as a trapezoid, as its retinal image would suggest. The Müller-Lyer and Ponzo illusions are explained in similar fashion (Gillam 1971; 1980; Gregory 1963; Tausch 1954), the converging lines in those figures suggesting a surface sloping in depth, so that features near the converging ends are measured against a more compressed reference grid than the corresponding feature near the diverging ends of those lines.

Several researchers have presented psychophysical evidence for a spatial interpolation in depth, which is difficult to account for except with a volumetric representation in which the interpolation is computed explicitly in depth (Attneave 1982). Kellman et al. (1996) have demonstrated a coplanar completion of perceived surfaces in depth in a manner analogous to the collinear completion in the Kanizsa figure. Barrow and Tenenbaum (1981, p. 94 and Fig. 6.1) showed how a two-dimensional wire-frame outline held in front of a dynamic random noise pattern stimulates a three-dimensional surface percept spanning the outline like a soap film, and that the perceived surface undergoes a Necker reversal together with the reversal of the perimeter wire. Ware and Kennedy (1978) showed that a three-dimensional rendition of the Ehrenstein illusion, constructed of a set of rods converging on a circular hole, creates a three-dimensional version of the illusion that is perceived as a spatial structure in depth, even when rotated out of the fronto-parallel plane, complete with a perception of brightness at the center of the figure. This illusory percept appears to hang in space like a faintly glowing disk in depth, reminiscent of the neon color spreading phenomenon. A similar effect can be achieved with a three-dimensional rendition of the Kanizsa figure. If the Ehrenstein and Kanizsa figures are explained by spatial interpolation in models such as that of Grossberg and Mingolla (1985), then the corresponding three-dimensional versions of these illusions must involve a volumetric computational matrix to perform the interpolation in depth.

Collett (1985) investigated the interaction between monocular and binocular perception using stereoscopically presented line drawings in which some features were presented only monocularly – that is, their depth information is unspecified. Collett showed that such features tend to appear perceptually at the same depth as adjacent binocularly specified features, as if under the influence of an attractive force in depth generated by the binocular feature. In ambiguous cases the percept is often multistable, jumping back and forth in depth, especially when monocular perspective cues conflict with the binocular disparity information. The perceived depth of the monocularly specified surfaces is measured psychophysically using a three-dimensional disparity-specified cursor whose depth is adjusted by the subject to match the depth of the perceived surface at that point. Subjects reported a curious interaction between the cursor and the perceived surface, which was observed to flex in depth toward the cursor at small disparity differences, in the manner of the attraction and repulsion in depth reported by Westheimer and Levi (1987). This dynamic influence is suggestive of a grouping by proximity mechanism, expressed as a fieldlike attraction between perceived features in depth; the flexing of the perceived surface near the three-dimensional cursor, as well as the multistability in the presence of conflicting perspective and disparity cues, is suggestive of a Gestalt Bubble model.

Carman and Welch (1992) employed a similar cursor to measure the perceived depth of three-dimensional illusory surfaces seen in Kanizsa figure stereograms, whose inducing edges are tilted in depth in a variety of configurations, as shown in Figure 16A. Note how the illusory surface completes in depth by coplanar interpolation defining a smooth curving surface. The subjects in this experiment also reported a flexing of the perceived surface in depth near the disparity-defined cursor. Equally interesting is the porthole illusion seen in the reverse-disparity version of this figure, where the circular completion of the portholes generates an ambiguous unstable semitransparent percept at the center of the figure, which is characteristic of the Gestalt Bubble model. Kellman and Shipley (1991) and Idesawa (1991) reported the emergence of more complex illusory surfaces in depth, using similar illusory stereogram stimuli as shown in Figures 16B and 16C. It is difficult to deny the reality of a precise high-resolution spatial interpolation mechanism in the face of these compelling illusory percepts. Whatever the neurophysiological basis of these phenomena, the Gestalt Bubble model offers a mathematical framework for a precise description of the information encoded in these elaborate spatial percepts, independent of the confounding factor of neurophysiological considerations.

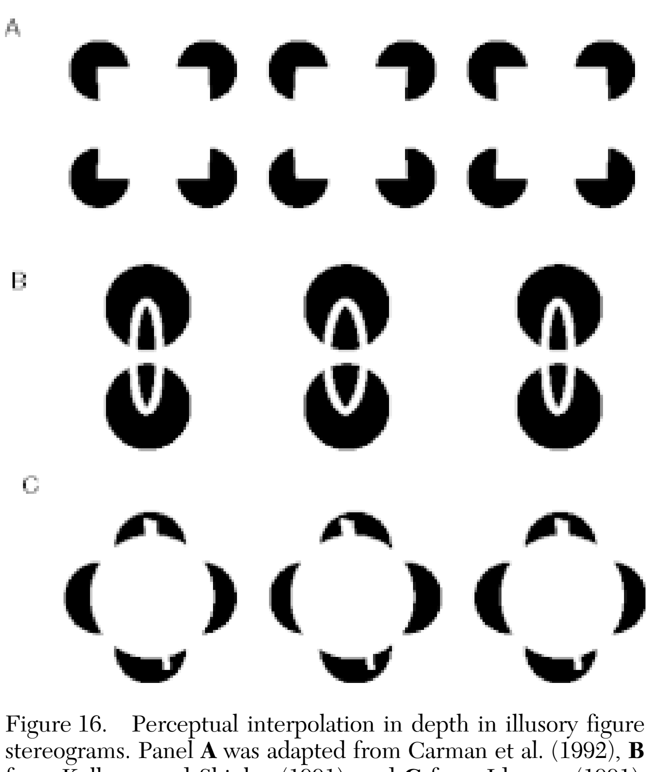

Figure 16. Perceptual interpolation in depth in illusory figure stereograms. Panel **A** was adapted from Carman et al. (1992), **B** from Kellman and Shipley (1991), and **C** from Idesawa (1991). Opposite disparity percepts are achieved by binocular fusion of either the first and second, or the second and third columns of the figure.

The sophistication of the perceptual reification capacity is revealed by the apparent-motion phenomenon (Coren et al. 1994), which in its simplest form consists of two alternately flashing lights that generate a percept of a single light moving back and forth between the flashing stimuli. With more complex variations of the stimulus, the illusory percept is observed to change color or shape in midflight, to carry illusory contours, or to carry a texture region bounded by an illusory contour between the alternately flashing stimuli (Coren et al. 1994). Most pertinent to the discussion of a spatial representation is the fact that the illusory percept is observed to make excursions into the third dimension when that produces a simpler percept. For example, if an obstacle is placed between the flashing stimuli so as to block the path between them, the percept is observed to pass either in front of or behind the obstacle in depth. Similarly, if the two flashing stimuli are in the shape of an angular feature like $<$ or $>$, the angle is observed to rotate in depth between the flashing stimuli, preserving a percept of a rigid rotation in depth, in preference to a morphological deformation in two dimensions. The fact that the percept transitions so readily into depth suggests the fundamental nature of the depth dimension for perception.

Although apparent-motion effects reify whole perceptual gestalts, the elements of this reification, such as the fieldlike diffusion of perceived surface properties, are seen in such diverse phenomena as the perceptual filling-in of the Kanizsa figure (Takeichi et al. 1992), the Craik-O'Brien-Cornsweet effect (Cornsweet 1970), the neon color spreading effect (Bressan 1993), the filling-in of the blind spot (Ramachandran 1992), color bleeding due to retinal stabilization (Heckenmuller 1965; Yarbus 1967), the motion capture effect (Ramachandran & Anstis 1986), and the aperture problem in motion perception (Movshon et al. 1986). In all of these phenomena, a perceived surface property (brightness, transparency, color, motion, etc.) is observed to spread from a localized origin, not into a fuzzy ill-defined region but, rather, into a sharply bounded region containing a homogeneous perceptual quality; and this filling-in occurs as readily in depth in a perspective view as in the frontoparallel plane. The time has come to recognize that these phenomena do not represent exceptional or special cases, nor are they illusory in the sense of lacking a neurophysiological counterpart. Rather, these phenomena reveal a general principle of neurocomputation that is ubiquitous in biological vision.

Evidence for the spherical nature of perceived space dates back to observations by Helmholtz (1925). A subject in a dark room is presented with a horizontal line of point-lights at eye level in the frontoparallel plane, and instructed to adjust their displacement in depth, one by one, until they are perceived to lie in a straight line in depth. The result is a line of lights that curves inward toward the observer, the amount of curvature being a function of the distance of the

https://doi.org/10.1017/S0140525X03410091

line of lights from the observer. Helmholtz recognized this phenomenon as evidence of the non-Euclidean nature of perceived space. The Hillebrand-Blumenfeld *alley experiments* (Blumenfeld 1913; Hillebrand 1902) extended this work with different configurations of lights, and later mathematical analysis of the results (Blank 1958; Luneburg 1950) characterized the nature of perceived space as Riemannian with constant Gaussian curvature (see Foley 1978; Graham 1965; and Indow 1991 for a review). In other words, perceived space bows outward from the observer, with the greatest distortion observed proximal to the body, as suggested by the Gestalt Bubble model. Heelan (1983) presented a more modern formulation of the hyperbolic model of perceived space, and provided further supporting evidence from art and illusion.

It is perhaps too early to say definitively whether the model presented here can be formulated to address all of the phenomena outlined above. What is becoming increasingly clear, however, is the inadequacy of the conventional feed-forward abstraction approach to account for these phenomena, and that novel and unconventional approaches to the problem should be given serious consideration. The general solution offered by the Gestalt Bubble model to all of these problems in perception is that the internal perceptual representation encodes properties of the distal object rather than of the proximal stimulus, and that the computations of spatial perception are most easily performed in a fully spatial matrix in a manner consistent with the subjective experience of perception.

10. Conclusion

I have presented an elaborate model of perception that incorporates many of the concepts and principles introduced by the original Gestalt movement. Though the actual mechanisms of the proposed model remain somewhat vague and poorly specified, a number of prominent aspects of visual experience which are generally ignored by other models are accounted for by this approach to modeling perception. These are summarized as follows.

1. When we view a three-dimensional surface, our subjective experience of that surface simultaneously encodes every point on that surface in three dimensions at a high resolution; in other words, our subjective experience of the world around us is perceived not as a flattened "2-D sketch," nor a nonspatial abstraction, but as a solid spatial world that appears to surround us in all directions.

2. Volumes of empty space are perceived with the same geometrical fidelity as volumes of solid matter.

3. Multiple transparent surfaces can be perceived simultaneously as distinct spatial structures at high resolution.

4. The infinity of external space is perceived as a finite, but fully spatial, representation that appears near-Euclidean near the body but becomes progressively flattened with distance from the body, the entire percept being bounded by a spherical shell representing perceptual infinity.

5. Parallel lines are perceived to meet at perceptual infinity, but at the same time they are perceived as parallel and with uniform separation throughout their entire length.

6. An illusory entity, like the Kanizsa figure or the apparent-motion illusion, is not experienced as a cognitive abstraction but is experienced perceptually as a solid spatial surface at high resolution, virtually indistinguishable from a real physical surface or object.

7. The subjective reversal of a multistable percept is not experienced as a change in a cognitive interpretation or the flipping of a single cognitive variable but, rather, it is vividly experienced as an inversion of a perceptual data structure, changing the perceived depth of every point on the perceived structure.

These phenomena are so immediately manifest in the subjective experience of perception that they need hardly be tested psychophysically. And yet, curiously, these most obvious properties of perception have been systematically ignored by neural modelers, even though the central significance of these phenomena was highlighted decades ago by the Gestaltists. There are two reasons why these prominent aspects of perception have been consistently ignored. The first results from the outstanding success of the single-cell recording technique, which shifted theoretical emphasis from fieldlike theories of whole aspects of perception to pointlike theories of the elements of neural computation. Like the classical Introspectionists, who refused to acknowledge perceptual experiences that were inconsistent with their preconceived notions of sensory representation, the Neuroreductionists of today refuse to consider aspects of perception that are inconsistent with current theories of neural computation, and some of them are even prepared to deny consciousness itself in a heroic attempt to save the sinking paradigm.

There is another factor that has made it possible to ignore these most salient aspects of perception, which is that perceptual entities, such as the solid volumes and empty spaces we perceive around us, are easily confused with real objects and spaces in the objective external world. The illusion of perception is so compelling that we mistake the percept of the world for the real world itself. And yet this naïve realist view that we can somehow perceive the world directly is inconsistent with the physics of perception. If perception is a consequence of neural processing of the sensory input, a percept cannot in principle escape the confines of our head to appear in the world around us, any more than a computation in a digital computer can escape the confines of the computer. Therefore, we cannot in principle have direct experience of objects in the world itself but only of the internal effigies of those objects generated by mental processes. The world we see around us can only be an elaborate, though very compelling, illusion, which must in reality correspond to perceptual data structures and processes occurring actually within our own heads. As soon as we examine the world we see around us, not as a physical scientist observing the physical world, but as a *perceptual* scientist observing a rich and complex internal percept, only then does the rich spatial nature of perceptual processing become immediately apparent. It was this central insight into the illusion of consciousness that formed the key inspiration of the Gestalt movement, from which all of their other ideas were developed. The central message of Gestalt theory is that the primary function of perceptual processing is the generation of a miniature, virtual-reality replica of the external world inside our head, and that the world we see around us is not the real external world but is exactly that miniature internal replica (Lehar 2003b). It is only in this context that the elaborate model presented here begins to seem plausible.

APPENDIX

A1. The coplanarity field

The mathematical form of the coplanarity interaction field can be described as follows. Consider the field strength F due to an element in the opaque state at some point in the volume of the spatial matrix, with a certain surface orientation, depicted in Figure A1, panel A, as a vector, representing the normal to the surface encoded by that element. The strength of the field F should peak within the plane at right angles to this normal vector (depicted as a circle in panel A) as defined in polar coordinates by the function $F_{\alpha} = \sin(\alpha)$, where α is the angle between the surface normal and some point in the field, that ranges from zero, parallel to the normal vector, to π, in the opposite direction. The sine function peaks at $\alpha = \pi/2$, as shown in panel B, producing an equatorial belt around the normal vector as suggested schematically in cross section in panel C, where the gray shading represents the strength of the field. The strength of the field should actually decay with distance from the element, for example with an exponential decay function, as defined by the equation $F_{\alpha r} = e^{-r2} \sin(\alpha)$ as shown in panel D, where r is the radial distance from the element. This produces a fading equatorial band, as suggested schematically in cross section in panel E. The equatorial belt of the function described so far would be rather fat, resulting in a lax or fuzzy coplanarity constraint, but the constraint can be stiffened by raising the sine to some positive power P, producing the equation $F_{\alpha r} = e^{-r2} \sin(\alpha)^P$, which will produce a sharper peak in the function as shown in panel F, producing a sharper in-plane field depicted schematically in cross section in panel G. To control runaway positive feedback and suppress the uncontrolled proliferation of surfaces, the field function should be normalized, which will project inhibition in directions outside the equatorial plane. This can be achieved with the equation $F_{\alpha r} = e^{-r2}\, 2 \sin(\alpha)^P - 1$, which has the effect of shifting the equatorial function halfway into the negative region, as shown in panel H, thereby producing the field suggested in cross section in panel I.

The field described so far is unoriented – that is, it has a magnitude but no direction at any sample point (r,α). What is actually required is a field with a direction; such a field would have maximal influence on adjacent elements that are oriented parallel to it, elements that are coplanar with it in both position and orientation. We can describe this orientation of the field with the parameter θ, which represents the orientation at which the field F is sampled, expressed as an angle relative to the normal vector. In other words, the strength of the influence F exerted on an adjacent element located at a point (r,α) varies with the deviation θ of that element from the direction parallel to the normal vector, as shown in Figure A2, such that the maximal influence is felt when the two elements are parallel (i.e., when $\theta = 0$, as in Fig. A2, panel A) and falls off smoothly as the other element's orientation deviates from that orientation, as in Figure A2, panels B and C.

This can be expressed with a cosine function, such that the influence F of an element on another element in a direction α and at separation r from the first element, and with a relative orientation θ, would be defined by

$$F_{\alpha r\theta} = e^{-r2}\, [2 \sin(\alpha)^P - 1]\, |\cos(\theta)^Q|. \qquad \text{(Eq. 1)}$$

This cosine function allows the coplanar influence to propagate to near-coplanar orientations, thereby allowing surface completion to occur around smoothly curving surfaces. The tolerance to such curvature can also be varied parametrically by raising the cosine function to a positive power Q, as shown in equation 1. So the in-plane stiffness of the coplanarity constraint is adjusted by parameter P, and the angular stiffness is adjusted by parameter Q. The absolute value on the cosine function in equation 1 allows interaction between elements when θ is between $\pi/2$ and π.

A2. The occlusion field

The orthogonality and occlusion fields have one less dimension of symmetry than does the coplanarity field, and therefore they are defined with reference to two vectors through each element at right angles to each other, as shown in Figure A3, panel A. For the orthogonality field, these vectors represent the surface normals to the two orthogonal planes of the corner; for the occlusion field one vector is a surface normal and the other vector points within that plane in a direction orthogonal to the occlusion edge. The occlusion field G around the local element is defined in polar coordinates from these two vector directions, using the angles α and β respectively, as shown in Figure A3, panel A.

The plane of the first surface is defined as for the coplanarity field, with the equation $G_{\alpha\beta r} = e^{-r2} \sin(\alpha)^P$. For the occlusion field this planar function should be split in two, as shown in panel B, to produce a positive half and a negative half, so that this field will promote surface completion in one direction only and will actually suppress surface completion in the negative half of the field.

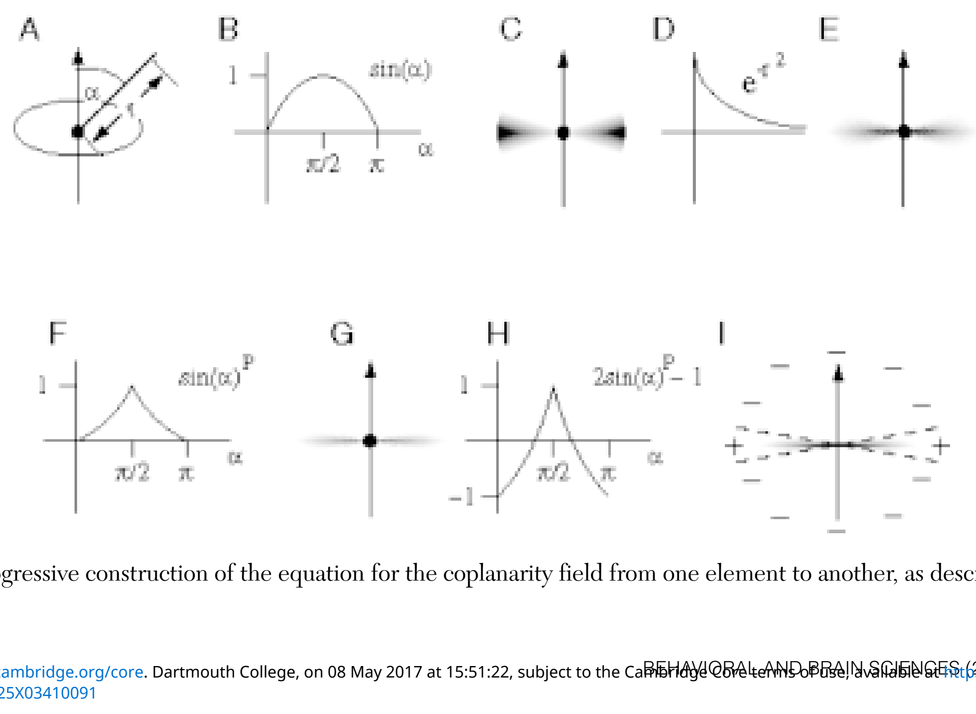

Figure A1. Progressive construction of the equation for the coplanarity field from one element to another, as described in the text.

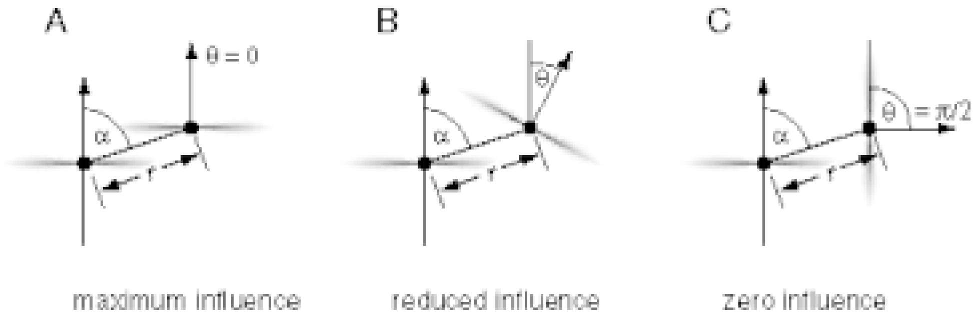

Figure A2. Orientation of the field of influence between one element and another. For an element located at polar coordinates (r,θ), the influence varies as a cosine function of θ, the angle between the normal vectors of the two interacting elements.

This can be achieved by multiplying the above equation by the *sign* (plus or minus, designated by the function sgn( )) of a cosine on the orthogonal vector; that is, $G_{\alpha\beta r} = e^{-r2} \sin(\alpha)^P \operatorname{sgn}(\cos(\beta))$. Because of the negative half-field in this function, there is no need to normalize the equation. However, the oriented component of the field can be added as before, resulting in the equation

$$G_{\alpha\beta r\theta} = e^{-r2} [\sin(\alpha)^P \operatorname{sgn}(\cos(\beta))] \mid \cos(\theta)^Q \mid. \qquad \text{(Eq. 2)}$$

Again, the maximal influence will be experienced when the two elements are parallel in orientation, when $\theta = 0$. As before, the orientation cosine function is raised to the positive power Q to allow parametric adjustment of the stiffness of the coplanarity constraint.

A3. The orthogonality field

The orthogonality field H can be developed in a similar manner, beginning with the planar function divided into positive and negative half-fields – that is, with the equation $H_{\alpha\beta r} = e^{-r2} \sin(\alpha)^P \operatorname{sgn}(\cos(\beta))$ – but then adding another similar plane from the orthogonal surface normal, producing the equation $H_{\alpha\beta r} = e^{-r2} [\sin(\alpha)^P \operatorname{sgn}(\cos(\beta)) + \sin(\beta)^P \operatorname{sgn}(\cos(\alpha))]$. This produces two orthogonal planes, each with a negative half-field, as shown schematically in Figure A3, panel C. Finally, this equation must be modified to add the oriented component to the field, represented by the vector θ, such that the maximal influence on an adjacent element will be experienced when that element is either within one positive half-plane and at one orientation or is within the other positive half-plane and at the orthogonal orientation. The final equation for the orthogonality field is therefore defined by

$$H_{\alpha\beta r\theta} = e^{-r2} [\sin(\alpha)^P \operatorname{sgn}(\cos(\beta)) \mid \cos(\theta)^Q \mid + [\sin(\beta)^P \operatorname{sgn}(\cos(\alpha)) \mid \cos(\theta)^Q \mid]]. \qquad \text{(Eq. 3)}$$

A4. Edge consistency and inconsistency constraints

There is another aspect of the fieldlike interaction between elements that remains to be defined. Both the orthogonal and occlusion states are promoted by appropriately aligned neighboring elements in the coplanar state. Orthogonal and occlusion elements should also feel the influence of neighboring elements in the orthogonal and occlusion states, because a single edge should have a tendency to become either an orthogonal corner percept or an occlusion edge percept along its entire length. Therefore, orthogonal or occlusion elements should promote like states and inhibit unlike states in adjacent elements along the same corner or edge. The interaction between like-state elements along the edge will be called the *edge-consistency constraint,* and the corresponding field of influence will be designated E; the complementary interaction between unlike-state elements along the edge is called the *edge-inconsistency constraint,* and its corresponding edge-inconsistency field will be designated I. These interactions are depicted schematically in Figure A4.

The spatial direction along the edge can be defined by the product of the two sine functions, $\sin(\alpha)\sin(\beta)$, defining the orthogonal planes, denoting the zone of intersection of those two orthogonal planes, as suggested in Figure A4, panel E. Again, this field can be sharpened by raising these sine functions to a positive power P, and it can be localized by applying the exponential decay function. The edge consistency constraint E therefore has the form $E_{\alpha\beta r} = e^{-r2} [\sin(\alpha)^P \sin(\beta)^P]$. As for the orientation of the edge-consistency field, this will depend now on two angles, θ and ϕ, representing the orientations of the two orthogonal vectors of the adjacent orthogonal or occlusion elements relative to the two normal vectors respectively. Both the edge-consistency and the edge-inconsistency fields, whether excitatory between like-state elements or inhibitory between unlike-state elements, should peak when both pairs of reference vectors are parallel to the normal vectors of the central element – that is, when θ and ϕ are both

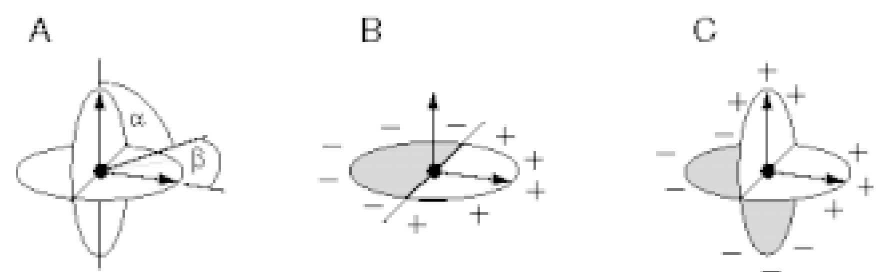

Figure A3. **A.** Polar coordinate reference vectors through each element. **B.** Occlusion field. **C.** Orthogonality field.

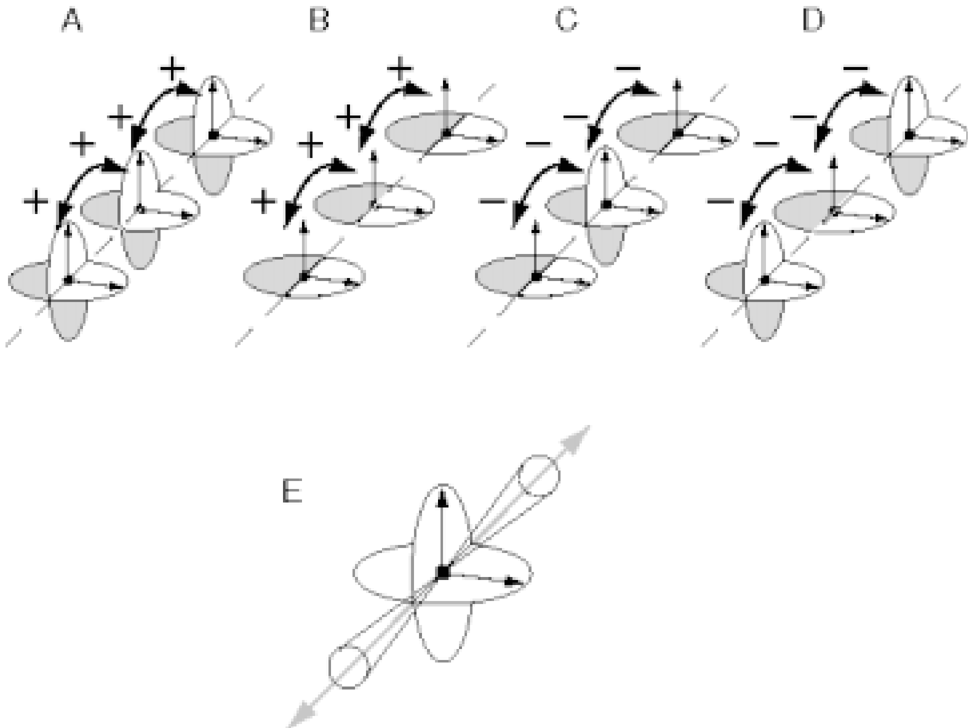

Figure A4. **A, B.** Edge consistency constraint as an excitatory influence between like-state elements along a corner or edge percept. **C, D.** Edge inconsistency constraint as an inhibitory influence between unlike-state elements along a corner or edge percept. **E.** The direction along the edge expressed as the intersection of the orthogonal planes defined by the sine functions on the two orthogonal vectors.

equal to zero. The full equation for the edge-consistency field E would therefore be

$$E_{\alpha\beta r\theta\phi} = e^{-r2}\,[\sin(\alpha)^P \sin(\beta)^P]\cos(\theta)^Q \cos(\phi)^Q, \qquad \text{(Eq. 4)}$$

where this equation is applied only to like-state edge or corner elements. The edge-inconsistency field I would be given by

$$I_{\alpha\beta r\theta\phi} = e^{-r2}\,[\sin(\alpha)^P \sin(\beta)^P]\cos(\theta)^Q \cos(\phi)^Q \qquad \text{(Eq. 5)}$$

applied only to unlike-state elements. The total influence R on an occlusion element therefore is calculated as the sum of the influence of neighboring coplanar, orthogonal, and occlusion state elements as defined by

$$R_{\alpha\beta r\theta\phi} = G_{\alpha\beta r\theta\phi} + E_{\alpha\beta r\theta\phi} - I_{\alpha\beta r\theta\phi}, \qquad \text{(Eq. 6)}$$

and the total influence S on an orthogonal state element is defined by

$$S_{\alpha\beta r\theta\phi} = H_{\alpha\beta r\theta\phi} + E_{\alpha\beta r\theta\phi} - I_{\alpha\beta r\theta\phi}. \qquad \text{(Eq. 7)}$$

A5. Influence of the visual input

A two-dimensional visual edge has an influence on the three-dimensional interpretation of a scene because an edge is suggestive of either a corner or an occlusion at some orientation in three dimensions whose two-dimensional projection coincides with that visual edge. This influence, however, is quite different from the local fieldlike influences described above, because the influence of a visual edge should penetrate the volumetric matrix with a planar field of influence to all depths and should activate all local elements within the plane of influence that are consistent with that edge. Subsequent local interactions between those activated elements serve to select which subset of them should finally represent the three-dimensional percept corresponding to the two-dimensional image. For example, a vertical edge as shown in Figure A5, panel A, would project a vertical plane of influence, as suggested by the light shading, into the depth dimension of the volumetric matrix, where it stimulates the orthogonal and occlusion states that are consistent with that visual edge. It would stimulate corner and occlusion states at all angles about a vertical axis, as

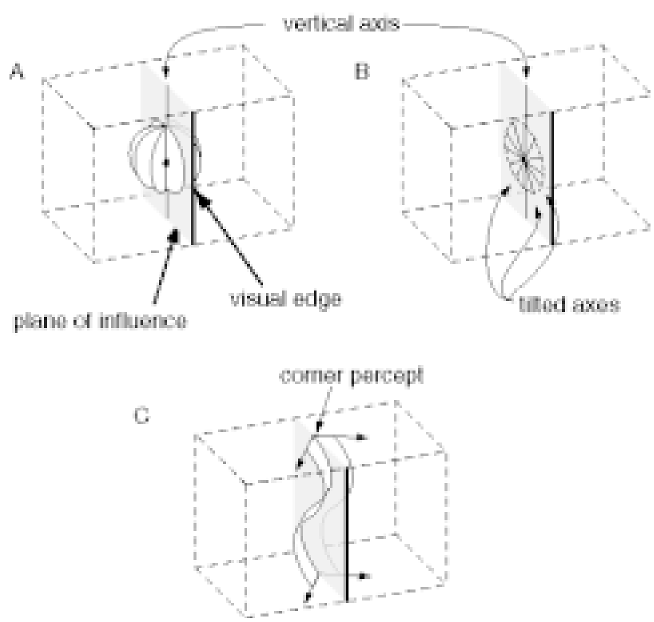

Figure A5. The influence of a visual edge, in this case a vertical edge, is to stimulate local elements in the occlusion or corner percept states at orientations about a vertical axis, panel **A,** or about a tilted axis, panel **B,** within the plane of influence of the edge. At equilibrium, panel **C,** a single unified percept emerges, in this case of a perceived corner at some depth and tilt in the volume of the matrix.

shown in panel A, where the circular disks represent different orientations of the positive half-fields of either corner or occlusion fields.

However, a vertical edge would also be consistent with corners or occlusions about axes tilted relative to the image plane but within the plane of influence, as depicted in panel B. The same kind of stimulation would occur at every point within the plane of influence of the edge, although only one point is depicted in the figure. When all elements consistent with this vertical edge have been stimulated, the local fieldlike interactions between adjacent stimulated elements will tend to select one edge or corner at some depth and at some tilt, thereby suppressing alternative edge percepts at that two-dimensional location at different depths and at different tilts. At equilibrium, some arbitrary edge or corner percept will emerge within the plane of influence as suggested in panel C, which depicts only one such possible percept, and edge consistency interactions will promote like-state elements along that edge, producing a single emergent percept consistent with the visual edge. In the absence of additional influences, for example in the isolated local case depicted in panel C, the actual edge that emerges will be unstable; it could appear anywhere within the plane of influence of the visual edge through a range of tilt angles and could appear as either an occlusion or a corner edge. However, when it does appear it propagates its own fieldlike influence into the volumetric matrix. In this example the corner percept would propagate a planar percept of two orthogonal surfaces that will expand into the volume of the matrix, as suggested by the arrows in panel C. The final percept therefore will be influenced by the global pattern of activity; that is, the final percept will construct a self-consistent perceptual whole whose individual parts reinforce one another by mutual activation by way of the local interaction fields, although that percept would remain unstable in all unconstrained dimensions. For example, the corner percept depicted in panel C would snake back and forth unstably within the plane of influence, rotate back and forth along its axis through a small angle, and flip alternately between the corner and occlusion states, unless the percept is stabilized by other features at more remote locations in the matrix.

ACKNOWLEDGMENT
Kind thanks to Eli Peli and the Schepens Eye Research Institute for help and support.

Open Peer Commentary

Commentary submitted by the qualified professional readership of this journal will be considered for publication in a later issue as Continuing Commentary on this article. Integrative overviews and syntheses are especially encouraged.

Phenomenology is art, not psychological or neural science

David A. Booth

School of Psychology, University of Birmingham, Edgbaston, Birmingham, B15 2TT, United Kingdom. **D.A.Booth@Bham.ac.UK**

Abstract: It is tough to relate visual perception or other achievements to physiological processing in the central nervous system. The diagrammatic, algebraic, and verbal pictures of how sights seem to Lehar do not advance understanding of how we manage to see what is in the world. There are well-known conceptual reasons why no such purely introspective approach can be productive.

To see something is an achievement. That is to say, the claim to have performed correctly can be tested. Indeed, we can investigate how that task of visual recognition was successfully carried out. We can try to infer the information-transforming (cognitive) processes mediating the performance by varying what is visible and observing changes in response (i.e., doing psychophysics); this is an example of psychological science.

The physical "engineering" of these processes of seeing can also be studied by varying the optical input, but this time observing what is projected onto the retina and activity in the central nervous system (CNS), from the rods and cones to V1 and beyond. Considerable progress has been made in relating cellular neurophysiology to the psychophysics of elementary features of the visible world. It is not so easy to get psychophysical evidence that distinguishes between a cognitive process being in consciousness and transiently out of consciousness (Booth & Freeman 1993), although it is clear that some visual information processing never enters consciousness. When we cannot specify a mental process as conscious, there cannot be a theory of the neural basis of that process. Lehar's complaint that neuroscience fails to explain visual consciousness is vacuous.

Furthermore, what we know to be the case through use of our senses is a very different kettle of fish from the contents of consciousness, in the sense of how things seem to us while we discount our beliefs about how they actually are. By definition, how things seem cannot be checked against how things are. So the systematisation of expressions of subjective experience is an art form. Lehar's diagrams, his field equations, and his verbal exposition are sophisticated elaborations of the sort of thing that I draw when I wake up and try to sketch the visual imagery that I was experiencing as I woke. His and my graphic, algebraic, and verbal efforts cannot be wrong or right; they merely express how it appeared to be.

Lehar says that his visual experience is holistic. I can empathise with that impression. Yet I also have visual experiences that are not holistic. I bet that he does too but chooses to ignore them. Any artist may do that, on the grounds that it would spoil the picture or detract from the story. However, that is aesthetics, not science.

I am not being positivistic. On the contrary, it is Lehar who commits the empiricists' and rationalists' epistemological fallacy of trying to build public knowledge on the basis of impressions or ideas that seem indubitable because they are private and so cannot be wrong – but then neither can they be right. Lehar writes: "These phenomena are so immediately manifest in the subjective experience of perception that they need hardly be tested psychophysically" (target article, sect. 10, para. 2). In words of one or two syllables: "What appears seeming to seem in seeing is so clearly clear that there is no need to test it against success at seeing."

Lehar's paper is built on equivocation in use of the word "perception" between the objective achievement and subjective experience. (The word "conscious" in his title is redundant: experiencing subjectively is the same as being conscious.) Like most philosophers, mathematicians, and physicists who expatiate on consciousness, he shows no sign of having considered what was shown, and how it was shown, by any psychological experiment on the perceiver's achievement in a visual task. He also ignores the philosophical advances following the later Wittgenstein's debunking, 60 years ago, of the pervasive fallacy of supposing that when a patch that is red (in the world that we all live in) is seen as red, this is a "seeming" in another world (Lyons 1983). Worse, because these appearances, subjective experiences, conscious qualia, or whatever, are part of each of us, Lehar (like many) locates them in our heads, or as neurocomputations if we are foolish enough to look for consciousness among the brain cells (Booth 1978). This is all a big mistake about the grammar of the verb "to seem." When we are viewing something but have reason to doubt that we perceive it correctly, then we may retreat to a claim that it seems to be so. We are not looking at a world inside our minds; we are having problems in seeing the colour of the patch out there.

The grammar of "seeming as though" or "seeing as" also shows what the subjective experience is isomorphic to. The syntax of "as" is the figure of speech known as simile. Subjective visual experience is holistic, at least at times, because the world in which we operate is "holistic" in its optics; black holes are pretty uncommon in everyday life. Lehar actually says this in section 3, although he has hidden the point from himself by a tangle of the conceptual mistakes that Wittgenstein (1953) cut through. "The perceptual experience of a triangle cannot be reduced to just three phenomenal values but is observed as a fully reified triangular structure that spans a specific portion of perceived space" (sect. 3, para. 2). Delete the reference to a contrary and all the redundancies and we get: "The perceptual experience of a triangle . . . is . . . as [*sic*] . . . triangular. . . ."

Furthermore, a triangle is not a triangle in any world unless it "emerges" "whole," "real," and "invariant." If a Gestalt is taken to be a subjective experience (rather than a perceptual performance), then it is consciousness simply of "seeing the world as it is."

There is no space in this commentary to dissect out the multitudinous errors built on this fundamental misorientation. Suffice it to deal with the absurdity of the target article's Figure 2. Lehar shows phenomenological slapdash if not downright dishonesty. You know and I know that he has never looked one way down a road at the very same moment as looking the other way. So it is rank self-deception to write (sect. 6.3, para. 1) that "the two sides of the road must in some sense be [subjectively] perceived as being bowed" as in the diagram. His Bubble bursts.

Double, double, toil and trouble – fire burn, and theory bubble![1]

Birgitta Dresp

Etude des Mécanismes Cognitifs et Dynamique du Langage, UMR 5596 CNRS–Université Lumière, Lyon 2, France. **birgitta.dresp@univ-lyon2.fr dresp@convergence.u-strasbg.fr**

Abstract: Lehar's Gestalt Bubble model introduces a computational approach to holistic aspects of three-dimensional scene perception. The model as such has merit because it manages to translate certain Gestalt principles of perceptual organization into formal codes or algorithms. The mistake made in this target article is to present the model within the theoretical framework of the question of consciousness. As a scientific approach to the problem of consciousness, the Gestalt Bubble fails for several reasons. This commentary addresses three of these: (1) the terminology surrounding the concept of consciousness is not rigorously defined; (2) it is not made evident that three-dimensional scene perception requires consciousness at all; and (3) it is not clearly explained by which mechanism(s) the "picture-in-the-head," supposedly represented in the brain, would be made available to different levels of awareness or consciousness.

In this target article we are told that "the most serious indictment of contemporary neurophysiological theories is that they offer no hint of an explanation for the subjective experience of visual consciousness" (sect. 1, para. 2). Lehar attacks "good old" Neuron Doctrine by stating that as a theoretical approach to visual perception, it has reached a dead end because he (Lehar) finds it "hard to imagine how . . . an assembly of independent processors [neurons] could account for the holistic emergent properties of perception identified by Gestalt theory" (sect. 1, para. 3). He then proposes his own doctrine, the Gestalt Bubble model. The Gestalt Bubble is presented as a computational approach to the perceptual representation of three-dimensional visual space using a volumetric matrix of dynamic elements, each of which can exist in one of several states: transparent for the representation of void space, opaque coplanar for the representation of smooth surfaces, opaque orthogonal for the representation of corners, and opaque occlusion for the representation of surface edges. The supposed transformation of the physical world outside by a perceptual process taking place inside the brain is defined as the turning on of the appropriate pattern of elements in the volumetric matrix of the model in response to visual input. The Gestalt Bubble thereby replicates the three-dimensionality of visual objects as they are experienced in the subjective percept. The principal merit of this model resides in the fact that it translates some major Gestalt laws of visual perception such as emergence, reification, multistability, and invariance into computational codes.

What the author fails to make clear in his target article is the supposed link between his Gestalt Bubble model and general theories of consciousness. All he does here is demonstrate that modern computer technology produces algorithms that allow us to translate the laws of perceptual organization formulated in Gestalt theory into formal codes within the framework of a computational model. What the model has to do with consciousness, however, remains totally unclear. Neither the fact that we are able to consciously experience and describe three-dimensional shapes as entities and wholes, nor the fact that we can find laws or codes describing how these emerge perceptually, implies or proves that consciousness is necessary to see and move around in three-dimensional space. In addition, although Lehar seems to imply that his Gestalt Bubble provides a ready model of what he refers to as visual consciousness, he fails to provide clear definitions of what we are supposed to understand by visual consciousness, phenomenal awareness, subjective perceptual experience, or consciousness in general. In the title of the target article, he uses the term "subjective conscious experience." Does this suggest that there should be an objective conscious experience as well? Moreover, the author readily assumes the existence of a "visual consciousness" as a particular form of consciousness. This assumption needs to be justified. How would a visual consciousness operate in comparison to an auditory, tactile, or olfactory consciousness, for example? In fact, by using ambiguous terminology in his text (terminological danglers?), switching readily from one level of explanation to another, the author fails to convince his readers that he knows what he is talking about when he discusses the question of consciousness.

Moreover, the fundamental difference between Lehar's "picture-in-the-head" model and the concept of isomorphism from classic Gestalt theory is not discussed in a satisfactory manner. After a lengthy introduction that confronts the reader with odds and ends of numerous general theories of mind and consciousness, the author all of a sudden pops up his own version of the Gestalt hypothesis of isomorphism by suggesting that we see the outside world as we do because that is and has to be the way the world is represented in the brain. This "picture-in-the-head" view goes far beyond the classic Gestalt concept of isomorphism because it assumes not only a functional but also a structural correspondence between the visual percept and its brain representation. It is introduced here as the only rightful answer to Koffka's question "Why do we see things as we do?"; the original Gestalt viewpoints (e.g., Kohler 1961; Metzger 1936; von Ehrenfels 1890; among others) on isomorphism are not discussed.

Interestingly, the author seems to have overlooked that his "picture-in-the-head" hypothesis (structural isomorphism) stands or falls on the validity of the assumption that one of the key principles formulated by Gestalt theory, that of the common fate of parts (*Ganzbestimmtheit der Teile;* Metzger 1936), reflects the result of a neurophysiological mechanism. In the early sixties, some psychophysicists questioned the neurophysiological validity of precisely this principle of perceptual organization. Pritchard (1961) presented figures as stabilized images on the retina and showed that the constituent elements of these figures disappeared from phenomenal awareness one by one – not all at once, as the principle of common fate of parts would predict if it reflected the result of a neurosensory mechanism (see also Pritchard et al. 1961). In any case, even if the "picture-in-the-head" view could be proven right, Lehar would still have to come up with an explanation of the mechanism(s) by which the picture in the head is made available to consciousness. Also, a rigorous distinction between "awareness," such as awareness of the emergent properties of a visual object at a given moment, for example, and "consciousness," such as the consciousness of being aware of the emergent properties of a visual object and its significance within a general context, for example, would then have to be made.

Lehar writes that it is of central importance for psychology to address what "all that neural wetware" is supposed to do and to determine which of the competing hypotheses presented in the introduction of his target article "reflects the truth." Who said that science has to bother with metaphors such as "truth"? As far as I understand it, science is all about facts and measures collected within a specific context of boring constraints, usually called "conditions," and therefore inevitably requires a diversity of methods and hypotheses. The concept of "truth" does not appear to be of much use here. Are we not often enough reminded to take care not to get trapped by the metaphors we use to construct hypotheses and explanations? The overwhelming "*Unsumme*" (as defined by Metzger 1936) of bits and pieces of philosophy and phenomenological "brain teasers" we are confronted with in this target article somehow shows how easily we can end up like the Sorcerer's Apprentice in Goethe's poem, who tries all sorts of curses and invokes all sorts of spirits, but is finally unable to take control.

In conclusion, whether theories based on or derived from the Neuron Doctrine will ultimately fail to provide a satisfactory approach to the question of consciousness, remains to be seen. The Gestalt Bubble model, as a scientific approach to consciousness, can be filed DOA (Dead on Arrival).

NOTE

1. After Shakespeare, *Macbeth.*

Just bubbles?

Włodzisław Duch

School of Computer Engineering, Nanyang Technological University, Singapore; Department of Informatics, Nicholaus Copernicus University, Torun, Poland. **wduch@phys.uni.torun.pl**
http://www.phys.uni.torun.pl/~duch

Abstract: Lehar misrepresents the Neuron Doctrine and indirect realism. His conclusions on consciousness are unjustified. The Bubble Gestalt perceptual modeling disconnected from neuroscience has no explanatory power.

1. Perception has not evolved for our enjoyment; it serves action, exploration of the world (see O'Regan & Nöe 2001). Although the richness of visual perception may partially be an illusion, sensory data should elicit brain states that reflect important features of perceptual organization. Such functional representation would be very useful, facilitating information retrieval from visual and auditory cortex, stored in attractor neural networks after termination of direct sensory inputs (Amit 1994). Persistent brain activity may be responsible for visual imagery, filling in, illusory contours, and other such phenomena. This internal representation, being a physical state of the brain, is focused and interpreted by other brain areas, gating it to the working memory and facilitating conscious perception. It is constructed from sparse information obtained from eye fixations between saccades (as is evident in the change blindness experiments; O'Regan & Nöe 2001) and hence may not be as faithful and rich as it seems. Because for many people endowed with visual imagination (individual variance seems to be quite large in this respect) visual experiences are rich and vivid, filling in of missing information must be strong.

2. Construction of the inner perspective is a difficult task. Lehar does not even attempt to enumerate the dimensions required for perceptual modeling that could replace (or at least complement) neural modeling. I have argued myself (Duch 1997) that an intermediate level of cognitive modeling should be useful. It should represent mental events in a way that is closer to our inner perspective, acceptable to psychologists, but should also facilitate reduction, at least in principle, to the neural level. Complex neural systems reveal emergent processes (responsible, as Lehar has noticed, for Gestalt phenomena) requiring a higher level of description characterized by new laws and phenomena. The usual approximation to neural activity misses the perceptual level by going from states of recurrent networks (such as Grossberg's adaptive resonant states; Grossberg 1995) to states of finite automata (cf. Parks et al. 1998 for neural models in psychiatry). A shortcut from neuroscience via neural networks to behavior is satisfactory only to behaviorists. Mind states and mental events may emerge as "a shadow of neurodynamics" in psychological or perceptual spaces (Duch 1997). This is in accord with the ideas of Shepard (1987; 1994), who believes that universal laws of psychology may be found in appropriate spaces. Psychological spaces are spanned by subjective dimensions (such as color, shape, and motion), and one may use them to explain subjective perception and to talk about mental events implemented at the neurodynamical level. Therefore, I sympathize with Lehar's goal, although details of his proposal are not satisfactory.

3. Trivializing the "Neuron Doctrine," Lehar writes about neural networks as the "quasi-independent processors," and "an assembly of independent processors" (target article, sect. 1, para. 3). The whole essence of neural networks is in the interaction of their elements, cooperative computational abilities that facilitate their holistic emergent properties. Recurrent neural networks are certainly not "the atomistic feed-forward model of neurocomputation" (target article, Abstract; cf. Parks et al. 1998). The Neuron Doctrine paradigm has been completely misinterpreted in the target article.

4. The arguments evoked against indirect realism are strange to say the least. Lehar mixes mental and physical levels freely, writing statements like "the world that appears to be external to our head is actually inside our head" and "beyond those perceived surfaces is the inner surface of your true physical skull encompassing all that you perceive" (sect. 2.2, para. 1). How can the physical skull encompass the nonphysical, inner world? "The world inside the head" is a metaphor, and it does not make much sense to invert it, unless one believes that there is some kind of physical world squeezed inside the skull.

Indirect realism claims that we perceive and comment upon the states of our own brain. These states reflect properties of the environment, but interpretation of the spatial structure of the states of the visual system has nothing to do with their physical location. There is nothing strange about it, as there is nothing strange about transmission of the voice and images via wires and radio waves. The spatial world inside the head is there in the same sense as a panoramic image in the integrated circuit of a computer graphic chip. Subjective reversal of a multistable percept follows the change of neural dynamics. It has to be experienced vividly as an inversion of a perceptual data structure, because visual experiences are a reflection of neural dynamics – how else could changes of visual cortex states be experienced?

5. It is certainly not clear "that the most fundamental principles of neural computation and representation remain to be discovered" (target article, sect. 2.4, para. 3). Churchland (1984) had already argued against it 20 years ago, and since that time computational neuroscience has made a lot of progress. It may very well be that Hebbian learning is the only fundamental principle that is needed and that sufficiently complex models of the brain will be able to simulate its emergent functions.

6. It is quite probable that "our own conscious qualia evolved from those of our animal ancestors" (sect. 6.5, para. 3). But certainly the "conclusion" (sect. 6.5, para. 6) "that all matter and energy have some kind of primal protoconsciousness" is not inescapable. In fact, I regularly lose my consciousness in sleep, and anesthetics and damage to the reticular formation lead to coma, obliterating consciousness. Complex organization of matter is not sufficient for consciousness. Instead of looking for conditions necessary for manifestation of consciousness – a fruitful way is to use

https://doi.org/10.1017/S0140525X03410091

here a contrastive approach between perception and reception (Taylor 1999) – Lehar goes down the beaten track of thinking about consciousness as some kind of a substance that is present in all matter, although sometimes in watered-down form. The conclusion of this line of reasoning is absurd: protoconsciousness of soap bubbles.

Of course, because the concept of consciousness is not defined, one may try to extend it to all matter, but talking about stomachs being "conscious" leaves no semantic overlap with the word "conscious" applied to a baby, or to a cat. If consciousness is a function and plays a functional role, as Lehar seems to believe ("It seems that conscious experience has a direct functional role" – sect. 6.5, para. 10), the inescapable conclusion is rather that not all brains are equal. Language is unique to humans, and even though one can extend the concept of language to some more primitive forms of communication, interaction between internal organs of the body or messages passing between components of a computer system is not the same "language" as natural languages. The difference between a "field" in agriculture and "field" in physics is comparable to the difference between animal "consciousness" and "consciousness" of a soap bubble due to the physical forces that determine its shape. We should not be deceived by words.

7. It remains to be seen if the main contribution of the target article, the Gestalt Bubble model, will be useful for understanding or even for a description of perception. The goal of science is not modeling per se but rather explaining and understanding phenomena. Modeling perception should not become an exercise in computer graphics, creating volumetric representations of space and objects. Bubbles of neural activity, as presented by Taylor (1999), have real explanatory power and are amenable to empirical tests. The perceptual modeling proposed by Lehar promises a new language to describe high-level visual perception. Any language that is useful in design and analysis of experiments must reflect more basic neural processes. Nothing of that sort has been demonstrated so far, and it is doubtful that the Gestalt Bubble model can explain observations that have not been hidden in its premises.

Empirical constraints for perceptual modeling

Charles R. Fox

Whitely Psychology Laboratories, Franklin and Marshall College, Lancaster, PA 17604-3003. **charlie.fox@fandm.edu** **http://www.fandm.edu/departments/psych_new/faculty/fox/**

Abstract: This new heuristic model of perceptual analysis raises interesting issues but in the end falls short. Its arguments are more in the Cartesian than Gestalt tradition. Much of the argument is based on setting up theoretical straw men and ignores well known perceptual and brain science. Arguments are reviewed in light of known physiology and traditional Gestalt theory.

Steven Lehar's article purports to present a new model of perception based on Gestalt principles. Lehar raises some interesting issues but in the end falls short of his claims. His heuristic model is more Cartesian than Gestalt and much of his argument is based on setting up straw men. He ignores much of what is known in perceptual and brain science. I will confine myself to these issues, although there are others.

Lehar maintains the Cartesian mind-body distinction and assumes internal representation as a requirement. He also ignores the distinction between conscious perception as active construction and the perception/action continuums implied by physiology and direct perception data. Lehar recycles the Cartesian machinelike body now inhabited by the "ghosts" of mental representations and computations. This dualism is at odds with traditional Gestalt theory (Köhler 1969).

The target article ignores the contemporary distinction between (1) perceptual mechanisms that subserve action; and (2) the cognitive mechanism of recall and analysis; instead, it suggests the latter as the sole perceptual mechanism. This emphasis stems from Lehar's belief that "introspection is as valid a method of investigation as is neurophysiology" (sect. 2.3, last para.). This is not the position of traditional Gestalt theory, which states that "a satisfactory functional interpretation of perception can be given only in terms of biological theory" and warns that "The value of biological theories in psychology is not generally recognized." Gestalt psychology adopted the program of building bridges between psychological rules and the activities of the central nervous system (Köhler 1940; 1947; 1961). Köhler recognized this task as "beyond present *technical* possibilities." These purely technical limits are being overcome today, yet the target article ignores a large body of empirical physiological evidence, some of which is presented below (see also Milner & Goodale 1995 and Gallese et. al. 1999 for summary of some areas). Although we should not limit our theories to physiology, theory must account for known physiology. The target model does not. To take a specific example, the model ignores the important role of eye movements even though they were of concern to the early Gestalt theorists (Koffka 1935) and are a critical part of contemporary perceptual theory (Ebenholtz 2001). More generally, there is ubiquitous evidence, collected over many decades, for the important role of physiological systems in perception. Simply consider the differential perceptions resulting from anatomical and physiological states of sensory end organs. Visual perception in the myopic, dark-adapted, or macular-degenerated eye is more influenced by anatomy and physiology than by computations on a mental image.

Lehar emphasizes computational neuroscience at the expense of known physiology despite his assertion that "most fundamental principles of neural computation and representation remain to be discovered" (sect. 2.4, para. 3). This leads to oversimplification to the point of error. For example, he dismisses direct perception because "No plausible mechanism has ever been identified neurophysiologically which exhibits this incredible property" (sect. 2.2, para. 3) and "all that computational wetware" (sect. 2.1, para. 2) must serve some "purpose" (i.e., "produce an internal image of the world"; sect. 2.1). Yet there is growing physiological evidence to the contrary. As I have discussed elsewhere (Fox 1999), area MST in monkeys (similar to area V5 in humans) shows cells that are responsive to three-dimensional motion information that is characteristic of the type of flow field emphasized by direct perception theory (Duffy & Wurtz 1995; 1997a; 1997b). More recently, direct perception theorists have examined the relation of neural information systems to Tau, a property of environmental optics (Grealy 2002; Lee et al. 2002). Hence, contemporary physiology supports an emerging model suggestive of an environmentally adapted physiology rather than the metaphor of representational/computational "wetware."

Lehar further misrepresents direct perception theory as describing perception "as if perceptual processing occurs somehow out in the world itself rather than as a computation in the brain" (sect. 2.1, para. 1). Using the term "perceptual processing" or "computation" is a serious misrepresentation of direct perception (Gibson 1966; 1979), regardless of where one attributes it. Gibson contends that the perceptual system is sensitive to "affordances" that are naturally occurring and require no processing but rather are directly perceived. The exact characteristics of affordances are disputed, but a recent paper (Chemero 2003) provides a critical analysis and comprehensive definition of the concept of affordances and makes it very clear that affordances are perceived relations that are dynamic but neither computed nor components of computations. This is consistent with the physiology described above.

Gestalt psychology is also misrepresented as a representational/computational approach. I content that a key – perhaps *the* key – insight of Gestalt theory is that adequate knowledge of wholes, such as objects, comes from observing wholes. Such understanding does not come from a "humpty-dumpty" approach that tries to put the object "back together again" through computation. The target model is reductionist/empiricist and, as such, is contrary to Gestalt theory (Koffka 1935; Köhler 1947). The relevant properties of things are not computational properties superimposed on the object system, but rather, the intrinsic relational properties within the object and between the object and the perceiver/actor (Köhler 1947). For example, Köhler certainly did not suggest that perception is a mental computation when he wrote: "While climbing once in the Alps I beheld . . . a big dark cloud . . . nothing could be more sinister and more threatening. . . . the menace was certainly in the cloud." The menace stems not from computations on mental images but from physiological sensitivity to relations among environmental physical energies, and between these relations and the state system of the observer/actor. I suggest a dynamic, person-environmental mechanism rather than internal representation and computations. This is consistent with the Gestalt statement: "rules in which we formulate (functional, psychological) relationships imply occurrences of certain functions in a realm that is surely not the phenomenal realm" (Köhler 1940).

A final, critical point concerns isomorphism: Isomorphic relations are ubiquitous, so one needs to be specific. Gestalt "Psychophysical Isomorphism" is a hypothesis that rejects Cartesian dualism and is informed by physiology (Köhler 1969). Lehar, using a digital computer metaphor, suggests a point-to-point isomorphism between the internal image and external objects/space. However, this is not supported by physiology. Cells in the supplementary eye field of the monkey show firing patterns (Olson & Gettner 1995) that do not encode visual space in any one-to-one manner. Rather, they incorporate higher dimensions of information such as attention or purpose (Fox 1999). Hence, even if we accept isomorphic, internal representations, there is neurophysiologic evidence that such representations are more complex than suggested in Lehar's model.

The target model does not accomplish its ambitious goals of presenting a modern Gestalt perceptual model. A more fruitful heuristic for understanding perception is a physiology that has evolved a sensitivity to meaningful environmental relational information, or, as suggested by Clark (1998), one that represents action-oriented systems.

ACKNOWLEDGMENT
The preparation of this commentary was partially supported by a grant to the author from the Franklin and Marshall College, Office of the Provost.

Linking visual cortex to visual perception: An alternative to the Gestalt Bubble

Stephen Grossberg

Department of Cognitive and Neural Systems, Boston University, Boston, MA 02215. **steve@bu.edu** **http://www.cns.bu.edu/Profiles/Grossberg**

Abstract: Lehar's lively discussion builds on a critique of neural models of vision that is incorrect in its general and specific claims. He espouses a Gestalt perceptual approach rather than one consistent with the "objective neurophysiological state of the visual system" (target article, Abstract). Contemporary vision models realize his perceptual goals and also quantitatively explain neurophysiological and anatomical data.

Lehar describes a "serious crisis," "an impasse," and a "theoretical dead end" (target article, sect. 1, para. 1) in contemporary models of vision and advances as a possible alternative his Gestalt Bubble approach, "which is unlike any algorithm devised by man" (Abstract). He also claims that "Gestalt aspects of perception have been largely ignored" (sect. 1, para. 3) by neural models of vision, and then goes on to describe presumed dichotomies between equally desperate attempts to understand how the brain sees. Lehar particularly comments about modeling work by my colleagues and myself, noting that "the most serious limitation of Grossberg's approach . . . is that, curiously, Grossberg and his colleagues did not extend their logic to . . . three-dimensional spatial perception [and] . . . no longer advocated explicit spatial filling-in" (sect. 3, para. 5). He also says it is "impossible for Grossberg's model to represent transparency" (sect. 3, para. 5). These general and specific claims unfortunately do not accurately represent the published literature about neural vision models. Lehar seems motivated to trash neural vision models because his own model makes no contact with neurophysiological and anatomical data about vision.

In reality, there is an emerging neural theory of three-dimensional vision and figure-ground perception called the FACADE theory, for the multiplexed Form-And-Color-And-DEpth representations that the theory attempts to explain (Grossberg 1987; 1994; 1997). Lehar refers to my 1994 article in summarizing the deficiencies of our models. However, this article explains many three-dimensional figure-ground, grouping, and filling-in percepts, including transparency, and uses an explicit surface filling-in process. Later work from our group has developed these qualitative proposals into quantitative simulations of many three-dimensional percepts, including three-dimensional percepts of da Vinci stereopsis, figure-ground separation, texture segregation, brightness perception, and transparency (Grossberg & Kelly 1999; Grossberg & McLoughlin 1997; Grossberg & Pessoa 1998, Kelly & Grossberg 2000; McLoughlin & Grossberg 1998).

These studies laid the foundation for a breakthrough in understanding how some of these processes are organized within identified laminar circuits of cortical areas V1 and V2, notably processes of cortical development, learning, attention, and grouping, including Gestalt grouping properties (Grossberg 1999a; Grossberg & Raizada 2000; Grossberg & Seitz 2003; Grossberg & Williamson 2001; Grossberg et al. 1997; Raizada & Grossberg 2001; 2003; Ross et al. 2000).

This LAMINART model has been joined with the FACADE model to develop a three-dimensional LAMINART model that quantitatively simulates many perceptual data about stereopsis and three-dimensional planar surface perception, and functionally explains anatomical and neurophysiological cell properties in cortical layers 1, 2/3A, 3B, 4, 5, and 6 of areas V1 and V2 (Grossberg & Howe 2003; Howe & Grossberg 2001), using three-dimensional figure-ground and filling-in concepts to do so. More recently, the three-dimensional LAMINART model has been generalized to explain how three-dimensional percepts of slanted and curved surfaces and of two-dimensional images are formed, and to clarify how three-dimensional grouping and filling-in can occur over multiple depths (Grossberg & Swaminathan 2003; Swaminathan & Grossberg 2001). This work includes explanations of how identified cortical cells in cortical areas V1 and V2 develop to enable these representations to form, how three-dimensional Necker cube representations rival bi-stably through time, how slant aftereffects occur, and how three-dimensional neon color spreading of curved surfaces occurs even at depths that contain no explicit bottom-up inputs. All these studies are consistent with the grouping interpolation properties that Kellman et al. (1996) have reported (p. 51), and with the three-dimensional grouping properties summarized in Lehar's Figure 16, which he seems to think cannot yet be neurally explained.

These modeling articles show that many of the perceptual goals of Lehar's Gestalt Bubble model are well handled by neural models that also provide a detailed account of how the visual cortex generates these perceptual effects. In summary, we do not need analogies like the soap bubble (sect. 8.2), or rod-and-rail (sect. 8 and Fig. 6), or different local states to represent opaque or transparent surface properties, as Lehar proposes. The brain has discovered a much more interesting solution to these problems, which links its ability to develop and learn from the world with its ability to see it.

Lehar makes many other claims that are not supportable by present theoretical knowledge. He claims that "we cannot imagine

https://doi.org/10.1017/S0140525X03410091

how contemporary concepts of neurocomputation . . . can account for the properties of perception as observed in visual consciousness [including] hallucinations" (sect. 2.4, para. 3). Actually, current neural models offer an explicit account of schizophrenic hallucinations (Grossberg 2000) as manifestations of a breakdown in the normal processes of learning, expectation, attention, and consciousness (Grossberg 1999b).

Contrary to Lehar's claims in section 8.7, recent neural models clarify how the brain learns spatial representations of azimuth, elevation, and vergence (see Lehar, Fig. 14) for purposes of, say, eye and arm movement control (Greve et al. 1993; Guenther et al. 1994). Lehar defends "the adaptive value of a neural representation of the external world that could break free of the tissue of the sensory or cortical surface" (sect. 8.8). Instead, *What* stream representations of visual percepts should be distinguished from *Where* stream representations of spatial location, a distinction made manifest by various clinical patients.

Lehar reduces neural models of vision to capacities of computers to include navigation as another area where models cannot penetrate (see sect. 6.1 and sect. 9). Actually, neural models quantitatively simulate the recorded dynamics of MST cortical cells and the psychophysical reports of navigating humans (Grossberg et al. 1999), contradicting Lehar's claim that "the picture of visual processing revealed by the phenomenological approach is radically different from the picture revealed by neurophysiological studies" (sect. 9, para. 1). In fact, a few known properties of cortical neurons, when interacting together, can generate emergent properties of human navigation.

Lehar ends by saying that "curiously, these most obvious properties of perception have been systematically ignored by neural modelers" (sect. 10, penultimate para.). Curiously, Lehar has not kept up with the modeling literature that he incorrectly characterizes and criticizes.

Steven Lehar's Gestalt Bubble model of visual experience: The embodied percipient, emergent holism, and the ultimate question of consciousness

Keith Gunderson

Department of Philosophy, University of Minnesota, College of Liberal Arts, Minneapolis, MN 55455-0310. **gunde002@umn.edu**

Abstract: Aspects of an example of simulated shared subjectivity can be used both to support Steven Lehar's remarks on embodied percipients and to triangulate in a novel way the so-called "hard problem" of consciousness which Lehar wishes to "sidestep," but which, given his other contentions regarding emergent holism, raises questions about whether he has been able or willing to do so.

Steven Lehar's Gestalt Bubble model (GBM) is said to emphasize the often ignored fact "that our percept of the world includes a percept of our own body within that world, . . . and it remains at the center of perceived space even as we move about in the external world" (sect. 6.4). I offer here a friendly, if folksy, example of a *simulation* of shared first-person subjectivity designed to reinforce Lehar's brief but interesting claims concerning the prominence of the embodied percipient in visual perception. This example leads to other questions regarding his analysis. I have labeled the example elsewhere, and with variations, the Cinematic Solution to the Other Minds Problem, and invoked it earlier against B. F. Skinner's view of subjective privacy and scientific inquiry, also objected to by Lehar for his own reasons (Gunderson 1971; 1984).

Suppose a film director wishes to treat us to the subjective perceptual experiences of another person, say Batman, as he gazes on the traffic far below from some window perch. How is this best done? Not, to be sure, by simply showing us the whole scene: the superhero perched on the ledge with the traffic moving by on the street below. This would not be anything like being privy to Batman's subjective perceptual experience. It would only amount to our own visual experience coming to include Batman. Instead, what is characteristically done is that Batman's filmed body (or at least the better part of it) is somehow (gradually or suddenly) subtracted from the screen in such a manner that we become insinuated into roughly whatever space and orientation Batman's body occupies and are thereby made party to the visual field (sense of height, traffic passing below, etc.), which we can assume would be Batman's from that perspective. We cannot, of course, literally occupy (even cinematically) exactly the same space that Batman does – a prerequisite to having his visual experience – but the tricks of the art permit us to enjoy a simulation of such an occupancy. It is the sleight-of-camera with respect to our seemingly ubiquitous embodied presence in visual perception that carries with it tactics for conjuring a sense of the usual "subjectivity barrier" between us and another percipient being breached. And here it occurs in a florid phenomenological manner, obviously different from the "relational information" that can cross that barrier, as described by Lehar (sect. 5.1). Notice too, that a "preset" feature of the whole typical movie experience involves the darkened theater and no focused sense of our own body being either present in the audience or included in the screen action. The effect is that where we are not assuming specifically Batman's perspective, we are assuming one belonging to no one in particular, or rather one "belonging" to anyone in the vicinity, as it were.

So the possibility of the cinematic simulation of shared subjectivity seems to presuppose the inclusion of an embodied percipient in our visual perceptions, along lines suggested by Lehar. But the apparent friendliness of the example has a complicated provocative side as well. For if what it takes to create the illusion is the clever collapsing of our perspective (or someone else's) into another's, then the epistemic-ontic primacy of the first-person point of view becomes obvious, and the "hard problem" of consciousness can be rephrased with respect to it this way: There is no analogous thought experiment that would render subjectivity or a point of view (one's own or another's) as being somehow manifest in any set of neurophysiological processes to begin with, such that another consciousness might appear as somehow insinuated into it. But there should be, if consciousness is to be modeled (displayed, illustrated) within any third-person physicalistic conceptualization. This rather flat and crude-sounding point is not, I think, irrelevant or naïvely realistic. In a nutshell, that there can be no cinematic-type simulation of a solution to the mind-body problem parallel to another mind's, can be seen to stem from our inability to cling to our sense of experiencing a point of view while being in some neurophysiological locus (however this is represented).

For Lehar, the salient residual problem(s) is this: Although the contents of all our subjective visual experience for the GBM are subsumed under the subjective, we lack *any* vivid demonstration of how *having* a first-person point of view in itself, which is a prerequisite to there being any such phenomenal contents, lies within that experience. Simply specifying *underlying* neurophysiological conditions for consciousness takes us nowhere we have not already unsatisfactorily been. That there is, and how there is, any locus at all for our perceptions remains unexplained within any micro or macro frame of reference. We think, of course, that the *locus of our locus* of perceptions lies in some way within the embodied. But to be apprised of all this does not thereby help us to see how any subjective perspective occurs in the first place, or why it is uniquely ours! (See Nagel 1965.) The problem of explaining it arises independently of whatever type of metaphysical substance the perceiver is believed to be embodied in, even as part of a *panpsychic* or *panexperientialist* scheme such as Chalmers' (as in sect. 6.5). And it can be reiterated with respect to any type of substance of any kind of complexity, as far we can tell.

Now, Lehar wishes to "sidestep" these latter matters by casting the GBM *wholly within the subjective.* Our perceived worlds –

our pattern-recognizing activities, including, of course, our total physical natures – will then supposedly lie within the range of what his subjectively rendered model is a model of. But I do not see how this really matters, even when naïve realism such as Skinner's is deleted from the picture for the (laudable) reasons Lehar provides. One might, of course, wish, out of other considerations, simply to set the mind-body problem aside and concentrate on refining taxonomic characterizations within phenomenal experience. (Nagel 1974 is cited as having suggested something like this.)

But more puzzling to me is why Lehar's concluding remarks about Koffka's and Köhler's views on emergence (sect. 7.1), which Lehar finds more satisfying than Davidson's anomalous monism, are not a way of directly addressing "the hard problem." The pivotal demystifying image in the "bottom-up" aspect of Lehar's summary of the mind-body relationship is that of perception characterized along Gestalt lines as being related to neurophysiological processes, in the way that a soap bubble holistically emerges from "a multitude of tiny forces acting together simultaneously" to produce a final perceptual state by way of a process that cannot be reduced to simple laws (sect. 7.1). But whatever other, if any, purposes this no doubt interesting image may serve, the relationship between bubble and tiny forces is not in any discernible way similar to whatever the connection between subjective states of conscious perceptual awareness and neurophysiological states is like. Both bubbles and tiny forces are happily in the world, as it were, whether as macro-bubblistic ones, or micro-force-istic ones, or as something like the pop-out dog example (sect. 7.1). These all involve one set of "out there" aspects being related to other "out there" aspects, whether within the subjectivized purview of the GBM or some other one. The bugbear of consciousness still seems to turn on the point that first-person conscious perspectival states cannot yet be even imagined as either macro or micro anythings to begin with, much less as popping up from micro ones.

Backdrop, flat, and prop: The stage for active perceptual inquiry

Julian Hochberg

Columbia University, 170 West End Avenue, New York, NY 10023.
Hochberg@columbia.edu

Abstract: Lehar's revival of phenomenology and his all-encompassing Gestalt Bubble model are ambitious and stimulating. I offer an illustrated caution about phenomenology, a more fractured alternative to his Bubble model, and two lines of phenomena that may disqualify his isomorphism. I think a perceptual-inquiry model can contend.

Steven Lehar's ambitious Bubble metaphor is highly stimulating, assuming a unified phenomenal visual world that explains and predicts our perceptual experience. Herewith are a cautionary reminder about phenomenology as such, an alternative to Lehar's specific enclosing Bubble model, and two lines of phenomena that Lehar ignores but that are difficult to reconcile with the particular isomorphism he espouses.

Phenomenology should indeed guide psychophysics and neurophysiology. But phenomenology is certainly not incontestable. For example, Lehar cites the CIE chromaticity diagram as a description of phenomenological color space. The Helmholtzean dogma – that the experience of yellow consists of red plus green experiences – lurked within mainstream sensory physiology until after World War II (and was often attributed to the CIE). Then, following Hering instead, Hurvich and Jameson's (1957) phenomenologically guided opponency-oriented psychophysics and model explained to neurophysiologists what their microelectrodes later revealed, thereby changing our view of neurophysiology and liberating our relevant phenomenology. (In fact, Jameson & Hurvich showed later [1967] that the CIE is no phenomenological

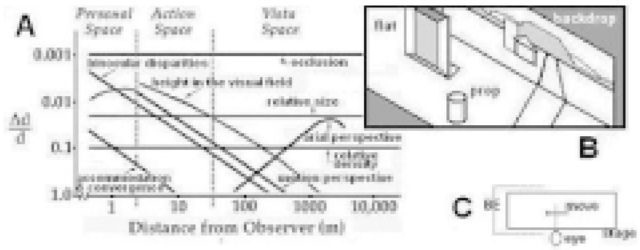

Figure 1 (Hochberg). Onstage and backdrop scenery. **A.** The strength of the major depth cues with egocentric distance, adapted from Cutting and Vishton 1995 (with permission). To the eye as actor, the backdrop usually lies between 10 and 50 feet upstage. **B.** The experienced stage in which visual inquiry proceeds. The viewer's normal actions provide no distance information beyond the plane labeled "backdrop" and they can readily generate and therefore incorporate information about the downstage prop. The curves of Figure 1A account for but are not salient in the experience of B. **C.** Attention extends the stage. When the inquiring eye visits a scene, its boundaries are remembered as further out than they were (see Intraub 1997); this is not merely memory, because such Boundary Extension (BE) is a function of where the viewer plans to look (Intraub et al. 2001).

summary – two very different colors come out at the same point on the graph.) Phenomenology must be both consulted and contested. Accordingly, a different model follows.

Lehar's tackling of encompassing space is an important step, but other phenomenological details might support a different, less holistic model – a *stage* or *set,* not a bubble: Several quite different aspects of our visual ecology afford distance information. Their zones of efficacy, as in Figure 1A (after Cutting & Vishton 1995), are surely important for any account of our encompassing visual world. Assume that the furthest zones form an essentially equidistant region like the *backdrop* on the stage in Figure 1B. Railroad tracks visible in those zones appear to converge. In nearer zones, the depth information effectively specifies the tracks as parallel and holds the backdrop in its place upstage.

This implies discontinuities (e.g., between backdrop and stage) that are not firmly fixed, because where the viewer attends, and with what intentions, affects what information is recovered and used (cf. Fig. 2B, C). Figure 1A can therefore serve only as a conditional account; and as Figure 1C implies, the phenomenal layout itself varies somewhat with the viewer's perceptual intentions. In this model, therefore, distance to the end of the internal world is not a continuous variable nor continuously defined. Why aren't the discontinuities spontaneously evident?

Is there evidence of such overlooked discontinuity? Figure 2A seems to reverse as a whole and has been offered as one example of how a minimum principle (including Lehar's version) leads to perceiving an entire three-dimensional structure (Hochberg & MacAlister 1953; Kopfermann 1930). But Figure 2B shows that, when tested, *perfectly possible* objects display the same dependence on what the viewer attends as was previously shown by the Penrose and Penrose (1958) *impossible* figures. *Perceptual consequences* (Hochberg 1998; in press), such as the effects of rotation described in Figure 2B and the surface-lightness effect in Figure 2C, attest that these are perceptual phenomena. They also share some aspects of Lehar's isomorphism. (And the absence of any salient break between the different spatial zones of the environment in Figures 1A and 1B, and in the apparently-continuous bubble that Lehar describes, merely parallels what happens within objects.)

Such phenomena raise difficulties for any holistic proposed isomorphism powered by the physical relationships as perceived. Gestaltist visions of isomorphism were of course concerned mostly with flat shapes, not three-dimensional structures (see Hochberg 1998). The fact that Peterson and her colleagues (see

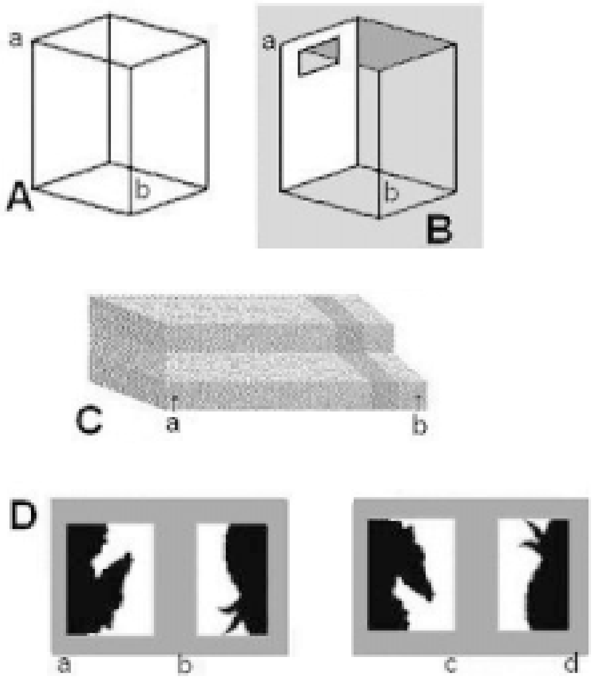

Figure 2 (Hochberg). Some shapes isomorphism must take. **A.** The reversible Necker cube. Sometimes offered as an example of how a minimum principle (or something like it, in Lehar's version) leads to perceiving an entire three-dimensional structure (Hochberg & McAlister 1953; Kopfermann 1930). **B.** The partly reversible Killer cube. When attended at (a), the present cube appears of definite and nonreversible three-dimensional structure; when attended at (b), it soon starts reversing, though the same Gestalt remains in view (though off attentional center). The reversals are attested by their perceptual consequences: When rotated clockwise around its vertical axis, the perceived motion is clockwise when (a) is attended; when (b) is attended and when it appears nearest the viewer, motion appears *counter*clockwise. Such perceptual consequences help validate one's otherwise unsupported phenomenology, as in the next figure (Hochberg, in press; Hochberg & Peterson 1987). **C.** Adelson's Impossible Staircase. With no discernible discontinuity, the right and left sides here are incompatible as three-dimensional structures; showing that they are actually seen that way. Note that the same print density appears of higher reflectance (lighter paint job) at (b) than at (a) – (after Adelson 2000, with permission); see text. **D.** Do configuration-based organizational factors first provide figure-ground segregation, which thereby offers a shape to be recognized? Not so you can tell: see text (see Peterson 1994; Peterson & Gibson 1993).

Figure 2D; cf. Peterson 1994; Peterson & Gibson 1993) have shown that meaningful (*denotative*) shapes preempt figural status when in their familiar orientations (Fig. 2Dc,d) but not when *the physically identical* configurations are inverted (Fig. 2Da,b), makes it hard even to imagine what an appropriate formulation of isomorphism would be like. A phenomenology centered on query-directed units of perceptual behavior, emulating the TOTE rubric offered by Miller et al. (1960), might be more effective (cf. Hochberg 1970; in press; O'Regan & Nöe 2001).

Does perception replicate the external world?

Donald D. Hoffman

Department of Cognitive Science, University of California at Irvine, Irvine, CA 92697. **ddhoff@uci.edu**
http://aris.ss.uci.edu/cogsci/personnel/hoffman/hoffman.html

Abstract: Vision scientists standardly assume that the goal of vision is to recover properties of the external world. Lehar's "miniature, virtual-reality replica of the external world inside our head" (target article, sect. 10) is an example of this assumption. I propose instead, on evolutionary grounds, that the goal of vision is simply to provide a useful user interface to the external world.

Lehar asserts that "The central message of Gestalt theory is that the primary function of perceptual processing is the generation of a miniature, virtual-reality replica of the external world inside our head, and that the world we see around us is not the real external world but is exactly that miniature internal replica" (target article, sect. 10, last para.). I wish to consider this assertion of indirect realism.

Suppose it is true. Then we do not see the real external world, nor do we hear, smell, taste, or in any other way perceive it. Instead, we perceive just the miniature virtual-reality (henceforth, mini VR) that we generate. Given this, what empirical grounds might we have for claiming that our mini VR replicates the external world? Perhaps we could compare objective measures of the external world against psychophysical measures of the mini VR. If mismatches are minor, we would have grounds for the replica claim. This process seems straightforward enough. The basic sciences measure the external world, and psychology the mini VR. So we simply compare data.

But this is too fast. It is not just psychologists who perceive only their mini VRs; all scientists, regardless of discipline, perceive only their mini VRs. So how do the basic scientists manage to measure the external world?

The trouble is that every time scientists try to measure the external world, whether they look through telescopes or microscopes, they see only their mini VRs. They extend their senses with countless technologies, but the technologies and their outputs are still confined to the mini VRs; for if they were not, then, according to indirect realism, the scientists could not perceive them. Hence, all scientists are confined to perceive only their mini VRs. If they wish to make assertions about the external world, even assertions that an external world exists, then these are necessarily, according to indirect realism, *theoretical* assertions. They are not direct measures. As Einstein notes, "physics treats directly only of sense experiences and of the 'understanding' of their connection. But even the concept of the 'real external world' of everyday thinking rests exclusively on sense impressions" (Einstein 1950, p. 17).

So indirect realism does not allow us incontrovertible empirical grounds to assert that our mini VRs replicate the external world. At best, it allows us to postulate an external world as a theoretical construct. Once we take the external world as a theoretical construct, then we have many options for the particular form of that construct. We can, as Lehar suggests, propose that our mini VRs are replicas of the external world. This is a particularly simple theory and, on the face of it, quite unlikely. Our best evidence suggests that mini VRs vary dramatically across species (Cronly-Dillon & Gregory 1991), and there are no evolutionary grounds to suppose that our species happens to be the lucky one that got it right. To assert otherwise would be anthropocentric recidivism.

Once we extend our gaze beyond the replica theory, many other possibilities arise. One class of possibilities is that there is little or no resemblance whatsoever between the external world and our mini VRs, but that instead our mini VRs are simply useful user interfaces to the external world, with no more need to resemble that world than a Windows interface needs to resemble the diodes, resistors, and software of a computer. Of course, we could not call a theory from this class an "indirect realist" theory because, by hypothesis, there is no realism. So indirect realism leads us to consider dropping indirect realism in favor of a broader and more likely class of theories. Let us call these new theories "user-interface" theories. For what they entail is that our mini VRs, rather than being replicas of the external world, are simply useful user interfaces to that world. Different species employ different user interfaces for their different purposes. The human user interfaces are simply a small set of the total, of special interest to us for only parochial reasons.

The move from indirect realism to user interface can be disconcerting, for it denies an anthropocentrism very dear to us: the assumption that our perceptions are privileged among all species. And it opens a Pandora's box of theoretical possibilities for the nature of the external world and its relation to our mini VRs. It has been convenient to assume that because there are neurons and synapses inside the heads that appear in our mini VRs, therefore there must be corresponding real neurons in real heads in the external world. But convenience rarely coincides with truth. It looked for millennia as though the sun and stars circled the earth, but we now know better. Even space and time themselves are not immune from this process, for as Einstein pointed out: "Time and space are modes by which we think and not conditions in which we live" (quoted in Forsee 1963, p. 81).

Moving from indirect realism to user interface does nothing to impede progress in modeling of the mini VR itself along the Gestalt lines proposed by Lehar. Nor does it impede progress in modeling the neural networks of the perceptual systems in our mini VRs. All this modeling can continue as it has. We simply realize that we are not modeling a replica of the external world; we are instead modeling our species-specific user interface to an external world. And in consequence we are far more cautious in our knowledge claims about the external world.

The move from indirect realism to user interface gives us more elbowroom in dealing with the hard problem of consciousness. The hard problem arises when we assume that neurons as we perceive them in our mini VRs are replicas of real neurons in the external world, and we must therefore figure out how those real neurons could possibly give rise to conscious experience. But if we drop the replica assumption, we now have a broader range of theoretical possibilities for what, in the external world, might correspond to neurons in our mini VRs. In this case our only limits in solving the problem are not the straitjacket of the replica assumption, but our imaginations.

ACKNOWLEDGMENT
This material is based upon work supported by the National Science Foundation under grant No. 0090833.

Psychological relativity

Donald Laming
Department of Experimental Psychology, University of Cambridge, Cambridge, CB2 3EB, United Kingdom. **drjl@cus.cam.ac.uk**

Abstract: "Psychological relativity" means that "an observation is a relationship between the observer and the event observed." It implies a profound distinction between "the internal first-person as opposed to the external third-person perspective." That distinction, followed through, turns Lehar's discourse inside-out. This commentary elaborates the notion of "psychological relativity," shows that whereas there is already a natural science of perceptual report, there cannot also be a science of perception per se, and draws out some implications for our understanding of phenomenal consciousness.

Lehar is lacking an essential idea. Physicists have it – "relativity" – but Lehar does not. Lehar mentions (sect. 1) "the internal first-person as opposed to the external third-person perspective" but fails to realise how that distinction impacts on his discourse. If the implications of that distinction are followed through, the entire body of problems addressed is turned inside-out. The overriding principle that Lehar is lacking is:

> an observation is a relationship between the observer and the event observed

and thereby depends on the observer as well as the event. So, two observers in motion relative to each other make different determinations of the velocity of a third object (Galilean relativity). Figure 2 sketches the set-up for Thouless's (1931a; 1931b) phenomenal regression to real size. The observer has a different view of the experiment to the experimenter.

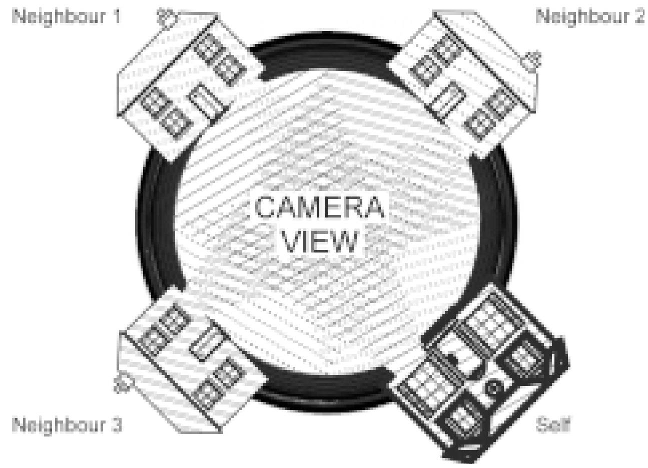

Figure 1 (Laming). The different views from four houses on a housing estate. (© 2004, Donald Laming. Reproduced with permission from D. Laming, *Understanding human motivation*, Blackwell.)

Figure 1 presents an analogy. Looking out from my window, I can see three other houses, separated from me by a road and a green sward. If there is a car in the road, my neighbour and I can readily agree that it is red. By agreeing on a suitable instrument for measurement, we can agree the colour of the car to whatever precision we desire. That arena outside our houses (*camera view*) is part of the public domain within which experiments can be conducted. But my neighbour and I cannot see into each other's houses. If I telephone my neighbour, I can only describe my interior furnishings by reference to what my neighbour will have seen elsewhere. The scope of experimental procedure can be extended to internal experience only by projecting that experience into the public domain. I might describe my curtains as scarlet, or carmine, or cerise – but my neighbour might think of a different colour referent to the one that I have in mind, and "seeing red" will then mean slightly different things to the two of us.

I can invite my neighbour into my house to see for himself but I cannot give him direct access to my visual experience. One might suppose that my internal visual experience could be measured, like the colour of the car in the road. But experimental psychologists have been trying to measure internal sensations for 150 years and have so far progressed nowhere (Laming 1997).

Some part of our visual experiences can be shared with others; the remainder is private. The Gestalt properties surveyed in sections 5 and 7 belong to that private part, which is why Gestalt psychology has not proceeded beyond verbal description. There is a boundary between experiences that can be shared and experiences that are essentially private. It is determined by what, within my field of view, my neighbour can also see (see Fig. 1). That is, the boundary is determined within my *neighbour's* field of view and is not to be found within my own visual experience. My own experience by itself contains no distinction between that which lies in camera view and that which is private. The junction is seamless. It is only too easy to confound subjective experience with objective observation; this is what Lehar has done.

It follows that there cannot be a natural science of perception. There is a science of perceptual report, a tradition that goes back to Fechner (1860/1966). But perceptual reports cannot be taken at their face value (here the Gestalt psychologists erred); rather, they must be evaluated by experiment. Lehar is aware of this (sect. 5.2), but asserts that perceptual experience is isomorphic to the neural substrate and thereby denies this distinction.

Lehar's stance is that "the world of conscious experience is accessible to scientific scrutiny after all, both internally through introspection and externally through neurophysiological recording" (sect. 2.3, para. 9). He envisages an isomorphism between perceptual experience as described by the observer and the observations of the natural scientist. Thouless's (1931a; 1931b) experiment on phenomenal regression to real size (Fig. 2) shows why such an isomorphism is not found in nature.

The observer's task is to select a disc set normal to the line of sight at distance *a* to match the *angular size* of the larger disc at distance *b.* Although people do choose a smaller disc from the alternatives at *a,* they systematically choose one too large to match (phenomenal regression to real size). Imagine that a neurophysiologist is making observations at the neural level of description relevant to understanding how and why this error of judgment occurs. If the observer's perceptions stand in the same relation to the neural substrate as the neurophysiological observations, then there has to be an internal "observer" looking at internal processes with the same objectivity as the neurophysiologist. The fact that Lehar has a mathematical model to replace the neurophysiological observations does not alter this requirement. This observer is represented by the "thinks bubble" in Figure 2. Philosophers will immediately identify this internal observer as Ryle's (1949) "ghost in the machine" (which is why the "thinks bubble" is decisively crossed out).

I next ask whether the hypothetical neurophysiologist can also observe the neural substrate of this "ghost." If so, the relationship of the ghost to the neural substrate is structurally different from that of the neurophysiologist; otherwise the "ghost" is pure mind-stuff. In fact, verbal descriptions of what is perceived are produced by the same system as that which does the perceiving, and the relationship of "observer" (if that term may still be used) to the neural substrate that is supposedly "observed" is essentially different from that of a third-party neurophysiologist. Several conclusions follow:

There need not be any useful isomorphism between neural process and perceptual experience.

Modelling perceptual experience is not an alternative to understanding the neural process.

There cannot be a natural science of perception, distinct from the study of perceptual report.

The idea of psychological relativity also impacts on consciousness (sect. 6). Because it is impossible to access any other person's subjective experience, it is not possible to observe any other person's consciousness. Even if the hypothetical neurophysiologist were to observe and record a substrate in the brain that subserved consciousness, there is no way in which the observations could be identified as such. However much one explores the brain, all that one finds is brain function. Phenomenal consciousness is simply the quality of subjective experience.

Lehar's discourse has neglected some real empirical relations between perceptual report and experimental observation. I give two examples. Rubin (1921) drew attention to the "figure-ground" phenomenon, the assertion that the first stage in visual perception was the separation of a figure from its background. Elementary neurophysiological study has revealed that sensory neurons are differentially coupled to the physical input (Laming 1986), so that they are specifically sensitive to boundaries in the visual field while responding with only a noise discharge to uniform illumination. This appears to match the "figure-ground" phenomenon. Second, the Necker cube is ambiguous as a visual stimulus. The ambiguity is temporarily resolved by factors from within the perceiver (sect. 7.3). But there is no reason why those internal factors should be consistent, comparing one instance with another, so that the project of constructing a consistent geometry of subjective perceptual space is not achievable.

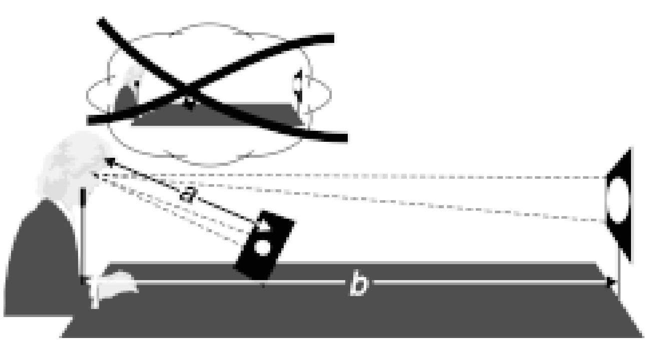

Figure 2 (Laming). Experimental set-up for the measurement of phenomenal regression to real size. (© 2004, Donald Laming. Adapted with permission from D. Laming, *Understanding human motivation,* Blackwell.)

Double trouble for Gestalt Bubbles

Dan Lloyd

Department of Philosophy, Trinity College, Hartford, CT 06101; Helsinki Collegium for Advanced Studies, University of Helsinki, Finland.
dan.lloyd@trincoll.edu

Abstract: The "Gestalt Bubble" model of Lehar is not supported by the evidence offered. The author invalidly concludes that spatial properties in experience entail an explicit volumetric spatial representation in the brain. The article also exaggerates the extent to which phenomenology reveals a completely three-dimensional scene in perception.

The real world is a place of many properties; so also is its presentation as a phenomenal world in the conscious brain. One way for a brain state to present in experience a worldly property *P* is to duplicate *P* itself. Like a painter striving for perfect mimesis, an embodied consciousness might use patches of red in the head to represent a red apple. Or, according to Lehar, a brain might use spatial properties to represent external spatial reality:

> The central message of Gestalt theory is that the primary function of perceptual processing is the generation of a miniature, virtual-reality replica of the external world inside our head, and that the world we see around us is not the real external world but is exactly that miniature internal replica. (target article, sect. 10)

Lehar's article makes the case for the internal replica, or "Gestalt Bubble," and then develops a model of how three-dimensional spatial modeling could occur in something like a neural medium. In this commentary, I suggest that the evidence in support of the Gestalt Bubble is in double trouble. It is both conceptually and phenomenologically flawed.

The coffee in the cup at my elbow is (to me) hot, brown, of a certain weight and size, and in a specific location. We cannot conclude, however, that the state of my brain that is my consciousness of the coffee replicates any of these properties itself. Yet this is an inference Lehar seems to make repeatedly in the target article. For example: "The fact that the world around us appears as a volumetric spatial structure is direct and concrete evidence for a spatial representation in the brain" (sect. 5.2).

This is a non sequitur, as can be seen by substituting "colored" for "spatial" in the passage. A slightly more elaborate argument is no less fallacious:

https://doi.org/10.1017/S0140525X03410091

> The volumetric structure of visual consciousness and perceptual invariance to rotation, translation, and scale offer direct and concrete evidence for an explicit volumetric spatial representation in the brain, which is at least functionally isomorphic with the corresponding spatial experience. (sect. 5.1)

Lehar is right that functional isomorphism between phenomenal experience and its implementation is required to avoid "nomological danglers," but once again, "explicit volumetric spatial representation" is in no way entailed – for "rotation, translation, and scale" substitute "hue, saturation, and brightness," and the fallacy will be apparent. Nor does Lehar's claim that phenomenal spatiality preserves the relational structure of spatial objects entail an internal replica, because (once again) a three-dimensional relational structure defines "color space" without in the least implying that the color solid appears somewhere in our brain. Functional isomorphism, meanwhile, is readily preserved between spatial objects/scenes and their representations without invoking replicas. For example, the World Wide Web is well stocked with virtual worlds that preserve functional isomorphism with spatial scenes, each of them encoded in some nonspatial computational idiom such as VRML.

In sum, the conceptual arguments in the target article do not support the author's main conclusion. Nonetheless, the brain does have properties, and some of its properties do determine the contents of conscious experience. Lehar's arguments do not establish that the brain must use space to represent space. Does phenomenality license any inferences at all about the neural medium? There are two ways to approach this question, beginning either with contingent generalities about perception or with its essential structures. The first approach begins with features of phenomenality (as revealed by perceptual psychology, including the Gestalt demonstrations of our perceptual capacities). The second analysis isolates essential or necessary structures of phenomenality. The second approach accords with classical phenomenology, as exemplified in the works of Husserl (e.g., Husserl 1974). In either case, the hope is that the analysis of phenomena will constrain the search for computational architectures sufficient to generate some or all of the features of phenomenality.

On neither approach is there compelling reason to posit the spatial virtual world proposed by Lehar. I do not doubt that I live in a spatial world, but my visual field – that is, what I see before me right now – conveys far less spatial information than Lehar's Gestalt Bubble encodes. At the focus of attention I am aware of surfaces, distance from my eyes, and edges, but outside of focal attention I experience only a very indefinite spatiality, which seems to me to be inconsistent with the continuously present three-dimensional models constructed in the Gestalt Bubble. The supposition that my experience specifies a full 360-degree diorama in my head arises from the "just-in-time" availability of spatial information with every attentional focus. The information is there when and where I need it, and experience presents an ordered sequence of focally attended presentations rather than a single wraparound replica of the spatial world. This seems to be phenomenologically "given" but it is also amply confirmed in psychological studies of "inattentional blindness" (Mack & Rock 1998) and "change blindness" (Simons 2000). (Sect. 8.8 briefly acknowledges the effect of successive gaze fixations in different directions, suggesting that parts of the replica fade while outside the visual field. This suggests either that the replica has an absolute spatial orientation and does not turn with the head or, if the replica does turn with the eyes, that only a small focal part of it has the spatial detail Lehar describes.)

This disagreement can be made more rigorous and more properly phenomenological. One essential property of the phenomenal world is expressed in our ability to distinguish properties by location. That is, I can be aware of a red circle and a green square at the same time without confusing the pairings of colors and shapes. Austen Clark refers to the problem posed by this pervasive perceptual ability as the "Many Properties" problem, and he argues that it can be solved only by coding places along with other perceptual properties (Clark 2000). So "red" and "circle" must be assigned a location, and "green" and "square," a second location. Experience, of course, solves the Many Properties problem easily, and arguably it is essential to the very concept of phenomenality that consciousness solve it. This argument so far provides support for Lehar's position but immediately raises the question: How many spatial dimensions are required? Lehar advocates three, Clark suggests two, but the argument necessitates just one, a linear dimension along which one point is tagged "red" and "circular," and another "green" and "square." The basic dimension, then, would be temporal, and experience would be an orderly ensemble of phenomenal leaps and bounds, a time line. Spatiality emerges from trajectories encoded in proprioception that orient each momentary percept to those before and after. This proposal conforms well with classical phenomenology (Husserl 1966; 1974), and in other work, I present evidence for its implementation in the brain (Lloyd 2002; 2003). This alternative cannot be defended here, but it does suggest that the Gestalt Bubble is not entailed by phenomenology.

It is important that theories of perception accommodate the Gestalt observations; Lehar brings forward an essential array of examples to consider, and exhibits the care and detail required to translate spatial perception into a computational model. But more evidence to support the model – from philosophy, phenomenology, psychology, and neuroscience – will be needed.

Isomorphism and representationalism

Riccardo Luccio

Dipartimento di Psicologia, Universita' degli studi di Firenze. 50123 Florence, Italy. **luccio@cesit1.unifi.it**

Abstract: Lehar tries to build a computational theory that succeeds in offering the same computational model for both phenomenal experience and visual processing. However, the vision that Lehar has about isomorphism in *Gestalttheorie* as representational, is not adequate. The main limit of Lehar's model derives from this misunderstanding of the relation between phenomenal and physiological levels.

Gestalt psychology has been fundamentally misunderstood in the United States (though the field too has to bear some responsibility; see Kanizsa 1995). After World War II, it had a meager destiny, cultivated only marginally in Germany and America though more intensively in peripheral countries such as Italy and Japan. However, mainly in the last few decades, some concepts of Gestalt psychology have appeared frequently in psychological debate, such as *prägnanz,* isomorphism, minimum principle, and so forth. The continuing debate demonstrates the inability of cognitive psychology to accept some highly significant aspects of our way of picking up the reality that is around us. Lehar's paper does not confine itself to stressing the importance of some classic Gestaltist ideas taken in isolation, as other scholars in the past have done, in an attempt, never completely successful, to integrate part of the *Gestalttheorie* into cognitive psychology. Instead, Lehar tries to build a computational theory that succeeds in offering the same computational model to both phenomenal experience and visual processing.

This highly interesting attempt deserves some comment, however. In my opinion, Lehar's vision of *Gestalttheorie* is not fully adequate, and this has some consequences for his theorizing. The point on which I disagree almost completely with Lehar is the following: He claims that there is a central philosophical issue that underlies discussions of phenomenal experience, as seen, for example, in the distinction between the Gestaltist and the Gibsonian view of perception. Is the world we see around us the real world itself or merely an internal perceptual copy of that world generated by neural processes in our brain? In other words, this is the question of *direct realism,* also known as *naïve realism,* as opposed to *indirect realism,* or *representationalism.* I note parenthetically that although Gibson (1966; 1979) called himself a naïve realist, this was only a provocation. The theory of direct perception is neither naïve nor realistic. As Michaels and Carello (1981, p. 90) clearly put it, "the test of the veridicality of perception concerns the mutual compatibility of the action of the actor/perceiver with the affordances of the situation." Here we are very far from the veridicality requested by genuine naïve realism.

More important is the picture of Gestalt psychology that Lehar offers to us. It is well known that in *Gestalttheorie* there was a strong Spinozian attitude. For example, Wertheimer (greatly impressed by Spinoza's *Ethica* from childhood on: see Luchins & Luchins 1982) remained in this orientation all his life. So we can speak in terms of an indifference or "indifferentism" about the problem of representation. In general, Gestaltist isomorphism has to be considered as a variant of psychophysical parallelism (see Boring 1942; 1950, mainly Ch. 13; for a recent survey of this issue, see Luchins & Luchins 1999). But the same could be said about almost all other Gestalt psychologists. Lehar quotes Köhler extensively. But Köhler never said that "the world we see around us . . . (is) . . . generated by neural processes in our brain" (target article, sect. 2, para. 1). Köhler, indeed, was in some instances a little ambiguous on this topic (e.g., Köhler 1969). But he was absolutely clear when he had to address the mind-body problem directly. He conceived the Gestalt position as a variant of parallelism (Köhler 1960, pp. 20–21), and said: "The thesis of isomorphism as introduced by the Gestalt psychologists *modifies* the parallelists' view by saying that the *structural* characteristics of brain processes and of related phenomenal events are likely to be the same" (emphasis added).

Lehar, quoting Köhler (1969), insists that the isomorphism required by Gestalt theory is not a strict *structural* isomorphism but merely a *functional* isomorphism. But Köhler always spoke of structural isomorphism. He was very clear in stating (Köhler 1940, Chs. 2 and 3) that the processes that run in our brain do not have any necessary correlate in our phenomenal experience. What is *structurally* identical is their interaction with what happens in bordering areas of the brain and the interaction that there is in the phenomenal field: Their dynamics and the dynamics of the phenomenal field.

The structural identity between the phenomenal world and physiological processes does not imply any causal relationship between the two levels. It means only that we are made up of one, and only one matter. The physical laws that rule matter lead to structurally identical outcomes when we consider the phenomenal level as well as the physiological one. In this sense, Gestalt psychology is neither representationalist nor antirepresentationalist; it is indeed indifferentist.

The main limit of Lehar's model derives, in my opinion, from this misunderstanding. His computational model, as I can assess it, works perfectly for a world that is organized in terms of soap bubbles (Koffka's metaphor [Koffka 1935], used too by Attneave 1982). A soap bubbles world is, in Gestalt terms, a world in which the forces of the perceptual field tend to dispose themselves to make an outcome that is maximally *good,* pregnant in the sense of *ausgezeichnet.* In Lehar's model, this happens at the phenomenological as well as the neurophysiological level. The fact is that – as I believe Kanizsa and I have demonstrated (Kanizsa & Luccio 1986; 1990) – a tendency of this kind does not exist in perception. These tendencies are instead well present in thinking, in memory, in all that Kanizsa (1979, Ch. 1) called "secondary processes," to distinguish them from primary processes of perception. But they are beyond the scope for which the concept of isomorphism is interesting – and relevant.

In recent years, a few other computational models have been presented to account for some typically Gestaltist phenomena, from information theory, to coding theory, to group algebra. However, Lehar is right when he says that they cannot account for both the phenomenal level and the neuropsychological level. I should stress that there is at least one exception: nonlinear dynamic systems, and, in particular, the synergetic approach. Apparently, we have not yet at our disposal a fully comprehensive theory; it should be interesting to test if the model proposed by Lehar could be integrated with other approaches.

The unified electrical field

William A. MacKay
Department of Physiology, University of Toronto, Toronto, ON M5S 1A8, Canada. **william.mackay@utoronto.ca**

Abstract: The electrophysiological perspective presents an electrical field that is continuous throughout the body, with an intense focus of dynamically structured patterns at the cephalic end. That there is indeed an isomorphic mapping between the detailed holistic patterns in this field and in perception (at some level) seems certain. Temporal binding, however, may be a greater challenge than spatial binding.

The independent processor model of individual neurons has given rise to the widespread impression, echoed by Lehar, that neurophysiology fails to deliver a unified basis for the holistic properties of perception. If there is any "illusion," it is not in the unity of perceptual awareness but in the portrayal of physical separation by techniques such as extracellular recording and fMRI. Overlooked is the axis of continuous activity stretching from the spinal cord to the cerebrum. The tonic activity in the brain stem activating systems (cholinergic, serotonergic, and noradrenergic), plus the histaminergic activating system of the hypothalamus, is responsible for our state of (un)consciousness (Pace-Schott & Hobson 2002). All sensory and motor activity feeds into this axis and influences the general distribution of activity. Also, the activating systems can directly trigger synchronization of activity within the cerebral cortex (Munk et al. 1996).

Furthermore, it is extremely doubtful that action potentials are of much significance in the direct link to perception. They are far too fleeting. It is the more sustained membrane potentials that are likely to correlate the best. Discrete neuronal activity in the brain, however isolated it may appear, is simply a local distortion in an unbroken continuum of electrical flux. All cells produce membrane potentials, even if static, such that an electrical field encompasses the entire body. The "panexperientialism" view would also suggest that perceptual awareness is linked to something like an electrical field. This is the only obvious property that is shared by both the atom and the organism, and it is increasingly elaborated as one ascends to the organism. One might postulate that the higher the degree of complexity in the electrical field, the higher the level of consciousness experienced. Using fMRI, it can be seen that the same cortical areas are active whether a stimulus is perceived or not. The difference in the case of perception is that the level of activation is greater (Moutoussis & Zeki 2002). This could mean either that more neurons are depolarized within the given area, or that the same synapses are active, but at a higher frequency, or both.

Neurons and their attendant glial cells manipulate membrane potentials like no other part of the body. This is their "game." Many attributes of neuronal electrical activity extend the range of information coding. No single one of them is the essence of conscious perception, but collectively they can raise (or lower) the level of consciousness. Spike synchrony is unquestionably relevant. For example, Riehle et al. (2000) have shown that unit pairs in the motor cortex synchronize activity to a very significant degree exactly at the moment of an expected signal. However, synchrony is not essential for "binding." In area MT, Thiele and Stoner (2003) recorded from pairs of units, one pair preferring the direction of motion of one visual grating, and the other preferring another grating direction. The units did not usually synchronize activity when the gratings were perceived as moving together in a coherent plaid. Synchrony elicited by coherent plaids was the same as for noncoherent ones. Again, it is probably not spiking activity per se that is ultimately important, but the associated changes in membrane potential and possibly phenomena such as depolarization fields manifested in superficial layers of cortex (Roland 2002).

The various states of Lehar's Gestalt Bubble model can easily be construed as hypothetical neuronal feature detectors. One could not ask for a better set of discriminators of planar properties in depth, and I suspect that something very similar lurks somewhere in the association areas between V1 and the inferotemporal cortex. The transformation from a two-dimensional image on the retina to a three-dimensional percept would follow a process as outlined by Lehar when the stimulus is an everyday, familiar experience with established expectations. For any unfamiliar object, whether presented to the eye or hand, exploratory movement is requisite to clarify ambiguities. Here, Lehar is correct to emphasize the translation/rotation invariance of the perception, divorced from the motion of the explorer. The target is perceived as it relates to its environment external to the viewer. This is the essence of the great transformation from egocentric (parietal cortex) to allocentric representation (presumably in the hippocampus or prefrontal cortex). The constancy of the percept over time as another data sample is added with each exploratory movement is also rightly highlighted.

It is essential that perception integrate over time as well as space. Even within one sampling episode, different sensory attributes such as color and motion are processed at slightly different times, although they are perceived as a unity. Hence, Zeki and Bartels (1998) postulate the existence of multiple "microconsciousnesses" in the brain, which are asynchronous with one another. This raises the problem of how they are integrated. A simple possibility is that everything processed within a finite window is integrated, just as two colors flashed within less than 40 milliseconds are blended together. But it cannot be that simple, because haptic exploration of an object can continue for hundreds of milliseconds.

Figure-ground designation also involves time constraints. Neurons in the inferotemporal cortex that are selective for shape maintain that shape preference when light-dark contrast is reversed (negative image) but not when a figure-ground reversal is made. Just as the perception of shape depends on whether a visual region is assigned to an object or background, so the visual analysis of form depends on whether a region is perceived as figure or ground (Rubin 2001). One cannot relegate the problem of resolving border-ownership of edges to earlier stages in the visual stream. It occurs quickly, within 10–25 milliseconds of response onset and really requires feedback from higher cortical areas. Hence, it is an instantaneous, holistic decision of the entire visual system, presumably selecting the most probable choice.

Lehar's excellent model of perceptual processes gives neurophysiology some precise goals and direction. Hopefully, the outcome will be convincing evidence that every percept is associated with a unique distribution of neuronal activity. An immediate problem, however, is the elucidation of the mechanism for binding elements of a percept in time.

The soap bubble: Phenomenal state or perceptual system dynamics?

Slobodan Marković

Laboratory of Experimental Psychology, Faculty of Philosophy, University of Belgrade, Belgrade, 11000, Serbia and Montenegro. **smarkovi@f.bg.ac.yu**

Abstract: The Gestalt Bubble model describes a subjective phenomenal experience (what is seen) without taking into account the extraphenomenal constraints of perceptual experience (why it is seen as it is). If it intends to be an explanatory model, then it has to include either stimulus or neural constraints, or both.

While presenting the theoretical background of his approach, Lehar attempts to keep a critical equidistance toward both indirect and direct realism. However, instead of a radically new approach, he offers a combination of some constructivist and some Gibsonian premises. On the one hand, like many constructivists (e.g., Gregory 1971; Hochberg 1978; Marr 1982; Rock 1983), Lehar adopts a representational paradigm that defines perception as a subjective conscious description or as an internal virtual copy of the external world. On the other hand, inconsistent with the constructivists' perspective and more similar to the views of proponents of direct realism (e.g., Gibson 1979; Shaw & Bransford 1977; Shaw & Turvey 1981), Lehar does not postulate any mediating mechanisms that process the representations within a perceptual system.

Moreover, Lehar's exact position concerning the question of direct perception of distal objects is not quite clear. At one point he explicitly claims that "the internal perceptual representation encodes properties of distal objects rather than of a proximal stimulus" (sect. 9, last para.). At another point he states that "the direct realist view is incredible because it suggests that we can have the experience of objects out in the world directly, beyond the sensory surface, as if bypassing the chain of sensory processing" (sect. 2.1, para. 1). Why would the thesis that distal objects are mapping onto the phenomenological domain without neural intervention be incredible and mysterious, while the idea about the projection of internal representation onto the external perceptual world not be incredible and mysterious? How is it possible that perception is partially indirect (representational), and partially direct (distally oriented)?

In his criticism of neurophysiologic modeling, Lehar rejects not only the classical Neuron Doctrine, but also some recent holistic approaches (cf. Crick & Koch 1990; Eckhorn et al. 1988; Singer 1999). Hence, for Lehar, the greatest shortcoming of neural models is not atomism, but rather, the problem of neurophenomenal decoding. That is, how can a fully spatial (topographical) perceptual description be created from spatially less constrained (topological), or even completely abstract, symbolic, and nonspatial neural representation? I find that this epistemological question is a natural consequence of a hidden ontological dualism: How does one domain of reality (consciousness) know how to read and understand the codes coming from the other (neural) domain?

To paraphrase Koffka (1935), the ultimate task for perceptual science is to answer why things look as they do. In the case of Lehar's theory, this question might be formulated as the following: Why is the phenomenal volumetric space such as it is? Why is it nonlinear in a particular way? Implicitly, Lehar proposes that this is an intrinsic property of phenomenal space which is not in a causal relationship with any other domain of reality. My opinion is that without the precise specification of the extraphenomenological aspects of perception, such as the stimulus and neural domains, it is difficult to answer the question related to why the percept looks as it does. For instance, imagine the difficulty in explaining the path shape and velocity of the planet Earth's motion without taking into account the mass and motion of other cosmic objects (moon, sun, other planets, and so on). A description of the Earth's motion is not an explanation of its motion.

Even Gestalt psychologists, who widely utilized the phenomenological method, did not create pure phenomenological explanations of perception. For instance, Koffka (1935) used the soap bubble metaphor, not to describe some phenomenal bubblelike experience, but to point out some basic principles of perceptual (neural) system functioning. Attneave (1982) also used the metaphor "soap bubble system" to describe the economy of perceptual system behavior. Like the soap bubble, which tries to enclose the largest volume within the smallest surface, the perceptual system tends to reduce the global spending of energy (entropy, minimum tendency) while at the same time striving to increase its effective use (dynamics, maximum tendency) (cf. Köhler 1920/1938; 1927/1971; see also Hatfield & Epstein 1985; Marković & Gvozdenović 2001).

If Lehar intends to create a Gestalt-oriented theory of perception, he has to have in mind that according to the classics of Gestalt theory, the phenomenological Gestalten are the consequences of both internal (neural) and external (stimulus) constraints (Koffka 1935; Köhler 1920/1938; 1927/1971; 1947). Simply speaking, the perceptual system tends to attain the maximum efficiency with the minimum investment (internal neural economy), but the minima and maxima will always be relative to the given stimulus conditions (external stimulus organization). The effect of external "control" of a perceptual economy is an articulation of more or less *prägnant Gestalten,* or as Wertheimer stated in his famous Law of Prägnanz, the phenomenal organization of a percept will be as "good" as the prevailing conditions allow (cf. Koffka 1935).

Bursting the bubble: Do we need true Gestalt isomorphism?

Niall P. McLoughlin

Department of Optometry and Neuroscience, University of Manchester, Institute of Science and Technology (UMIST), Manchester M60 1QD, United Kingdom. **n.mcloughlin@umist.ac.uk**

Abstract: Lehar proposes an interesting theory of visual perception based on an explicit three-dimensional representation of the world existing in the observer's head. However, if we apply Occam's razor to this proposal, it is possible to contemplate far simpler representations of the world. Such representations have the advantage that they agree with findings in modern neuroscience.

Lehar proposes to model visual perception using his subjective visual experience as his source of data. He proposes a perceptual modeling approach because "conventional concepts of neural processing offer no explanation for the holistic global aspects of perception identified by Gestalt theory" (target article, Abstract). This allows him conveniently to ignore current research in visual neuroscience while concentrating on the central issues of the representation of the visual field and of our subjective visual experiences. As he correctly points out, the world we see and experience surrounding us exists only as nerve impulses within our head. Lehar proposes that because our subjective experience of the world is that of a high-resolution three-dimensional volume, and because this representation must exist in our heads, it must therefore be some form of a high-resolution three-dimensional structure. However, this does not necessarily follow. For example, on a computer system it is possible to generate a sparse representation of the world into which it is placed so that the computer could interact with objects in the world in a meaningful manner. Objects could be represented as tokens at such-and-such x, y, and z location, and so forth. There would be no explicit representation of empty space within this sparse representation. Who is to say what the subjective experience of the computer might be?

There is no doubt that my subjective experience of the world is that of a three-dimensional solid environment which I perceive in equal detail in all directions. Yet, as visual scientists and practiced observers, we know that this is patently not the case. Each of our eyes responds to incoming photons in a non-uniform manner and this non-uniformity is further exaggerated in the cortex. The over-representation of the fovea is magnified between the retina and cortex, and the multiple interconnected cortical regions amplify this distinction even further. Most naïve observers are surprised to discover that they have a fovea and amazed that they have a blind spot in each eye. How do we fool ourselves?

The very fact that we are genuinely fooled (until we make careful observations) calls into question the use of subjective experience as the basis for theories of visual perception. Furthermore, although the Neuron Doctrine is indeed the foundation for most modern neuroscience research, I refute the notion that this doctrine implies purely feed-forward models of neurocomputation. Certainly, recent findings in both neuroanatomy (e.g., Angelucci et al. 2002; Bosking et al. 1997) and neurophysiology (Kapadia et al. 2000; Levitt & Lund 1997) emphasize the roles played by feedback and lateral connections in visual processing. Likewise, a number of popular modern computational theories make use of feed-forward, feedback, and lateral connections (e.g., Grossberg 1994). If a Gestalt Bubble model subserves perception, then why do we have so many visual areas, each containing a retinotopic map of visual space?

Is there any evidence for Gestalt-like processes at work neurophysiologically? Recent electrophysiological recordings from as early as the lateral geniculate and V1 have found interactions well outside the classical receptive field (e.g., Blakemore & Tobin 1972; Felisberti & Derrington 2001; Jones et al. 2000; 2001; Kapadia et al. 2000; Levitt & Lund 1997; Solomon et al. 2002; Stettler et al. 2002). Although the source of these interactions (whether they are mediated by feedback or by lateral connections) remains to be elucidated, it is clear that many aspects of grouping, completion, and emergence may well arise from such nonlocal interactions. In addition, recent neurophysiological studies in the primate (e.g., Livingstone & Hubel 1988) suggest that different aspects of a visual scene are represented primarily in different visual streams and areas. Although there is some disagreement as to the amount of segregation of function, numerous neuropsychological studies in humans back up the suggestion that multiple representations exist for different attributes and/or functional roles. One such patient studied by Humphrey and Goodale (1998) suffered from visual-form agnosia (Farah 1990). She was unable to discriminate between visual forms, let alone recognize her friends and family, yet her color vision was close to normal and she could recognize shapes when placed in her hands. Such case studies suggest that the brain encodes the external world using multiple representations, each one perhaps subserving a different role or task rather than a single isomorphistic one.

What Lehar seems to have forgotten is that the high-resolution representation is generated only when we pay attention to the input and focus our eyes on the object or texture under inspection. We need not represent even our immediate environment in high resolution unless we need to interact directly with it. Why waste time and space representing the world in vivid detail when we interact with only a small part of it at any one time? Surely our central representations should be goal-directed. We can always direct our vision to different locations in a scene to find out what is there, and given that most useful scenes are dynamic, why waste effort representing space in high resolution when it is constantly changing? O'Regan (1992) argued along a similar line when he suggested that "seeing constitutes an active process of probing the environment as though it were a continuously available *external memory*" (p. 484, emphasis in original). He suggests that seeing does not involve the reification of a three-dimensional spatial representation of the external world in the observer's head but rather depends on one's ability to interrogate the environment through directed eye movements. It may well be that we have a fuzzy three-dimensional representation of the external world in our heads that we use to help direct eye movements, but I remain to be convinced that we would need or want anything more complex. If we need the detail, we look.

Given the lack of physiological evidence for such a complex and computationally expensive representation, coupled with the lack of necessity for such a complete representation, Occam's razor suggests we burst this Gestalt Bubble model.

Relations between three-dimensional, volumetric experiences, and neural processes: Limitations of materialism

Axel Randrup
International Center for Interdisciplinary Psychiatric Research, Bygaden 24 B, Svog. DK-4000 Roskilde, Denmark. **arandrup@mobilixnet.dk**
http://www.mobilixnet.dk/~mob79301

Abstract: Certain features of perception – the quale red, for example, and other qualia – must be regarded as additions to the materialist neurophysiological picture of perception. The perception of three-dimensional volumetric objects can also be seen as qualitative additions to the neurophysiological processes in the brain, possibly without additions to the information content.

In the history of science and philosophy, the world has been regarded as material, mental (idealist philosophy), or dualist (both material and mental). Like many people today, Lehar has chosen the materialist view, and he attempts to avoid dualism by assuming the mind-brain identity position ("consciousness is a physical process taking place in the physical brain" – sect. 2.3, para. 5). Still, he writes that there remains a subjective quality (or *quale*) to the experience of red, for example, which is not in any way identical to any physical variable in the brain. I think this must mean that the experience of qualia *adds* something to the assumed material world and that Lehar therefore does not stay consistently within the materialist frame of reference. Lehar also writes (sect. 2.3) that sense data, or the raw material of conscious experience, are the *only* thing we can know actually exists, and that all else, including the entire physical world, is informed conjecture based on that experience. To me this statement appears as a departure from materialism; it is actually close to the idealist view.

I now suggest that the perceptual experience of three-dimensional, volumetric objects, and of empty space is also something that "subjective conscious experience" adds to the assumed material electrochemical processes in the brain, possibly without changing the information content – a qualitatively different representation. Lehar thinks that the gap between the materialist descriptions of neurophysiology and the phenomenological descriptions of Gestalt features of perception may be due to the present "embryonic" state of neurophysiology, but I regard this as a promissory belief rather than an explanation.

Analogously (and staying within the materialist frame of reference) I believe that a computer can produce a three-dimensional, volumetric figure, namely, if it is connected with a device that can construct that figure. The figure will then be another representation of the information content which is represented inside the computer by electrical processes. Of course, a human person can also construct a three-dimensional figure with his hands or describe it in words and drawings, as Lehar does. In this case, it is the connection with the body, particularly with the muscles and the hands, that enables the brain to make these constructions and descriptions from its information content.

I think that materialism has served science well within a rather large domain, but with studies of cognition such as Lehar's, we move into a domain where materialism reveals significant shortcomings. I find that such shortcomings appear in Lehar's work.

Hence, on his materialist background, Lehar rejects direct (naïve) realism which suggests that we can have experience of objects out in the world directly, as if bypassing the chain of sensory processing. Provided that the materialist background is retained, I agree with this rejection. But if we apply an idealist worldview, our perceptions are of course experienced directly, and based on these perceptions we form concepts, such as the concepts of a "material" object, a "material" world, and perceptual models such as Lehar's Gestalt Bubble model. I see these concepts and models as mental constructs representing features of the perceptual reality, such as quantitative features and three-dimensional Gestalt features. These constructs are of course also experienced directly, and they can be made unambiguous and precise. Here I agree with Lehar, who thinks that perceptual models remain "safely on the *subjective* side of the mind/brain barrier" (emphasis in original) and writes about "objective phenomenology" leading to "perceptual modeling" (sect. 4). It is when we accord "material" concepts a special existence of their own, principally different from the existence of conscious experiences, that is, when we move to materialism, that we run into trouble with direct realism.

Lehar finds troubles with indirect realism as well but eventually accepts this view on the premise that the world we see around us is not the real external world but a miniature virtual-reality replica, an internal data structure within our physical brain. I think this view gives only an incomplete, imprecise conception of the "external world," including our "physical brain." This incompleteness and imprecision are shared with other philosophies assuming indirect realism, such as "hypothetical realism" (Löw 1984; Randrup, submitted; Wuketits 1984), "commonsense realism" (Ruse 1986), and Kant's concept of "the things in themselves" versus "the things for us." According to Kant's philosophy, we actually know nothing about things in themselves, except that they are supposed to exist. I think that this uncertainty or renunciation of knowledge compares unfavorably with the precision of the "material" concepts based directly on perceptual data in the idealist worldview.

Another shortcoming of materialism in relation to the study of cognition is that it is difficult consistently to avoid dualism, as appears from Lehar's views about qualia mentioned above. And if dualism is admitted, it is hard to see how conscious experiences can be generated by material processes in the brain, as Lehar thinks they are (sect. 2.4). In the alternative idealist view of the world, it is not so hard to see, conversely, how "material" *concepts* are generated by the mind; the history of science shows how such concepts have been created (e.g., quanta, superstrings) or deleted (impetus, phlogiston, the ether) following the advent of new perceptual (observational) experiences. The special material type of *existence* is not a part of the idealist philosophy. (For a more extensive discussion of the mind-matter and mind-brain problems in relation to cognition, see Knight 2001; Randrup 1997; 2002.)

Actually I think that Lehar's study, based on "the primacy of subjective conscious experience" and leading to a model of phenomenal perception, is most readily understood within the idealist worldview, and within this view his troubles with direct and indirect realism, with materialist monism, and with mind-matter relations will be significantly reduced. For more about the idealist worldview proposed here, see Randrup (1997; 2002).

Consciousness as phenomenal ether?

Antti Revonsuo
Department of Philosophy, Academy of Finland, Center for Cognitive Neuroscience, University of Turku, FIN-20014 Turku, Finland.
revonsuo@utu.fi

Abstract: The Gestalt Bubble model of visual consciousness is a courageous attempt to take the first-person perspective as primary in the study of consciousness. I have developed similar ideas as the Virtual Reality Metaphor of consciousness (Revonsuo 1995; 2000). I can, hence, only agree with Lehar about the general shape of a proper research strategy for the study of consciousness. As to the metaphysical basis of the research program, I have, however, several reservations about panexperientialism.

I agree with Lehar on several points but disagree about the ultimate metaphysical nature of consciousness. I shall first describe points of agreement and then proceed to a criticism of panexperientialism. First, any research program on consciousness should start by taking the *explanandum* seriously, constructing a systematic description of it. This is Lehar's "objective phenomenology." In the context of the biological sciences, this is the initial, descriptive stage of inquiry. All branches of biology have begun with the descriptive stage, and the study of consciousness should be no exception.

Second, in the study of consciousness the top-down approach should be of at least as much importance as the bottom-up approach. Once we have a detailed description of the structure, organization, and dynamics of a higher level of organization (in this case, subjective phenomenology), it will impose significant constraints on the possible lower-level (neural) mechanisms that could account for the higher-level features. The lower-level mechanism must be capable of supporting exactly the kind of structure, organization, and dynamics as is found at the higher level of phenomenology; otherwise the proposed mechanism is not a plausible candidate to explain the phenomenon. The bottom-up strategy is important too, but it should be combined with the top-down strategy. Otherwise bottom-up approaches may lead either to the elimination of consciousness (because it is so difficult to see how single-neuron activity could add up to holistic features of consciousness), or to the search for the mere neural *correlates* of consciousness (rather than the directly underlying constitutive mechanisms that explain the phenomenon), because the signals that are collected from the brain usually originate nowhere near the higher levels of organization where consciousness itself resides (Revonsuo 2001).

Third, indirect realism as a theory of perception seems to be the only alternative that can give a plausible explanation of dreams and other hallucinations. Dream experiences show that the brain in rapid eye movement (REM) sleep can bring about a fully convincing simulation of the perceptual world and a simulated self embodied inside this virtual world. Dreams are temporally progressing "being-in-the-world" experiences generated inside the brain. During dreaming, phenomenal consciousness is causally isolated from the stimulus environment, from the concurrent state of the physiological body, and from behavioral output systems. As I have argued in my previous *Behavioral and Brain Sciences* commentaries on Pessoa et al. (1999) and O'Regan & Noë (2001), their theories of visual consciousness cannot account for our vivid visual experiences in dreams.

Although I therefore largely agree with Lehar as to what the proper approach to the study of consciousness should be, there is one core issue on which we seem to have differing views. His fundamental metaphysical commitment is to panpsychism (or panexperientialism), according to which (a simple form of) consciousness is a fundamental property of physical matter. According to this view there is no *radical* discontinuity between any physical systems as to the possession of consciousness; it is just a matter of degree. Everything is more or less conscious; simple physical systems to a lesser degree, the human brain perhaps to the highest possible degree. This smooth continuum of consciousness across all physical entities is supposed to have the following explanatory strengths: (1) consciousness is a fundamental property of physical matter and therefore need not be explained in terms of (nonconscious) physical matter; (2) there is no radical conscious/nonconscious dichotomy to be found anywhere in the natural order (e.g., in phylogeny or ontogeny).

This approach raises some severe problems. There are clear, well-demonstrated dichotomies between the presence and the absence of the state of consciousness (caused by anesthesia, epileptic seizures, fainting, coma) and between the presence and absence of particular contents of consciousness even though the stimuli are implicitly processed (as in blindsight or neglect). Any theory of consciousness should be able to explain these radical subjective differences between the presence and absence of consciousness. The panexperientialist is, however, forced to say that these are not really cases where the presence and total absence of consciousness in the brain could be strictly contrasted. The contrast is only between *primitive* and *more sophisticated forms* of consciousness. According to the panexperientialist, the primitive form may be something so simple that we would hardly recognize it as consciousness at all. Hence, what we thought was the total absence of experience is actually the presence of a primitive form of experience; we just cannot recognize it as experience.

Unfortunately, this move will not help us to understand the radical contrast between the presence and absence of conscious experience in the above cases. Regarding everything as conscious (to some degree) does not remove the radical conscious/nonconscious contrasts. In fact it leads to a position as difficult as (but the exact opposite to) the eliminativist position defended by Dennett. If we take either the panexperientialist position that phenomenal consciousness is *everywhere* in the world or the eliminative position that it is *nowhere*, we are no closer to explaining the radical empirical differences that we want to understand.

Furthermore, panexperientialism smacks of a misuse of the concept of experience. It is difficult to see why the postulated "primitive form" of consciousness – which we might not even recognize as experience – should be placed in the same category as our vivid phenomenal experiences. There seems to be no clear idea of what "protoconsciousness" could be, whether it exists at all, or how the claims for its existence could be empirically tested or theoretically modeled; and how exactly the primitive form of consciousness relates to our ordinary, vivid, phenomenal consciousness.

Hence, I do not regard panexperientialism as an advisable metaphysical commitment for a research program on consciousness. I would rather postulate that the sphere of subjective experience is a higher level of biological organization in the brain. Phenomenal experience exists only at that level and in those creatures whose brains can realize that level. Otherwise, the physical universe is devoid of phenomenal consciousness. When we totally lose consciousness, as we do during anesthesia, for example, our brain is temporarily incapable of supporting the phenomenal level of organization. The radical difference between the presence and the absence of phenomenal experience is to be described and explained in terms of biological levels of organization in the brain. Physical matter at lower levels of organization perhaps may be said to contain the *potentiality* of being conscious, but only in the weak sense in which all physical matter contains the *potentiality* to be alive. The mere potentiality does not make simple physical systems (say, carbon atoms or diamonds) alive, and it would be a waste of time to study the microphysical structure of diamonds in order to understand the biology of living systems. In a similar vein, I fear that the assumption that all physical systems (diamonds, toothbrushes, bacteria, and so on) are conscious (or "protoconscious") is going to be a useless, untestable hypothesis for the science of consciousness.

Protoconsciousness seems to be comparable to "ether," the invisible form of matter that was once believed to fill all physical space. The idea of a vacuum devoid of physical matter was unimaginable. Perhaps the idea of a "phenomenal vacuum" or the total absence of conscious experience is equally difficult to accept. But although there were genuine empirical phenomena that the ether models tried to account for, there seem to be no phenomena (either nonconscious physical or conscious phenomenal) that the phenomenal ether of panexperientialism accounts for. Furthermore, as far as we know there *are* total phenomenal vacuums, total absences of phenomenal experience, and we should not try to fill them by postulating a phenomenal ether that pervades all physical matter. Instead, our theories of consciousness should explain the definitive differences, both phenomenal and biological, between the total presence and the total absence of consciousness in the brain.

ACKNOWLEDGMENT
This research was supported by the Academy of Finland (project 45704).

Gestalt Bubble and the genesis of space

Victor Rosenthal[a] and Yves-Marie Visetti[b]

[a]*INSERM XR-324, Centre Paul Broca, 75014 Paris, France;* [b]*LATTICE-CNRS-ENS, 92120 Montrouge, France.* **victorro@broca.inserm.fr**
yves-marie.visetti@ens.fr

Abstract: Lehar (rightly) insists on the volumetric character of our experience of space. He claims that three-dimensional space stems from the functional three-dimensional topology of the brain. But his "Gestalt Bubble" model of volumetric space bears an intrinsically static structure – a kind of theater, or "diorama," bound to the visual modality. We call attention to the ambivalence of Gestalt legacy and question the status and precise import of Lehar's model and the phenomenology that motivates it.

Lehar should be applauded for making a strong case for the fundamental character of volume and depth in our experience of space. The originality of his proposal resides inter alia in the radical claim that a three-dimensional experience of space stems from the functional three-dimensional topology sustained by the human brain (not to be naïvely equated with brain topography). He posits that subjective spatial experience criterially requires a three-dimensional topological substratum – a device lacking a three-dimensional topological-dynamical structure could never account for the volumetric experience of space. In other words, the only viable option for a functionalist indifferent to brain physiology is three-dimensional topological-dynamical functionalism.

Lehar depicts his model as an outgrowth of Gestalt tradition. Indeed, one can easily recognize two essential features of Gestalt theory: its phenomenological approach to subjective experience, and the postulate of psychophysical isomorphism. Phenomenological space, its emergence, and its scientific explanation as a brain process are, according to Lehar, grounded in pregiven, continuous, and coherent topology, specifically a three-dimensional functional topology.

Lehar may not be aware that the way Gestalt psychologists treated space was in reality quite equivocal. Although they were in principle cognizant of the fundamental status of volumetric space, they granted it low priority in their scientific agenda and tended "provisionally" to treat space as a series of transparent/opaque surfaces, if not as ambient ground against which to set a figure. On the other hand, it is true that Köhler's theory of psychophysical isomorphism explicitly referred to three-dimensional functional brain topology to construe not only three-dimensional geometrical static structures but also two-dimensional structures evolving in time (see Koffka 1935). The theory combined empirical and phenomenological constraints with speculative brain physics (e.g., the theory of cortical fields) so as to represent both brain process and phenomenological experience in a single dynamical scheme (see Rosenthal & Visetti 2003).

Several attempts have been made to model Gestalt principles of perception in accordance with neurophysiology and in particular with the doctrine of *neural coding* (e.g., of perceptual microfeatures). For example, models of *neural fields* or *neural repertoires* feature a two-dimensional functional topology that corresponds to a topographic two-dimensional arrangement of units in primary areas (e.g., retinotopy) (e.g., Hoffman 1989; Koenderinck 1990; see Petitot 1999 for a review). Less discrete models, unconstrained by brain physiology, were developed in the context of image processing, and sometimes resorted to fairly complex mathematics, but maintained set to a bidimensionality of their input (retinal or pictorial; see Morel & Solimini 1995). The very idea of three-dimensional functional topology was hardly taken into consideration in the few attempts to account for depth (e.g., Grossberg 1994), which therefore had to resort to hosts of specialized coding units: a patently implausible solution, as Lehar rightly noted.

The solution advocated by Lehar is original and certainly deserves attention. He defines a three-dimensional topological *milieu* where any local element can be in one of four states (corresponding to local surface elements). Each individual element (or point in a perceptual matrix) exerts a field influence on adjacent elements for them to take on a similar state (or to be prevented from this by inhibition). Reciprocal determination between surface elements is assumed to generate equilibrium in which the relevant features are stabilized. The input to the model is an image set in the frontal plane (much like a retinal image). The output (actually the first step in "geometrization" of space) is a distribution of geometrical surface microfeatures in a three-dimensional space. Although Lehar does not mention this issue, one can readily deduce that unit formation or individuation is assumed to take place in this three-dimensional visual matrix. The originality of this proposal should be highlighted: Whereas the majority of rival models first individuate two-dimensional units (from two-dimensional image input), then categorize them as faces of three-dimensional units, Lehar sets his three-dimensional structure *ab initio,* and whatever is to populate this three-dimensional distribution of geometrical microfeatures supposedly comes next.

It is not clear, however, which scientific question Lehar has set out to answer. He does not seem to attempt another *perspectival* reconstruction of the visual field, for his model, in contrast to its alleged purely phenomenological motivation, builds on a physicalist metaphor. Although Lehar dismisses neurophysiological concerns, the analogy between his model and neural net models jumps to the eye: Traditional "neurons" with their receptor fields are replaced by elements or points in perceptual matrix, and neural connections are supplanted by fields of influence. Moreover, Lehar alludes to the possibility that the model may take a discrete or granular form (see target article, Fig. 7A). Why, then, does he hammer so loudly his physicalist credo? It seems that Lehar believes that the process by which space is *constituted* necessarily sheds light on the way we perceive space. Then why does he not try to motivate his model *genetically?* Clearly, Lehar needs to tell us the rules of the scientific game he plays more explicitly (does he want to model the constitution of space from a purely phenomenological viewpoint or does he attempt a free mathematical reconstruction of subjective experience?).

Lehar could have mentioned that during the past century other theorists put forth elaborate proposals concerning the constitution of space experience (e.g., Gibson 1950; Husserl 1907/1997; Poincaré 1905/2001). Instead of sticking to neurophysiology, they referred to the structure of the organism or the *lived body.* These were strongly dynamic, sensorimotor "models" of constitution of phenomenological space that assumed a multimodal origin of volumetric space and explicitly related its dimensionality to repertoires of self-generated movements. Although none of these "models" can be regarded as fully effective, they account for the ontogenesis of space in a *dynamic* fashion and for a variety of phenomena of adaptation (e.g., to distorting or inverting goggles). We suggest that considering the dynamics of the genetic, multimodal, and sensorimotor character of the constitution of space is as important in modeling perceived space as neurophysiology and the kind of static geometry on which Lehar elaborates. What comes along with such dynamics is the constitutive relationship between *external* and *bodily space.* Lehar appears to be aware that perception of space involves one's own body, but instead of taking this as a constitutive relation, he treats the body as just another object in space.

Finally, we have strong reservations with respect to Lehar's phenomenology. The field of vision he refers to neglects readiness for prospective action, and the phenomenological subject is not immersed in the *practical field* of ongoing activity with its qualitative, praxeological, and prospective dimensions (see Rosenthal & Visetti 2003). What about the nonisotropy of perceived space and the resulting potential heterogeneity in the constitution of regions of space? Is it advisable to consider phenomenological space as a mere deployment (be it three-dimensional), independent of the engaged or prospective actions to which it gives stage?

Neurological models of size scaling

Helen E. Ross

Department of Psychology, University of Stirling, Stirling FK9 4LA, United Kingdom. **h.e.ross@stir.ac.uk**
http://www.stir.ac.uk/psychology/staff/her1

Abstract: Lehar argues that a simple Neuron Doctrine cannot explain perceptual phenomena such as size constancy but he fails to discuss existing, more complex neurological models. Size models that rely purely on scaling for distance are sparse, but several models are also concerned with other aspects of size perception such as geometrical illusions, relative size, adaptation, perceptual learning, and size discrimination.

Lehar argues (sect. 2.2 and elsewhere) that there are no adequate neurological models to explain why we see the world the way we do, and that theorists have ignored the discrepancies between the proximal stimulus and our perceptual experience. He then presents a computational model to describe our perceptual experience of hyperbolic space. He rightly complains about the shortage of neurological models for size and shape constancy but he fails to discuss the models that do exist.

Psychologists have long been interested in size scaling, or discrepancies between perceived size and image size: The phenomena include size constancy, geometrical illusions, optical distortions, adaptation, and aftereffects. The classical account of size constancy maintains that size is scaled for distance in a quasi-geometric manner (the size-distance invariance hypothesis); this account is not productive of neurological models because it assumes that retinal image size is "correctly" encoded in the visual cortex and that the image is then scaled for distance in some unexplained "cognitive" manner. Kirschfeld (1999) argues that the image representation has to be scaled for distance neurologically before it enters consciousness and that this might be done in area V4. He notes that Dobbins et al. (1998) found that some neurons in this area varied their response to the angular size of lines depending on viewing distance.

The idea that image size is transformed at some preconscious stage of visual processing by mechanisms other than distance scaling (e.g., McCready 1985) may be more fruitful. Stuart et al. (1993) proposed a computational model based on broadly tuned layers of size detectors, which could account both for Weber's law in size discrimination and for the biasing effects of geometrical illusions; however, they did not extend the model to include scaling for distance. The main alternative approach to size constancy – generally supported by Gibsonians – is that object sizes are scaled in relation to the surrounding spatial scale. This approach has the advantage of embracing other size illusions in addition to size constancy and it is more productive of neurological models. Size contrast illusions have been attributed to adaptation of cells that detect spatial frequency or to other neural interactions in the brain (see Gillam 1998). However, spatial frequency is not the same thing as image size (the distance across an image), so spatial frequency models are unhelpful for general models of size perception.

Andrews (1964) proposed a perceptual learning model of size calibration in which the brain corrects the metric of the visual field according to the most recent information and attempts to equalize the spacing of contours. This would allow for learning in addition to explaining some illusions, aftereffects, and size constancy. Richards (1977) suggested that simple cells in the cortex might respond to relative rather than absolute size and he also discussed the properties necessary for the neural basis of size constancy.

Some authors have attempted to explain size constancy through the enlargement of perceived size for the central part of the visual scene, which occurs because the representation of the central part of the retinal image covers more cortical cells at later stages of analysis. Such an idea is based on the anatomical fact of cortical magnification, which enhances acuity for central vision. The fovea contains more densely packed cone cells than the surrounding area and it projects to a relatively larger region of the primary visual cortex. Schwartz (1980) incorporated this idea into his model of size constancy. When an observer fixates a distant object, it forms a small image in central vision, whereas close objects form larger images that spread further into the periphery: The small central image is therefore expanded neurologically relatively more than the larger image. Such a mechanism might contribute marginally to size constancy, but it fails to explain how objects of the same angular size can appear different in size even when both are viewed in central vision.

An example of this problem is the moon illusion (see Ross & Plug 2002). The moon illusion is the apparent enlargement of the sun or moon when low on the horizon compared with its size when higher in the sky; the effect is similar to size constancy but is hard to explain by the usual "scaling for distance" account. The difficulty is that the low moon appears nearer than the high moon, whereas size-distance invariance requires it to appear further. Trehub (1991, pp. 242–47) developed the "retinoid" model, which could account for both size constancy and the moon illusion. He argued that size magnification is expensive in neurological terms because it involves the use of more networks of cells. The brain husbands its resources by magnifying only the most "ecologically relevant" parts of the scene – that is, objects in the near distance when looking horizontally, and close overhead when looking up. Humans cannot normally interact with celestial objects or with distant terrestrial objects, so the images for such objects can safely be left relatively small. Size constancy is therefore poor for far horizontal distances and even poorer when looking upwards. The three-dimensional representation of distance is also shrunk vertically in comparison with horizontally, again for the purpose of minimizing neural resources. Distance is computed within the three-dimensional retinoid system and is represented by "sheets" of cells; the extent of size magnification is linked to the distance plane onto which the image is mapped. This biased mapping of the visual scene onto brain structures is largely the result of human evolution, but it can be further modified by individual experience.

There are neuropsychological findings that support multiple representations of three-dimensional space (see Previc 1998). There are also findings on micropsia and hemineglect that give clues as to how and where size might be coded (see Kassubek et al. 1999). Lehar may be correct that a simple Neuron Doctrine cannot account for size scaling, but more complex neurological models show promise.

Spatial phenomenology requires potential illumination

James A. Schirillo

Department of Psychology, Wake Forest University, Winston-Salem, NC 27109. **schirija@wfu.edu**

Abstract: Collapsing three-dimensional space into two violates Lehar's "volumetric mapping" constraint and can cause the visual system to construct illusory transparent regions to replace voxels that would have contained illumination. This may underlie why color constancy is worse in two dimensions, and argues for Lehar to revise his phenomenal spatial model by putting "potential illumination" in *empty* space.

Lehar's phenomenological description of space neglects the fact that *empty* space is actually *full* of illumination. For example, if a cast shadow crosses half of this page and you move your finger from a word under shadow to one under full illumination, you are not surprised when your finger crosses the shadow, even though your finger is closer in depth than the page. This is because every voxel between your eye and the page contains some amount of light. It is unfortunate that Lehar overlooks this fact, because he correctly asserts that depth information is volumetric, whereas current neurological models fail to "represent transparency[,] with multiple depth values at every single (x,y) location, or to represent the experience of empty space between the observer and a visible object" (sect. 3, para. 5). These same models also ignore that every voxel of "empty" space contains light of some intensity and chromaticity.

This confusion probably results from naïvely accepting the popular notion that humans care only about the location and qualities of objects, making the perception of illumination irrelevant. This assumption is so prevalent that much of color research is devoted to determining how the visual system "discounts the illuminant." However, a viable solution to the Gestalt problem of color constancy will emerge only with a more complete description and understanding of how we subjectively *experience* illumination. Ironically, Lehar's aspiration to describe the subjective experience of spatial vision in terms comparable to those of color vision reveals that current color vision research is also in peril. That is, he claims that color phenomena are reducible to hue, intensity, and saturation because that is how the brain represents them physiologically (sect. 2.3). Yet models of hue, intensity, and saturation cannot be the "primitives of raw conscious experience" (sect. 4, para. 3), in that these qualities remain invariant as illumination changes across space.

This confound is apparent when Lehar discusses his Figure 1 as containing "explicit volumes, bounded by colored surfaces, embedded in a spatial void" (sect. 5.1, para. 2), where "every point can encode either the experience of transparency or the experience of a perceived color at that location" (sect. 6). His more accurate intuition is that there are also intermediate states between transparent and opaque "to account for the perception of semitransparent surfaces" (sect. 8.1, para. 1). I suggest that Lehar consider filling these semitransparent voxels with "potential illumination" "at multiple depth values at every single (x,y) location." This would also strengthen his second and third conclusions that "volumes of empty space are perceived with the same geometrical fidelity as volumes of solid matter" and that "multiple transparent surfaces can be perceived simultaneously" (sect. 10, points 2 and 3). Having semitransparent voxels contain "potential illumination" is a more parsimonious description of the void between your eyes and this page. You can actualize the "potential illumination" of these voxels by placing your finger in front of any shadow cast on the page. More accurately, Lehar's phenomenological model allows *only* the plane of voxels directly in front of a given surface to contain cast shadows (i.e., less illumination), because the voxels that compose the surface must be the color of the opaque surface itself (sects. 5.1, 8.1).

Note that this concept is not merely peripatetic (Aristotle 1976) or an ether explanation, in that we are always subjectively aware of the illuminant. For example, by looking from your illuminated reading room into a dark hallway, your subjective experience is not only that the hallway walls are under less illumination but also that the space itself contains less light. In this way, "potential illumination" can also address why color constancy differs in two- versus three-dimensional scenes. For example, Gilchrist (1977) had observers look through a pinhole into a room containing a doorway into a second room. The near room was dimly illuminated and the far room was highly illuminated. Attached to the door frame were several papers, arranged so that a mid-gray paper appeared either adjacent to the door frame or (with its corners removed) on the far room's back wall. The lightness of the paper shifted in the direction of lightness constancy depending on whether it appeared on the door frame or on the far wall. Schirillo et al. (1990) generated equivalent stimuli in two dimensions using a stereoscopic cathode ray tube (CRT) and stereoscope, yet this replication produced only a fraction of Gilchrist's constancy. I hypothesize that this occurred because stereoscopic space does not contain the actual voxels of either high (e.g., near room) or low (e.g., far room) illumination. In essence, Schirillo and colleagues failed to preserve Lehar's necessary condition of "volumetric mapping" (target article, Fig. 1D).

The ubiquitous use of CRT images reduces scenes to Alberti's window, which retains perspective cues but eliminates Lehar's requirement that space be volumetric. This obfuscates the color constancy paradox, in that these voxels contain illumination. For example, Adelson's (1993) famous wall-of-blocks illusion contains cubes of identical luminance that appear dissimilar in lightness and concomitantly under illusory transparent stripes. Logvinenko et al. (2002) eliminated both the appearance of transparency and the lightness illusion by constructing a three-dimensional version of Adelson's two-dimensional display. I hypothesize that the visual system does not add a transparent veil to Logvinenko's display because it already ascribes illumination to every voxel in space. However, when Adelson eliminates such volumes but retains the same spatial geometry via X-junctions, the visual system reconstructs the volume to contain regions of illusory transparency (i.e., illumination). Consequently, Lehar's improved spatial model requires a phenomenal description of *empty* space that includes "potentially illuminated" voxels.

If vision is "veridical hallucination," what keeps it veridical?

Peter Ulric Tse

Department of Psychological and Brain Sciences, Dartmouth College, Hanover, NH 03755. **Peter.Tse@dartmouth.edu**
http://www.dartmouth.edu/artsci/psych/faculty/tse.html

Abstract: If perception is constructed, what keeps perception from becoming mere hallucination unlinked to world events? The visual system has evolved two strategies to anchor itself and correct its errors. One involves completing missing information on the basis of knowledge about what most likely exists in the scene. For example, the visual system fills in information only in cases where it might be responsible for the data loss. The other strategy involves exploiting the physical stability of the environment as a reference frame with respect to which the eyes and body can move.

> *[S]pace and time are only forms of sensible intuition, and hence are only conditions of the existence of things as phenomena . . . we can have no cognition of an object, as a thing in itself, but only as an object of sensible intuition, that is, as phenomenon*
> —Immanuel Kant (*Critique of Pure Reason*, 1781)

Lehar develops the Kantian insight that perception is (1) entirely a mental construction; (2) lacks access to the world-in-itself to determine the accuracy of its representations; and (3) is only possible given an internal framework of space-time that permits sensory input to be interpreted as occurring in an external space-time. Here I focus on how the brain can construct true information about the world when there is no way to judge objectively whether that information is true by comparing that information to the world-in-itself.

To create veridical information, the visual system must compensate for errors, data loss, and processing bottlenecks imposed by its imperfect design. It has nothing but the ambiguous, incomplete, and noisy image to determine whether it has made an error. It must therefore know what types of image cues indicate errors and it must have strategies to correct errors. The visual system corrects itself only when it is responsible for errors or data loss. It compensates for its own likely errors using two strategies. One is to rely on world knowledge, and the other is to assume that the world is stable.

For example, when does the visual system fill in missing phenomenal features and when does it merely note that completion takes place without filling-in (see Fig. 1)? Filling-in occurs when the information that is missing from the image is missing because of the visual system's own failure to detect it. The visual system follows the principle "no news isn't necessarily bad news when there was no way to get the news in the first place." The visual system functions as if it knows that it does not always have adequate in-

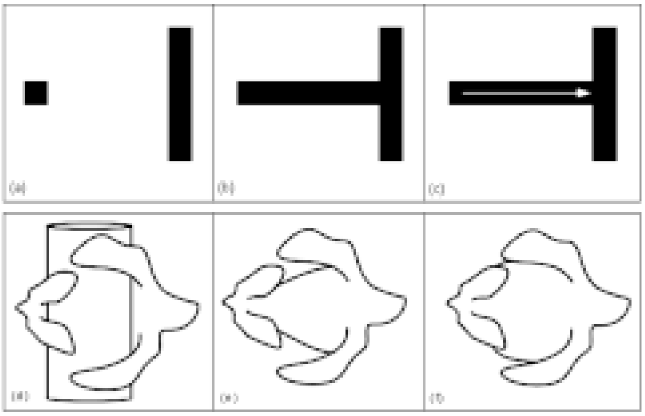

Figure 1 (Tse). When (a) is replaced all at once by (b) a smooth apparent motion (indicated by an arrow) is filled in (Tse & Logothetis 2002). No filling-in occurs in (d) cases of amodal completion (Tse 1999b). The shape behind the occluder, whether (e) or (f), is not completed.

formation in a particular domain to determine the structure of the world. Hence, when information is missing from an image, this is not necessarily regarded as contradictory information or information that the undetected thing does not exist in the world. The information might be undetected because of poor viewing conditions or because of the inherent limits on detection imposed by a noisy perceptual apparatus that has limited sensitivity. A more precise formulation is:

> In the absence of direct (but presence of appropriate indirect) image evidence for the existence of *x*, under viewing conditions where *x* would not be detectable in the image even if it were present in the world, the visual system may not only not reject *x*, it may assume *x* to be the case, and interpolate *x* so that *x* is seen as if it were visible.

Filling-in occurs because the visual system in effect blames itself. The sensitivity of the visual system under given viewing conditions can be too poor to permit detection of an entity that indirect image evidence implies exists. Under such conditions the visual system creates what it "knows" must be missing. In amodal completion there is no filling-in because the visual system does not blame itself. The shape or features of the occluded portions of an object are not filled in, because under no possible viewing conditions would shape or surface features be visible through an opaque occluder. No matter how insensitive the visual system might be, it cannot blame itself for not detecting entities that are in principle undetectable under any viewing conditions.

A second strategy to overcome potential errors is to analyze image data under the assumption that the world is stable. First, the visual system does not need to store detailed information about the world because it can always sample the world for more information (O'Regan 1992). Second, the visual system can stabilize perceptual space by relying on the presumed stability of the world. For example, the retinal image usually changes en masse only when the body or eyes move. The system can exploit this stability in order to maintain the eyes and body in a constant position with respect to the world. A classic demonstration of this is the "moving room" experiment (Lee & Aronson 1974), in which a person stands in a room that is set on rollers. When the walls move, rather than assume that the world has moved, the visual system assumes that the body has moved and corrects for this by changing the body's position. Sometimes subjects even fall over. It is as if the visual system blames itself for the discrepancy caused by the moving room and compensates by relying on a world that it wrongly assumed was stable.

Another example can be found in the recalibration of perceptual space that takes place after a saccadic eye movement. Deubel and colleagues (Deubel et al. 1998) have argued that the system seeks its saccade target immediately after a saccade. If this target is found within a certain spatial and temporal window, the visual system assumes the target object to have remained stable and uses it as a reference object to determine the positions of other objects. This is true even when the target object in fact moves during the saccade. Even more surprisingly, Deubel and colleagues find that if the target has moved to the right, and a neighboring distractor has not moved at all, the visual system creates a percept of a target that has remained stationary and a distractor that has jumped to the left. Because the visual system's initial saccade lands accurately on the position where the target was at the beginning of the saccade, the visual system should know that the target has changed position. But this is true only if it assumes its saccade was infallibly correct. Instead, a corrective saccade is automatically made to the new position of the target, and the object is assumed to have remained stable. The distractor's illusory jump to the left is filled in because it is the motion that must have occurred, assuming the stability of the target and the world. Again, it is as if the visual system blames itself for the discrepancy and relies on the stability of the world to correct its presumed error.

Because the visual system has no direct access to the world, it must rely solely on the image to judge whether it has made errors in specifying the image-to-world correspondence. Error correction is only possible based on assumptions about world structure and statistics. Completion may be phenomenal or not, depending on whether the visual system "blames itself" for the data loss. In addition, the visual system takes a world that it assumes to be stable as its frame of reference. These two strategies allow the visual system to overcome the handicap that the truth of perceptual information cannot be judged by comparing that information with the world-in-itself.

Is the world in the brain, or the brain in the world?

Max Velmans

Department of Psychology, Goldsmiths College, University of London, New Cross, London SE14 6NW, United Kingdom. **m.velmans@gold.ac.uk**
http://www.goldsmiths.ac.uk/departments/psychology/staff/velmans.html

Abstract: Lehar provides useful insights into spatially extended phenomenology that may have major consequences for neuroscience. However, Lehar's biological naturalism leads to counterintuitive conclusions, and he does not give an accurate account of preceding and competing work. This commentary compares Lehar's analysis with that of Velmans, which addresses similar issues but draws opposite conclusions. Lehar argues that the phenomenal world is in the brain and concludes that the physical skull is beyond the phenomenal world. Velmans argues that the brain is in the phenomenal world and concludes that the physical skull is where it seems to be.

Is the phenomenal world in the brain, or is the brain in the phenomenal world? As William James (1904) noted, "the whole philosophy of perception from Democritus's time downwards has been just one long wrangle over the paradox that what is evidently one reality should be in two places at once, both in outer space and in a person's mind." James defended the former view, and consequently developed a form of neutral monism in which the phenomenal world can be regarded as being either "mental" or "physical" depending on one's *interest* in it. If one is interested in how the appearance of the perceived world depends on perceptual processing, one can think of it as mental (as a psychological effect of that processing). If one is interested in how some aspect of the perceived world relates to other aspects of that world (e.g., via causal laws), one can think of it as physical. Lehar, by contrast, defends "biological naturalism" (a form of "physicalism") – the view that the experienced world is literally in the brain.

The difference is fundamental. But whatever view one takes about *where* to locate the perceived world, one fact is clear: the three-dimensional world that we see around our bodies, which we normally think of as the "physical world," is *part of* conscious experience, not *apart from* it. This perceived world is *related* to the unperceived world described by physics (in terms of quantum mechanics, relativity theory, etc.) but it is not identical to it. This is potentially paradigm-shifting for the reason that it redraws the boundaries of consciousness to include the perceived physical world, with consequences for our understanding of mind-body relationships, subjectivity versus objectivity, science, epistemology, and much else (see extensive discussions in Velmans 1990; 1991a; 1991b; 1993; 1996; 2000; 2001; 2002a; 2002b). As Lehar notes in his target article, this conceptual shift also has consequences for neurophysiology. An accurate phenomenology of consciousness is a prerequisite for an adequate understanding of the neural processes that support that phenomenology. In this, Lehar's Gestalt Bubble model provides an interesting, original, and potentially useful step forward.

Given the fundamental nature of the issues and the positive contributions of his paper, it is a pity that Lehar's review of preceding and competing positions is often inaccurate and unnecessarily dismissive. For example, I barely recognised my own work on these problems from his summary. I do not have space to correct these errors here,[1] but merely note that apart from a few crucial differences, Lehar's understanding of the consciousness-brain relationship in visual perception is virtually identical to my own.

What are the crucial differences? Consider the simple model of visual perception shown in Figure 1. Viewed from the perspective of an external observer *E* (not represented within the figure), light rays reflected from an entity in the world (which *E* perceives to be a cat) stimulate *S*'s eye and visual system. Neural representations of the entity, including the neural correlates of consciousness, are produced in *S*'s brain. In terms of what *E* can observe that is the end of the story. However, once the conditions for consciousness form in *S*'s brain, she also *experiences* a cat out in the world – so a full story of what is going on has to combine what *E* observes with what *S* experiences (see discussion of mixed-perspective explanations in Velmans 1996; 2000). If we combine *E*'s observations with those of *S,* an entity in the world (the initiating stimulus) once processed, is consciously experienced to be an entity in the world (a cat), making the entire process "reflexive."

But here's the puzzle: The neural representations of the cat (observed by *E*) are undoubtedly *in S*'s brain, so how can *S* experience the cat to be *outside* her brain? The *effect* is natural and ubiquitous, so there must be a natural explanation. Lehar's Gestalt Bubble model gives some indications of *what* is achieved but does not suggest *how* it is done – and at present, we just do not know. However, both virtual reality and holography might provide useful clues (Pribram 1971; Revonsuo 1995; Velmans 1993; 2000). Suppose, for example, that the information encoded in *S*'s brain is formed into a kind of neural "projection hologram." A projection hologram has the interesting property that the three-dimensional image it encodes is perceived to be out in space, in *front* of its two-dimensional surface, provided that it is viewed from an appropriate (frontal) perspective and it is illuminated by an appropriate (frontal) source of light. If the image is viewed from any other perspective (from the side or from behind), the only information one can detect about the image is in the complex interference patterns encoded on the holographic plate. In analogous fashion, the information in the neural "projection hologram" is displayed *as* a visual, three-dimensional object out in space only when it is viewed from the appropriate, first-person perspective of the perceiving subject. And this happens only when the necessary and sufficient conditions for consciousness are satisfied (when there is "illumination by an appropriate source of light"). When the image is viewed from any other, external perspective, the information in *S*'s "hologram" appears to be nothing more than neural representations in the brain (interference patterns on the plate).

The "projection hologram" is, of course, only an analogy, but it is useful in that it shares some of the apparently puzzling features of conscious experiences. The *information* displayed in the three-dimensional holographic image is encoded in two-dimensional patterns on a plate, but there is no sense in which the three-dimensional image is *itself* "in the plate." Likewise (contra Lehar), I suggest that there is no sense in which the phenomenal cat observed by *S* is "in her head or brain." In fact, the three-dimensional holographic image *does not even exist* (as an image) without an appropriately placed observer and an appropriate source of light. Likewise, the existence of the phenomenal cat requires the participation of *S,* the experiencing agent, and all the conditions required for conscious experience (in her mind/brain) have to be satisfied. Finally, a given holographic image exists only for a given observer and can be said to be located and extended only where that observer *perceives* it to be.[2] *S*'s phenomenal cat is similarly private and subjective. If she perceives it to be out in phenomenal space beyond the body surface, then, from her perspective, it *is* out in phenomenal space beyond the body surface.

But this does not settle the matter. To decide whether the phenomenal cat is really outside *S*'s head, we have to understand the relation of phenomenal space to physical space. Physical space is conceived of in various ways depending on the phenomena under consideration (e.g., as four-dimensional space-time in relativity theory, or as eleven-dimensional space in string theory). However, the physical space under consideration here and in Lehar's analysis is simply *measured space.* Lehar agrees, for example, that at near distances, phenomenal space models measure space quite well, whereas at far distances this correspondence breaks down (the universe is not really a dome around the earth). How do we judge how well phenomenal space corresponds to measured space? We measure the actual distance of an object within phenomenal space, using a standardised measuring instrument – at its simplest, a ruler – and count how often it has to be placed end to end to get to the object. Although rulers look shorter as their distance recedes, we know that their length does not significantly alter, and we conclude therefore that distant objects are really further than they seem.

Lehar and I agree (with Kant) that whether we are "subjects" or "external observers," we do not perceive things as they are in themselves – only phenomena that *represent* things themselves, and, together, such phenomena comprise our personal phenomenal worlds. In Figure 1, for example, the cat, the subject's head, and the neural representations in *S*'s brain (as they appear to *E*) are as much part of *E*'s phenomenal world as the perceived cat is

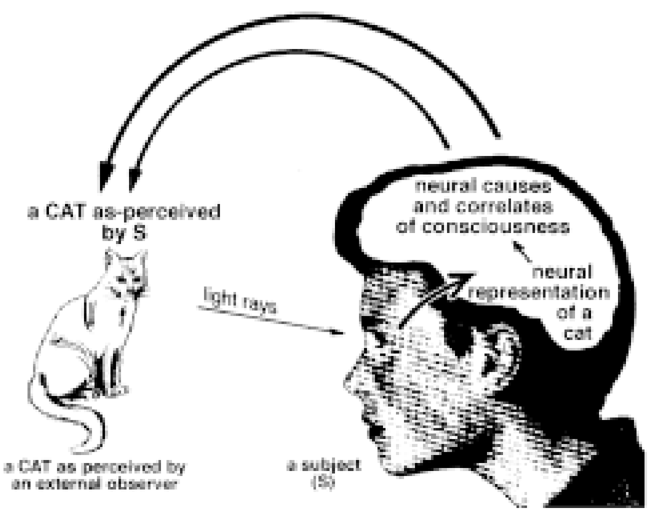

Figure 1 (Velmans).

part of *S*'s phenomenal world. This applies equally to rulers or other instruments that *E* might use to measure distance. In sum, to carry out his science, *E* does not have an observer-free view of what is going on, anymore than *S* does. *E* and *S* simply view what is going on from different third- and first-person perspectives. This has extensive consequences (worked out in Velmans 2000), but I have space to comment on only one of these here. According to Lehar, the three-dimensional phenomenal world in my own analysis is "undetectable externally by scientific means," does not "exist in any true physical sense," and is therefore "a spiritual entity to be believed in (for those who are so inclined), rather than anything knowable by, or demonstrable to, science" (target article, sect. 2.3, para. 8). Nothing could be further from the truth. Data in science consist entirely of observed *phenomena* that occur in a spatially extended phenomenal world, and the measurements that we make within that phenomenal world are the only ones we have on which to ground our science!

Where is this phenomenal world? Viewed from *E*'s perspective, it is outside his head, and the distance of the phenomenal objects within it can be measured, using standardised instruments that operate on phenomenal space (the distance of this phenomenal page from your eye, for example, can be measured with a ruler). Viewed from *E*'s perspective, the phenomenal world also appears to be represented (in a neural form) in *S*'s brain. Viewed from *S*'s perspective, things look the same: The phenomenal world appears to be outside her head, and, if she looks, a neural representation of that world appears to be encoded in *E*'s brain. Given that the evidence remains the same, irrespective of the perspective from which it is viewed, one can safely conclude (with James) that although a neural encoding of the world is within the brain, the phenomenal world is outside the brain. As this is how the natural world is formed, there must be a natural explanation (see above). I have shown (Velmans 2000) how this analysis can be developed into a broad "reflexive monism" that is consistent with science and with common sense.

Now consider Lehar's alternative: It is widely accepted that experiences cannot be seen in brains viewed from the outside, but Lehar insists that they can. Indeed, he insists that *E* knows more about *S*'s experience than *S* does, and *S* knows more about *E*'s experience than *E* does, as the phenomenal world that *S* experiences outside her brain, is nothing more than the neural representation *E* can see inside her brain, and vice versa. This has the consequence that the *real* physical skull (as opposed to the phenomenal skull) exists *beyond* the phenomenal world. As Lehar notes, the former and the latter are logically equivalent.

Think about it. Stick your hands on your head. Is that the real physical skull that you feel or is that just a phenomenal skull inside your brain? If the phenomenal world "reflexively" models the physical world quite well at short distances (as I suggest), it is the real skull, and its physical location and extension are more or less where they seem to be. If we live in an inside-out world, as Lehar suggests, the skull that we feel outside our brain is actually inside our brain, and the real skull is outside the farthest reaches of the phenomenal world, beyond the dome of the sky. If so, we suffer from a mass delusion. Our real skulls are bigger than the experienced universe. Lehar admits that this possibility is "incredible." I think it is absurd.

NOTES

1. Details are given in an unabridged version of this commentary (Velmans 2003).

2. The position of the image relative to the plate, for example, changes slightly as the observer moves around the plate. Nevertheless, the image is sufficiently clear for the observer to (roughly) measure its width and how far it projects in front of the plate (e.g., with a ruler).

Percepts are selected from nonconceptual sensory fields

Edmond Wright

3 Boathouse Court, Trafalgar Road, Cambridge CB4 1DU, United Kingdom.
elw33@hermes.cam.ac.uk http://www.cus.cam.ac.uk/~elw33

Abstract: Steven Lehar allows too much to his direct realist opponent in using the word "subjective" of the sensory field per se. The latter retains its nonconceptual, nonmental nature even when explored by perceptual judgement. He also needs to stress the evolutionary value of perceptual differences between person and person, a move that enables one to undermine the direct realist's superstitious certainty about the singular object.

With regard to the title of Steven Lehar's article, it is vital that the term "subjective" not be used of the structurally isomorphic sensory field. To acquiesce in its use is to yield ground to the direct realist opponent. It can be credibly argued that the ground of all sensory experience is thoroughly nonconceptual, beyond that of Gareth Evans's use of the term, in that recognisable entities and properties are *not* given in the initial stage of the process, not even that of a subject (Evans still took "nonconceptual" to include the perception of separable objects-as-unrecognised; Evans 1982, p. 228). The isomorphic field, because of its very isomorphism, is as brute as the input at the sensory organ, therefore as *non*mental as material (Wright 1996), whatever its nature may be as an emergence from complexity. How could it not be if it is, however indirectly, covariant with the input? As John Foster has put it, sensing is something that just "happens to us" rather than "something that we do" (Foster 2000, p. 123). Subjectivity does not enter into the equation until the establishment of perceptual judgement and memory has taken place upon that nonconceptual evidence at the behest of the motivational module. Therefore, it is going too far to attribute "protoconsciousness" at this level (sect. 6.5), for this correction regards sensing as always existing apart from judgement, merely evidence upon which a mind may or may not work (Wright 1996, pp. 24–28).

Lehar (sect. 2.4) justifiably uses the analogy with the television screen employed by Roy Wood Sellars, Barry Maund, and Virgil C. Aldrich (Sellars 1916, p. 237; Maund 1975, pp. 47–48; Aldrich 1979, p. 37), in that the distinction made between the screen-state (of the phosphor cells) and what is judged to be shown upon it is structurally similar to that between the sensory evidence within the brain and the percepts chosen from it. If he accepts the cogency of this comparison, then he ought to acknowledge that the radically nonconceptual nature of the sensory evidence is implied by this analogy. However much information-theoretic evidence there may be on screen/neural raster, it registers only covariations with light-wave frequencies and intensities at the camera/retina, not any information about recognisable entities and properties (if the TV set was upside down and one had just entered the room where it was, one would be unable to use one's memories to judge that, say, Ian McKellen as Gandalf was at that moment "visible," the screen thus revealing its permanently nonconceptual state). So Lehar should accept the criticism made above.

Those anti-qualia philosophers and psychologists who inveigh against the "picture-in-the-head" proposal (e.g., O'Regan & Noë 2002), have always opposed the television analogy. Lehar does not sufficiently defend himself against this attack (sect. 2.3). As I have pointed out (Wright 1990, pp. 8–11), there cannot literally be pictures in the head, for, if colours are neural events, actual pictures are *not coloured,* and the "picture" in the head is. Nor is an eye required for sensing neural colour, for eyes are equipped to take in *uncoloured* light-waves, and there are no light-waves in the head. Visual sensing is a direct experience for which eyes would be useless. Gilbert Ryle's attempt to maintain that one would have to have another sensation to sense a sensation remains as an argument, as Ayer described it, "very weak" (Ryle 1949/1966, p. 203; Ayer 1957, p. 107).

Once this radically nonconceptual nature of the fields is admitted, its *evolutionary* value can be brought out, which is precisely what Roy Wood Sellars and Durant Drake – the very philosophers that Lehar calls to his aid – insisted upon (target article, sect. 2.3; Drake 1925; Sellars 1922). Sellars particularly stressed the feedback nature of the perceptual engagement, which allows for the continual updating of entity selection from the fields (altering spatiotemporal boundaries, qualitative criteria, etc.), a claim that renders stances such as Gibson's which take the object as given (amusingly termed "afforded"; Gibson 1977), not so much as "spiritual," the term favoured by Lehar (sect. 2.3), but as literally *superstitious.*

What weakens the direct realist case is its unthinking reliance on the pre-existing *singularity* of "external things." If the feedback argument of Sellars *père* is correct (Sellars 1970, p. 125), then the perfectly *singular* "object" or "entity" is but a feature of the mode of perceiving and not ontological in its nature. The behaving *as if* it is singular, the trusting *assumption* that it is, is a necessary feature of the intersubjective cooperation, for we could not even roughly coordinate our differing percepts unless we did project a strictly imaginary perfectly common focus of them; but it is fatal to take the convergence as without residue, for that would cancel the possibility of feedback and hence of mutual correction.

Lehar adverts to the uncertainty of the object (sect. 6.1). The only basic ontology required under the theory above is of the material continuum: When human social perceiving is in operation, with its incessant *intersubjective* correction in action, then a very modest ontological further claim can be made, namely, that a community of correctors exists, and hence of selves and their sensory fields, but *not as fixed entities,* only as current tentative selections from sensory and motivational experience. The direct realist, by contrast, is committed to an indefinite number of separable singular entities (objects and persons), a superstition that is disconcertingly all too common in recent books on the philosophy of perception (from Millar 1991 to Thau 2002; there are very few exceptions, e.g., Maund 1995). The act of faith in singularity which is necessary to bring our differing percepts into some kind of working overlap, is taken by the direct realist as actual, which represents an insidious and dangerous move to the conviction that his own percept is the standard for all.

Author's Response

Alternative paradigmatic hypotheses cannot be fairly evaluated from within one's own paradigmatic assumptions

Steven Lehar

Peli Lab, Schepens Eye Research Institute, Boston, MA 02114-2500.
slehar@cns.bu.edu http://cns-alumni.bu.edu/~slehar

Abstract: To avoid endless and futile debate, critics of an alternative paradigmatic hypothesis cannot simply state their own paradigmatic assumptions as if they were plain fact while dismissing those of the opposition as self-evidently absurd, because it is exactly those initial assumptions that are brought into question by the paradigmatic proposal. Perceived incredibility is no valid ground for rejection of a paradigm whose alternatives are at least equally incredible, and arguably more so.

The energetic responses of the open peer commentaries indicate that the target article has touched a raw nerve; this is perhaps a harbinger of an interesting direction of investigation. The epistemological issue at the core of the debate is a paradigmatic question that challenges some of the foundational assumptions of psychology and neuroscience, which have remained so long unchallenged that they are generally held to be established fact. As is frequently the case in paradigm debates, the opposing camps often cite the selfsame evidence to support their opposite conclusions, because they are arguing from different foundational assumptions. To avoid endless debate, it is therefore essential for commentators to recognize the paradigmatic issue at the core of the debate, and not just state their own paradigmatic assumptions as if they were established fact – while dismissing those of the opposition as self-evidently absurd – because it is exactly those initial assumptions that are brought into question by the target article. If alternative paradigms are to be fairly evaluated, it is necessary to temporarily and provisionally suspend one's own paradigmatic assumptions, and accept the assumptions of the alternative paradigm as if they could actually be true. Only then can the competing paradigms be fairly compared, not on the basis of the perceived incredibility of their initial assumptions, but on the overall coherence and self-consistency of the world view they implicate in total.

R1. Rigor paradigmatis

Unfortunately, many of the commentators failed to grasp the paradigmatic nature of the proposal and restated their own paradigmatic assumptions as if they were plain fact, thus committing the error of *petitio principii,* assuming from the outset that which is to be proven.

Booth complains that it is "foolish" to look for consciousness among the brain cells. I contend that it is foolish to look for it anywhere else but in the brain! As in most paradigmatic debates, one man's "foolish" is another man's "obvious." But Booth says not a word about the epistemological difficulties, which were discussed at length in the target article, of the view that he defends. If the experience of a red surface, for example, is located anywhere else but in the brain, then it is a spatial structure that exists, but it does not exist in any space known to science. This makes Booth's hypothesis a religious or spiritual theory, because the experienced surface is in principle beyond detection by scientific means, and therefore it is a theory that is impossible to disprove. It's no good trying to dismiss the structure of consciousness in a trick of grammar, as Booth proposes, by claiming that the spatial structure of experience is a "seeming" rather than something real. That objection was addressed in the target article with the observation that visual consciousness has an *information content,* and information cannot exist independent of an actual physical mechanism or substrate in which it is registered. Booth seems to think that simply stating his own paradigmatic hypothesis as if it were plain fact ("We are not looking at a world inside our minds; we are . . . seeing the colour of the patch out there.") is an adequate response to the hypothesis that what we are seeing really is in our brain.

Dresp complains that I fail to make clear the link between the Gestalt Bubble model and general theories of consciousness.

> What the model has to do with consciousness . . . remains totally unclear. Neither the fact that we are able to consciously experience and describe three-dimensional shapes as entities and wholes, nor the fact that we can find laws or codes describing how these emerge perceptually, implies or proves that con-

https://doi.org/10.1017/S0140525X03410091

sciousness is necessary to see and move around in three-dimensional space.

The link to general theories of consciousness is through the philosophy of *identity theory* (Feigl 1958; Russell 1927), as explained in the target article, whereby mind and brain are not separate and distinct but are ontologically one and the same. The presence of spatial structures in the perceptual representation is identically equal to a conscious experience of those structures.

Dresp simply assumes as if it were plain fact, that the experiential component of consciousness is separate and distinct from the physical mechanism by which it is instantiated, and that, therefore, a model of the mechanism cannot possibly be a model of the experience, because the experience cannot be modeled in principle. But it is at least equally likely prima facie that experience is *not* separate and distinct from the mechanism that carries it, but that experience is a physical process taking place in the physical brain, so a model of the mechanism would automatically also be a model of the experience. In fact, this is *by far* the more parsimonious explanation because it employs a single *explanans,* the brain, to account for the properties of mind and brain. Identity theory is an equally valid paradigmatic alternative that cannot be dismissed without demonstrating why it is less credible than a mystical, nonphysical theory of experience beyond science. Furthermore, the Gestalt Bubble model is explicitly defined as a model of experience, rather than its neurophysiological correlate, so it is inconceivable how Dresp fails to see the connection between a model of experience and the experience it models.

Dresp objects to my exhortation to discover the real truth behind visual processing. "Who said that science has to bother with metaphors such as 'truth'?" she asks. But science is all about modeling *objective external reality,* a truth that science presupposes to exist. Either there are "pictures in the head," as explicit spatial structures, or there are not. And whether there are pictures in the brain is of primary importance for psychology, philosophy, and neuroscience. Simply defining those pictures as a mystical nonphysical entity brings us no closer to understanding how consciousness arises in brains.

Duch asks: "How can the physical skull encompass nonphysical, inner world? 'The world inside the head' is a metaphor, and it does not make much sense to invert it, unless one believes that there is some kind of physical world squeezed inside the skull." That is precisely the hypothesis I presented, although it is a perceptual world, expressed in physical form, which is squeezed inside the skull. This notion seems to Duch so absurd from the outset that he cannot believe that is what I am proposing. Duch states, as if it was plain fact, that "interpretation of the spatial structure of the states of the visual system has nothing to do with their physical location." This is exactly the issue brought into contention by the target article. If he is to contest this assumption, he must explain which incredible alternative he wishes to defend in its place. Does he claim that consciousness does not exist at all, as Dennett (1981) would have us believe? Or does he allow that it exists but in some hidden dimension inaccessible to scientific scrutiny, as **Velmans** proposes? If so, how does he address my critique that his view of consciousness is a religious rather than a scientific hypothesis?

In his conclusion, **Duch** says "It is doubtful that the Gestalt Bubble model can explain observations that have not been hidden in its premises." That is the nature of paradigmatic hypotheses, and is just as true of the hypothesis that Duch defends. If you begin from the outset with the assumption that consciousness has no location, then you are guaranteed never to find it located anywhere!

Fox begins his commentary with the complaint that "[Lehar] ignores much of what is known in perceptual and brain science." The truth is that Lehar *challenges* much of what is "known" in perceptual science. Far from ignoring, I have taken pains to point out the errors of what is known in perceptual and brain science. If Fox begins with the assumption that those supposed "facts" are indisputable, then he is bound to have trouble coming to grips with a hypothesis that they may perhaps be mistaken.

There is much in **Fox**'s commentary that is deeply mysterious. **Fox** accuses me of maintaining the "Cartesian mind-body distinction." But the central hypothesis of the target article is an identity theory whereby *mind* is nothing other than the functioning of the physical brain. This monistic view is diametrically opposed to Cartesian dualism. Furthermore, in section 2.3, I explicitly *refute* Cartesian dualism as a spiritual rather than a scientific hypothesis.

Laming also fails to grasp the paradigmatic nature of the proposed model, presenting as counterarguments the axioms of his own paradigmatic alternative as if they were plain facts, rather than unsubstantiated initial assumptions. Laming insists that some parts of visual experience can be shared with others while the remainder are private, and therefore there cannot be a natural science of perception. So a psychophysical report that, for example, a subject can or cannot see an extended red object, is a valid subjective report. But the redness and spatial extendedness of that perceived object are not validly reportable because they are private. Curiously, the very aspects of experience that Laming considers illegitimate, including all of the Gestalt properties surveyed in sections 5 and 7 of the target article, are exactly the aspects of experience that reveal the spatial structure that Laming insists have no physical reality. Laming must explain *why* the spatial aspect of perception is so private that it cannot be reported, when the Gestalt Bubble model clearly demonstrates how the spatial aspects of perception *can* be reported and quantified in a spatial model. If he contests my phenomenology and claims not to see the sky as a dome, and the sides of a road converge to a point, then he should tell us what he sees instead. It is a paradigmatic choice of Laming's, not a statement of incontestable fact, to call the spatial extendedness of perception unreportable. And if Laming chooses to believe that phenomenal consciousness is not a physical entity in the brain, he should address the clear objections to that paradigm outlined in the target article. In particular, conscious experience, according to Laming, is a spatial structure; it is a structure that exists, and yet it does not exist in any space known to science, and it is in principle undetectable by scientific means. This is a religious or spiritual hypothesis because it is impossible in principle to disprove. To accept Laming's view of consciousness is to declare *consciousness* in principle forever beyond explanation in scientific terms – which would then become a self-fulfilling prophecy.

R2. Perceived incredibility

Some commentators reject the representationalist thesis because they find it to be frankly incredible. **Velmans** complains

> Stick your hands on your head. Is that the real physical skull that you feel or is that just a phenomenal skull inside your brain? . . . If we live in an inside-out world, as Lehar suggests, the skull that we feel outside our brain is actually inside our brain, and the real skull is outside the farthest reaches of the phenomenal world, beyond the dome of the sky . . . Our real skulls are bigger than the experienced universe. Lehar admits that this possibility is "incredible." I think it is absurd.

And with this, I believe, Velmans touches on the principal reason this alternative has been given so little consideration.

I am viscerally sympathetic with this objection, so much so that for years I too refused to accept the conclusion toward which all of the evidence points. It is indeed incredible to think that your physical head is larger than the dome of the sky. But science has discovered many things which were initially considered to be at least equally incredible; like the vastness of the universe, and its cataclysmic genesis from a singularity in space and time, and the bizarre nature of black holes and of quantum phenomena. All of these theories were initially held to be incredible but have since been accepted as established fact. And the reason they were accepted is not that they became any less incredible. Scientific fact is accepted on the basis of the evidence, regardless of the incredible truth to which that evidence points. In fact, that is exactly what gives science the power to discover unexpected or incredible truth: When the obvious explanation is blocked by chronic paradoxes, it is time to give the seemingly incredible alternative a serious look.

In his conclusion, **McLoughlin** invokes Occam's razor against the notion of the world of experience being a picture inside the head. I believe McLoughlin reflects the consensus view in neuroscience: that the hypothesis seems frankly too incredible to deserve serious consideration. But before deploying Occam's razor we must first balance the scales and take a full accounting of the alternatives under consideration. For the alternative is that experience is a spatial structure; it is a structure that exists, but it exists in a separate space that is inaccessible to scientific scrutiny: It is a structure with a vast information content, but with neither mass nor energy nor spatial presence in the physical universe known to science, and the brain conducts a continuous two-way exchange of information with this phenomenal, semiexistent nothing. Alternatively, McLoughlin might prefer Dennett's (1981) eliminative alternative, that conscious experience, the spatial structure under discussion, simply does not exist – and that's the end of the problem of consciousness. If McLoughlin finds the idea of the world-in-your-head incredible, he must balance his rejection of it by telling us which of the other two incredible hypotheses he finds more credible, otherwise it is *his* bubble that bursts, not mine.

R3. Paradigmatic alternatives

Not every challenge to alternative paradigmatic hypotheses is a case of rigor paradigmatis. It is perfectly valid to challenge a paradigm based on the overall coherence or self-consistency of the world view that it implicates in total. In fact that is the only way that paradigms can be fairly compared, as seen in the commentaries discussed in this section.

Revonsuo accepts the representationalist view of the target article but challenges panexperientialism. Revonsuo cites phenomena, such as neglect and blindsight, that suggest perceptual information can be processed without necessarily entering consciousness. But, as in the case of most paradigm debates, both camps can usually muster an explanation for almost any phenomenon raised, although each explanation is consistent only within its own paradigm and sounds patently absurd from the point of view of the other.

For example, *blindsight,* the apparently unconscious processing of visual information, can be explained as a case of amodal perception. When the blindsight patient reports a vague suspicion of motion in the absence of an experience of a moving object, he is reporting a conscious experience of a vague suspicion of motion without actually seeing anything in motion. Even people with normal vision commonly have such experiences in peripheral vision, and many psychophysical experiments measure vague perceptual experiences at the threshold of detection which are not so much *seen* as suspected to have possibly appeared. Many philosophers deny that amodal percepts, or other forms of non-sensory knowledge, can be validly considered as conscious processes, insisting instead that only modal experience is experience. But if we exclude the amodal component of perception from conscious experience, then by definition, amodal perception is always blindsight. It seems more parsimonious to suppose that amodal perception is consciously experienced, even if only amodally, than to suppose that something experienced can be unconscious.

Other examples of apparently unconscious processing can be explained in the panexperiential view as separate, parallel, and largely independent conscious processes in the brain. The part of the brain that performs the processing is indeed conscious of its own performance, but it is not in touch with the part of the brain that reports on that processing, so no processing is reported. Similarly, unconscious processing in dichotic listening can be explained by separate, parallel streams of consciousness, one of which overwrites the other, which is therefore never recalled. This is not to say the evidence of neglect and blindsight favors panexperientialism, merely that it does not refute it. The motivation for panexperientialism lies elsewhere.

It is true, as **Revonsuo** suggests, that the protoconsciousness hypothesis probably cannot make testable predictions, but that is not why I invoked it in the target article. It was raised to plug up some otherwise serious holes in a purely monistic or physicalist explanation of conscious experience. This paradigmatic choice avoids a most subtle residual dualism hidden in Revonsuo's alternative.

As long as a sharp step or abrupt discontinuity is allowed between conscious and unconscious processes, there will always be an explanatory gap, because physically, the brain *can* be disassembled into ever-smaller pieces, all the way to atoms and molecules; whereas in **Revonsuo**'s view, consciousness does not have this ability to be disassembled but disappears abruptly as soon as the minimal conditions for it are no longer met. To be clear, I do not dispute that consciousness may exhibit, and indeed appears to exhibit, an abrupt cut-off – for example, when falling asleep or waking up – although intermediate semiconscious states are also known. If consciousness appears abruptly at some level of organization, then something else physically observable must come into existence at that point also.

For example, consciousness might be identified with a holistic process such as spatial standing waves of electro-

https://doi.org/10.1017/S0140525X03410091

chemical resonance in the brain (Lehar 2003). Resonance shares with consciousness the property of coming abruptly into being when the conditions for its formation are just right, as when blowing a note in a musical instrument. And yet, standing waves are not some supervenient spiritual entity but a real, physically measurable phenomenon that emerges holistically in a physical system. But if standing waves were indeed the physical substrate of consciousness, then that would suggest that musical instruments also necessarily possess some form of primal spatial consciousness. And when global consciousness breaks down in the brain, whether as a result of sleep or of anesthesia, the global synchrony does not disappear so much as it breaks up into a million fragments of locally isolated coherence. Would these fragments not each experience an isolated fragmentary consciousness? If not, then we have again an abrupt discontinuity that suggests a dualism between experience and its physical correlate. The more likely alternative is that parallel fragmented states of consciousness are indeed experienced during periods of unconsciousness but they never register in memory and are therefore quickly forgotten, as are many dreams.

A further problem with the abrupt discontinuity of consciousness is that it opens the problem of the "bridge locus" in the brain; that is, the question of why some very special parts of the brain become conscious, while other parts do not. It also leads to problems with partial or fragmented consciousness, as in split-brain patients, and in cases of multiple personality syndrome, and hypnotic or trance states, which all seem to indicate multiple parallel or alternating consciousnesses in a single brain. Whether the disassembled pieces of consciousness can be usefully considered conscious in any real sense is admittedly a semantic question. But what is *not* semantic is the question of whether consciousness can be disassembled into component pieces, as the matter and energy in the brain can be, or whether the mind operates by different laws than the matter that is sometimes its physical correlate. The only self-consistent physicalist explanation is that complex consciousness in a complex brain occurs by the same principle as simple consciousness in simpler brains; and the same argument propagates all the way down to the root of the phylogenic tree and beyond. Consciousness is what it feels like for matter to exist, and complex human consciousness is what it feels like to be the waves of energy resonating in a human brain.

Revonsuo complains that the panexperientialist position brings us no closer to explaining the radical empirical differences that we want to understand. Quite to the contrary: Until we bring consciousness fully into the realm of the physical world, one small corner of it will remain permanently trapped in a supervenient dimension forever, in principle, beyond scientific scrutiny. That is the modern "ether" theory that must be shown to be pure vacuum.

Gunderson observes that visual experience consists of more than just a spatial structure – it is a spatial structure that is experienced as being viewed from a particular point – and that this aspect of viewing from a point is not captured in the Gestalt Bubble model. In the first place, whether experience is viewed from a point or not, the fact remains: Visual experience is a spatial structure, and the spatial structure of experience is captured in the Gestalt Bubble model.

But in fact, the Gestalt Bubble model goes further, suggesting that the experience as if viewing from a point is itself an illusion. Once we recognize the world of experience for the internal model that it is, it becomes evident that our objective noumenal "self" is not the body-image homunculus observed at the center of our phenomenal world, but that in fact, the whole world of experience is part of our real self. The blue of the sky is not observed "from" the egocentric point, but it is experienced to exist out there where it lies at the surface of our perceptual sphere. The blueness of that azure dome is experienced to exist at a location relative to the egocentric point, but it is in no sense transposed or projected back to the egocentric point. In fact, phenomenally speaking, there is nothing special at all at the location of the egocentric point, which is experienced as an empty void of phenomenal space inside the phenomenal head, just like any other empty space in the phenomenal world.

Three factors contribute to the vivid illusion of viewing the world from a point. The first is the body-image homunculus, which we take to be our real "self" because that particular piece of the phenomenal world is under our direct volitional control. Under large doses of hallucinogenic drugs such as LSD the perceptual distinction between self and nonself tends to disappear, as the body image merges with the background, leaving the entire sphere of perceptual experience to be identified as "self," a common theme also of Buddhist phenomenology. The second contributory factor is the warped geometry of phenomenal space that is organized around a center, the point of highest perceptual resolution, marking that center of symmetry as a special location in phenomenal space. Finally, the illusion is bolstered by perceived surfaces being perceived modally only when they are exposed to the egocentric point, as if they were indeed being viewed from or by the egocentric point. Objects not exposed to the egocentric point are invisible to direct modal experience, and are therefore experienced in amodal fashion.

A similar phenomenon is observed on a radar scope, where radar "echoes" are registered only from those surfaces exposed to the central radar dish, for example, from the exposed front faces of nearby mountains. No echoes are registered from the hidden rear face of the mountains, nor from more remote surfaces occluded by the nearer mountains. As in perception, the center of a radar scope is not the "observer" of the rest of the image on the scope, and appears to be special only because the image on the scope is a veridical manifestation of the external, "nouminal" situation, where the radar echoes are indeed received or "viewed" from the location of the radar dish – which is why no echoes are received from surfaces that are not exposed to that point. Similarly, the phenomenal experience of viewing from a point is a veridical manifestation of the external nouminal situation where physical light from the external world is indeed received by the nouminal eye, no light being received from hidden or occluded surfaces, creating the illusion that the phenomenal world is being viewed from the location of the phenomenal eye.

R4. Reliability of phenomenology

A number of commentators challenged the reliability of phenomenological observation. **McLoughlin** pointed out that naïve observers are surprised to discover that they have a fovea and amazed that they have a blind spot in each eye. True enough, but the same naïve observers can be easily

https://doi.org/10.1017/S0140525X03410091

educated by the most convincing demonstration of all, phenomenological observation of their own loss of resolution in peripheral vision and of their own blind spots. Like any tool, phenomenology is useful only if employed with intelligence.

Booth complains that I commit the epistemological fallacy of trying to build public knowledge on the basis of private impressions. Booth objects that phenomenological observations are private, and so "they cannot be wrong – but then neither can they be right." But then he objects that my observations on phenomenal perspective are wrong! If phenomenological observation cannot be wrong, then how can Booth claim that my observations of phenomenal perspective are wrong? Booth complains that it is impossible to look one way down a road and at the very same moment be looking the other way, So it cannot be correct to write that the two sides of the road must be bowed, as in Figure 2 of the target article. In the first place, it is not necessary to be looking in opposite directions at the same time to see the curvature of the phenomenal world, all one needs to do is to look in one direction and observe that the parallel sides of a road meet at a point at a distance which is less than infinite, and that those parallel sides appear straight and parallel throughout their length. And in the opposite direction one sees exactly the same thing, and in between one sees a spatial continuum exactly as depicted in Figure 2. Phenomenological observation can indeed be right and it can be wrong, and Booth's phenomenology is just plain wrong if he can't see perspective foreshortening in the three-dimensional world around him!

Hochberg suggests an alternative, less holistic model – a stage set rather than an all-encompassing bubble, with an abrupt discontinuity at a certain depth, where a proximal percept of a full three-dimensional road with perfectly parallel sides changes abruptly to a flat two-dimensional experience at right angles to the view direction, in which the sides of the road converge to a point in the plane of the backdrop. No matter how hard I try, I cannot see the world this way; I always see the two extremes of near and far perception seamlessly connected through a continuous intermediate zone, wherein the sides of the road are perceived in full three dimensions, and yet they are also perceived to converge, and they are perceived to be parallel even as they converge.

I acknowledge that perception is somewhat more fragmented than the Gestalt Bubble model suggests. For example, every visual saccade presents a momentary experience perhaps somewhat like a stage set. But the most salient and immediate aspect of conscious experience is the way these individual theater sets are welded together into a unified sphere of spatial experience. Whatever direction we gaze, we are constantly aware of where that gaze is directed in the global sphere of surrounding space, and the objects perceived in that direction are perceived to be located in that part of global space. The experience is more that of a stable, structured surrounding space than a series of theater sets showing successively on the same stage.

As evidence to disqualify the Gestalt Bubble model, **Hochberg** cites visual illusions that vary as a function of where they are attended, because the same Gestalt is in view wherever it is attended. There are two aspects of spatial experience that must be carefully distinguished, we might call them *global* and *focal.* In the global experience our view is of a perfectly stable surrounding world, as suggested in the target article's Figure 2, whose entire surface is painted in modal colors, because whichever direction we look, that is the way it appears. The other aspect of experience is focal, the immediate experience of looking in a particular direction. The world appears at higher resolution in the direction of gaze than in the periphery, and the rear hemisphere behind our head is blank, as suggested in Figure 15. Both global and focal aspects are observed in our experience, so they should both be represented in a model of that experience. The combined experience is modal and focal in the direction of sight, but amodal in the hidden rear portion of the field, as indicated by the dashed lines in Figure 15, although successive saccades in different directions create an illusion of the complete modal sphere suggested in Figure 2.

The point of the Gestalt Bubble model is not to deny that there are localized focal processes active in perception, but merely to add that there is also a single, globally unified perceptual experience, and the localized focal experiences are perceived to be embedded at specific locations in this larger global framework of spatial experience.

Hochberg also cites examples of ambiguous or unstable percepts attributable to figures such as Adelson's (Adelson 2000) Impossible Staircase (Figure 2C of Hochberg's commentary). Hochberg argues that since these percepts are observed to be unstable and/or ambiguous, they would disqualify a Gestalt Bubble model of a globally unified perceptual world. But Hochberg need not have gone farther than Figures 3, 5, 6C, 11D, 12A, and 16A (porthole illusion variant) of the target article for examples of unstable, semistable, or ambiguous figures. The perceptual tendency towards a unified, globally consistent percept is a goal that the perceptual mechanism seeks, but does not always achieve, so unstable and multistable percepts are not counterexamples to the principle of emergence as described in the Gestalt Bubble model; they provide a more detailed look at the mechanism of that emergence.

Hoffman argues that the perceived world of the Gestalt Bubble model is not a veridical replica of the external world, but merely a useful "user interface" to the external world, with no more need to resemble that world than a Windows interface needs to resemble the diodes, resistors, and software of a computer. Therefore, there may be no real resemblance at all between the structure of our phenomenal world and the real external world that it represents. Hoffman proposes to replace indirect realism with this species-specific "user-interface" theory of perception. But that is exactly what I have proposed in the Gestalt Bubble model. Nowhere was it stated that the phenomenal world is in any sense identical to the external world. Phenomenal colors are very much more impoverished than the chromaticity of physical light, and phenomenal perspective shrinks the infinite external world into a finite bounded bubble. These are clearly species-specific "user-interface" simplifications of external reality. The "realism" in "indirect realism" is already modulated by the word "indirect," that is, the phenomenal world is a very real and direct view of processes taking place within our own brain, and those processes in turn represent indirectly the structures and surfaces presumed to be present in the more remote external world.

Immanuel Kant (1781/1991) anticipated **Hoffman**'s observation that the phenomenal world need not show any resemblance to the external physical world. We do not even

https://doi.org/10.1017/S0140525X03410091

know if it has three spatial dimensions and time; all we know is that those are the dimensions of the internal phenomenal world. It should also be noted, however, that in every other realm of human activity, from hunting and gathering, to business and finance, to politics and engineering, the assumption that the world of experience is an accurate representation of objective reality is so successful that the indirectness of perception can be readily ignored. And that in turn suggests that experience must accurately reflect some essential aspects of the external, although **Hoffman** is right that we cannot determine phenomenologically which aspects of the world are veridically replicated and which are not.

Lloyd presents a very clever argument by analogy that appears to punch a hole through the central premise of isomorphism. Consider the statement from the target article: "The fact that the world around us appears as a volumetric spatial structure is direct and concrete evidence for a spatial representation in the brain" (sect. 5.2, para. 6). Lloyd suggests that the absurdity of this statement can be revealed by substituting "colored" for "spatial" in this passage. Phenomenal color experience defines a three-dimensional relational structure of phenomenal color space. But, Lloyd correctly objects, the fact that we experience phenomenal color does not mean that the color solid appears anywhere in our brain, and by the same token, spatial experience does not imply a spatial structure in the brain.

This objection is already addressed in the target article by the specification of *information content* as the one quantity which is necessarily preserved across the mind/brain barrier. A color experience is indeed a three-dimensional relational structure, not so different in principle from the color phosphor dots on a television monitor. It takes three values of three phosphor dots to encode a single point of color on your screen, and those three values define a single point in Red-Green-Blue [RGB] space. However, a single point on a television screen does not define an entire color solid, but merely one point in that three-dimensional color space, representing the color currently represented at that point. The color solid is not explicitly present anywhere, but the relational structure that it encodes is implicitly present in the range of possible values of the three phosphor dots.

Spatial perception is different from color perception in this one significant aspect, that every point in perceived space can be perceived with a distinct color. That means that there are as many separate and potentially distinct color values in a perceived surface as there are resolvable points across that surface. The points in a perceived surface are perceived simultaneously and in parallel, and together they define a relational structure in which every point bears a specific spatial relation to every other point in that perceived surface. This is quite different from the implicit structure of color space that encodes only one color at a time, because spatial perception encodes a whole spatial array of color values, all of which are simultaneously present in experience.

Regarding the value and prospects for phenomenology, **Marković** says that "without the precise specification of the extraphenomenological aspects of perception, such as the stimulus and neural domains, it is difficult to answer the question related to why the percept looks as it does." This, however, is difficult only if one employs phenomenology merely to confirm theories of vision based on neurophysiology. Once we realize that what we are seeing in experience *is* the representation in our own brain, there is a *great deal* that can be learned about why things look the way they do, and how things are represented in the brain. Marković is right that scientific explanation must go beyond mere description. In Marković's example, the Earth's motion becomes comprehensible only when considering the influence of the Sun. But before science can propose explanations it must begin with description. The influence of the Sun on the Earth's motion would have never become clear had we not first observed and described that motion. Psychology too must begin with a description of experience before it can attempt a plausible explanation for it.

R5. Explicit volumetric representation

McLoughlin points out that a volumetric space can be expressed in a sparse, more symbolic code, without recourse to an explicit spatial array, with objects represented as tokens, with x, y, and z, location, and so forth. There are many aspects of mental function, such as verbal and logical thought, that are clearly experienced in this abstract manner. But visual consciousness has an information content, and that content is equal to the information of a volumetric scene in an explicit volumetric representation. Every point in the volume of perceived space is experienced simultaneously and in parallel. To propose that the representation underlying that experience is a sparse symbolic code is to say that the information content of our phenomenal experience is greater than that explicitly expressed in the neurophysiological mechanism of our brains.

Velmans' holographic analogy is very apt. There is indeed no "picture" as such on a holographic plate, just a fine-grained pattern of interference lines. But for the picture to be experienced by a viewer, or to be available for data access in an artificial brain, that picture must first be reified out of that pattern of interference lines into an actual image again; that is, the holograph must be illuminated by a beam of coherent light. After passing through the holographic plate, that beam of light generates a volumetric array of patterned light, every point of which is determined by the sum of all of the light rays passing through that point, and it is that volumetric pattern of light in space that is observed when viewing a hologram.

So if holography is to serve as a metaphor for consciousness, the key question is whether the metaphorical hologram is illuminated by coherent light to produce a volumetric spatial pattern of light or whether the hologram in experience is like a holographic plate in the dark. If it is the former, then conscious experience in this metaphor is the pattern of light waves interfering in three-dimensional space. It is a spatial image that occupies a very specific portion of physical space, and it requires energy to maintain it in that space. This is exactly the kind of mechanism we should be looking for in the brain. If it were the latter, as **Velmans** suggests, then why would the shape of our experience not be that of the interference patterns etched on the holographic plate, rather than the volumetric image they encode? What magical substance or process in conscious experience performs the volumetric reconstruction that in the real universe requires an actual light beam and some complicated interference process to reconstruct? If it is a spatial structure that we observe in consciousness, then it is a spatial structure that we must seek out in the brain, not a potentially spatial structure that remains stillborn in a nonspatial form. Otherwise, the spatial image-like nature that is so salient a property of subjective experience must remain a magical mystical entity forever in principle beyond the reach of science.

In the target article, I commended **Grossberg** for advocating explicit filling-in to account for Gestalt illusions, but chided him for not extending that same reasoning into the third dimension. To this, Grossberg responds that I have not kept up with the modeling literature; he cites the FACADE and LAMINART as models that explain many three-dimensional figure-ground, grouping, and filling-in percepts, including transparency, and that use an explicit surface filling-in process. This I do not doubt. But in both FACADE and LAMINART, depth is handled in a disparity based representation with left and right eye image pairs and disparity images to represent depth information. In neither of those models is there a three-dimensional volumetric spatial matrix with receptive fields at every location and every orientation in three dimensions, as would be required for a neural network model of spatial experience. While FACADE and LAMINART do perform explicit filling-in of both contours and surfaces, the filling-in itself does not propagate in the third dimension by diffusion as it does in the other two. The third dimension is handled very differently than the other two, and the result is a 2-D sketch rather than a full volumetric spatial matrix. Whatever their merits as neurophysiologically plausible models, these models leave something to be desired as *perceptual* models, because perceptual experience *is* fully volumetric and three-dimensional, and multiple depth values can be experienced in any direction.

If **Grossberg**'s argument that explicit filling-in is required to account for two-dimensional illusions has any validity at all, then it should apply just as well to three-dimensional perception as it does to two, at least for a *perceptual* model that models the experience rather than its neurophysiological correlate.

R6. World as external memory

McLoughlin endorses O'Regan's (1992) concept of seeing as an active process of probing the environment as though it were a continuously available external memory. But probing the world with visual saccades, especially in the monocular case, is nothing like accessing a memory, internal or external, because every saccade retrieves only a two-dimensional pattern of light. The three-dimensional spatial information of the external world is by no means immediately available from glimpses of the world but requires the most sophisticated and as yet undiscovered algorithm to decipher that spatial information from the retinal input. Furthermore, in the absence of a global framework to register the information from each saccade in its proper place, vision as described by O'Regan would be indistinguishable from apperceptive agnosia, a visual integration failure. In other words, the condition of apperceptive agnosia is the absence of a visual function whose existence O'Regan effectively denies. McLoughlin is right that the brain need not explicitly represent more than it requires at any particular time, and it can make do with a sparse or abbreviated representation of the world. But he misses the paradigmatic point that the world we observe in experience is already that sparse representation, the real world beyond experience being infinitely more complex than our experience of it. So the brain must explicitly encode exactly as much detail as we observe in experience, no less, and unbiased phenomenological observation clearly reveals a spatially structured world.

Fox complains that I refute direct perception on the grounds that no plausible mechanism has ever been identified neurophysiologically that accounts for the external nature of perception. "Yet," says Fox, "there is growing physiological evidence to the contrary," and he cites neurophysiological findings *in the brain.* But the kind of physical evidence required to support *direct perception* would have to be energy or information located *outside* the physical brain, out in external space where perception is supposed by direct realism to occur. Fox chides me that "Using the term 'perceptual processing' or 'computation' is a serious misrepresentation of direct perception." He is quite correct. But that is exactly what is wrong with Gibson's theory of direct perception, and that is exactly why modern proponents of Gibson's theories usually take care to disclaim his most radical views. For if perception is *not* a computation in the brain based on sensory input, then why does Fox cite evidence from the brain to explain that perception? Fox suggests "A more fruitful heuristic for understanding perception is a physiology that has evolved a sensitivity to meaningful environmental relational information or . . . action-oriented systems." And how would one build an artificial system with a "sensitivity to meaningful environmental relational information" that is *not* attained by way of input through sensory systems and internal representations? This "explanation" is every bit as mysterious as the property of consciousness it is supposed to explain.

Lloyd disputes the phenomenological basis of the Gestalt Bubble model and insists that outside of focal attention he experiences only a very indefinite spatiality, which seems to him inconsistent with the continuously present three-dimensional model constructed in the Gestalt Bubble. Instead, he proposes that the natural supposition that our experience specifies a full 360-degree diorama arises from the "just-in-time" availability of spatial information with every attentional focus. But the availability of spatial information is not only "just in time," but, more significantly, it is also "just in place," that is, the spatial percept appears at the point in the global experience of three-dimensional space that the percept is perceived to occupy in that space. Lloyd's Gibsonian view also fails to account for dreams and hallucinations, where the world as an external memory is no longer available for data access, and yet a structured world is experienced nonetheless. There is no question that there is a loss of resolution in peripheral vision – that too is easily confirmed phenomenologically. But if Lloyd's experience of each individual saccade appears separate and disconnected from any global whole, like a series of scenes on a television screen, then either he is suffering a form of apperceptive agnosia, or more likely, his theory of vision suffers from apperceptive agnosia, which in turn handicaps his phenomenological observations. This suspicion is supported by Lloyd's own analysis of the dimensions of conscious experience. The basic dimension, according to Lloyd, is temporal, and experience is an orderly ensemble of phenomenal leaps and bounds along a time line. Spatiality emerges from trajectories encoded in proprioception, which orient each momentary percept to

https://doi.org/10.1017/S0140525X03410091

those before and after. This is the consequence of designing a phenomenology based on one's theory of perception, rather than a theory of perception based on one's phenomenology!

R7. Neurophysiological issues

McLoughlin argues that the fragmented architecture of the visual cortex into separate retinotopic maps requires a fragmented model of vision. But that is true only for a purely neurophysiological model that cares nothing about phenomenology, where the unity of visual experience is its most salient feature. But a neuroscience that explains everything about the brain except how it generates consciousness, is a neuroscience that explains nothing, because it is consciousness that makes the brain interesting in the first place. To declare from the outset that the unity of consciousness requires no explanation is to guarantee that no explanation will ever be found.

Ross agrees that a simple neuron doctrine cannot account adequately for size constancy in perception, but contends that more complex neurological models show promise. She then cites a number of neurophysiological models that account for some aspect or other of size constancy. But curiously, none of the models that Ross cites accounts for the one aspect of size constancy that is the central focus of the target article; that is, the fact that objects in space appear as solid volumetric objects embedded in a volumetric surrounding space, and that space has the peculiar property that its size scale shrinks progressively in nonlinear fashion with distance from the egocentric point. Both Gibson and the Gestaltists complained about the trend in psychophysics of breaking the complex phenomenon of visual experience into very simple visual tasks that are then recorded as keypress data points in psychophysical studies. Neural network or other models are then devised to replicate those data points, and those models are then considered to be models of vision. Lost in the shuffle is the rich and complex volumetric spatially extended experience of visual consciousness, which never finds its way into those models of vision.

Grossberg is quite right when he says that the Gestalt Bubble model "makes no contact with neurophysiological and anatomical data about vision." This means either that the model is wrong, or that neuroscience is in a state of serious crisis because it offers no hint of an explanation for the observed properties of conscious experience. If the latter should happen to be the case, as the target article suggests, then limiting our observations of our phenomenal experience to that which is allowed by contemporary theories of neural representation will turn out to have been an exercise in futility.

Duch also complains that I misrepresent the neuron doctrine by omitting discussion of dynamic recurrent neural networks. He must have missed my discussion of the dynamic recurrent neural network models of Grossberg, and their fundamental difficulties with modeling spatial experience (sect. 3 of the target article).

MacKay bolsters the evidence for Gestalt processes in the brain by considering the web of continuous electrical activity stretching from the spinal cord to the cerebrum. MacKay proposes that the "panexperientialist" view suggests that awareness is linked to something like an electrical field of this sort. Indeed, that is exactly why Köhler was so interested in electrical fields. My own preference is for a *harmonic resonance* theory (Lehar 2003) involving patterns of electrochemical standing waves in the neural substrate. Standing waves inherit all the properties of static electric fields, and add to them an extraordinarily rich repertoire of spatiotemporal behaviors that are very Gestalt-like in nature. This hypothesis also resolves the issue of integration raised by MacKay, because it is in the very nature of different resonances in a mechanism to couple to each other and thereby produce a single larger integrated resonance, of which the original resonances become higher harmonics (Lehar 2003b).

R8. Various and sundry issues

Laming raises the homunculus objection, that if there were picture-like processes active in perception, then there would have to be an internal viewer of those picture-like processes. I refuted this objection in the target article with the argument that information encoded in the brain needs to be available only to other internal processes rather than to a miniature copy of the whole brain. Laming rejects this explanation with the statement "The fact that Lehar has a mathematical model to replace the neurophysiological observations does not alter this requirement." But the requirement for an internal observer of any spatial data is itself a paradigmatic assumption on Laming's part. He has not shown that it is necessary in the first place, and it is at least equally likely prima facie that it is not. Furthermore, we know for a fact that our experience is expressed in the form of a spatial structure, regardless of whether that structure requires an observer, and that experienced structure can be expressed in a perceptual model. There is no reason a model of perceptual experience should be invalid.

Luccio takes issue with my characterization of Gestalt theory as a representationalist theory; he claims that it is neither representationalist nor antirepresentationalist but is merely "indifferentist" to the epistemological question. There have been different schools of Gestalt, not all of which have shared the same philosophy. But at least Koffka and Köhler, and therefore by implication presumably Wertheimer, were definitely representationalists. Koffka makes the most clear representationalist case with his distinction between the "geographical environment" (the objective external world) and the "behavioral environment" (the phenomenal world), and he clearly stated that the behavioral environment is located inside the geographical body in the geographical environment (see Koffka 1935, p. 40, Fig. 2). Köhler expressed his representationalist views most clearly in Köhler 1971, p. 125.

That is not to say that one can't be a Gestaltist and a direct realist. One can profess, like Gibson, that illusions are not illusory at all, and that perceived illusory surfaces have a real objective existence out in the physical world, although that existence cannot be verified by scientific means, and the function of the sense organs becomes highly ambiguous. In my view, the message of Gestalt has been representationalist from the very beginning, with its focus on objects experienced vividly in phenomenal space that are known to have no objective existence.

Marković is puzzled that I can claim at one point "the internal perceptual representation encodes properties of the distal object rather than of the proximal stimulus" (sect. 9, last para.) while at another I state "the direct realist view is incredible because it suggests that we can have experience of objects out in the world directly, beyond the sensory surface, as if bypassing the chain of sensory processing" (sect. 2.1, first para.). Why, asks Marković, would the thesis that distal objects are mapping onto the phenomenological domain without neural intervention be incredible and mysterious, while the idea about the projection of internal representation onto the external perceptual world not be incredible and mysterious? How is it possible that perception is partially indirect (representational), and partially direct (distally oriented)?

Perception is entirely indirect; what we experience is in every sense inside our physical head. The "distal orientation" of perception is seen in the form or dimensions in which perceptual information is expressed. The perceived world is expressed not in terms of the proximal image on the sensory surface, that is, a two-dimensional pattern of brightnesses, but in terms of actual three-dimensional objects and surfaces in a world that we take to be reality. We do not see visual patterns and infer them to be a table, we experience a table, expressed in terms of volumes of perceived wood embedded in a volume of perceived space. But that world information does not enter experience directly in some magical mystical manner, but indirectly by the conventional route of sensory input, and that input is expanded out or reified in the brain to become the spatial percept that we experience.

Randrup complains that my position is not really materialist, because I say in the target article that "there remains a vivid subjective quality (or *quale*), to the experience of red" for example, which "is not in any way identical to any externally observable physical variable" (sect. 3, para. 6) in the brain. This quoted passage however represents not my own view, but my summary characterization of Chalmers' "hard problem" of consciousness, and why it is considered by some to imply a fundamental dualism. According to identity theory, the difference between subjective experience and its objective physical realization is a difference in viewpoint or perspective, rather than an ontological dualism, and that makes the Gestalt Bubble a materialist position.

Randrup himself favors an idealist position, and goes on to conclude that the Gestalt Bubble model is most readily understood within the idealist world view, whereby the troubles of direct or indirect perception are significantly reduced. It is true that the Gestalt Bubble model is useful even as a structural description of pure mind. But to deny the existence of an independent objective material world, of which that mind is a copy, strains my credulity beyond its elastic limit.

Rosenthal & Visetti are generally supportive of the perceptual modeling approach proposed in the target article. However, they are puzzled about whether the proposed mechanism of emergence in the model is motivated primarily by the emergent properties of perception, or whether it is a physicalist model, whose spatial matrix and fieldlike interactions represent physical space and physical forces in the brain.

In the first place, the model is explicitly defined as a model of experience, and the local elements in the model are defined as local perceptual experiences. The dynamic fieldlike forces are therefore defined as perceptual tendencies observed phenomenologically, for example, the tendency for perceived surfaces to fill in like a milky soap bubble, and the tendency for corner or occlusion percepts to link up to produce globally coherent edges. Although the dynamics of these experiences are usually so fleeting as to be impossible to observe, it is the configuration of the end result, or final stable percept, that implicates an emergent spatial filling-in, because no other mechanism could plausibly produce that result. So the Gestalt Bubble model is not a physicalist model of the brain, but a mathematical model of experience, although it is committed to an emergent spatial computational strategy as offering the best explanation for spatial experience, and that in turn sets constraints on the corresponding neurophysiological mechanism in the brain.

Schirillo proposes to extend the Gestalt Bubble model by adding the perception of illumination to that of spatial structure. Schirillo is an astute phenomenological observer; the perception of brightness, lightness, and illumination are indeed intimately coupled to the perception of visual structure. I have explored the interaction of spatial perception to the perception of illumination in Chapter 5 of my book (Lehar 2003b).

Tse is generally supportive of the Gestalt Bubble model, and offers a more general analysis of the why the visual system operates as it does. Tse identifies two general principles by which the visual system attempts to correct errors. One involves completing missing information on the basis of knowledge about what most likely exists in the scene, that is, perceptual "filling-in"; and the other involves exploiting the physical stability of the environment as a reference frame with respect to which the eyes and body can move. An interesting aspect of this view is that the visual system implicitly understands its own limitations, and attempts perceptual filling-in only when it "knows" that it has failed to detect something that it believes must be present.

I take issue, however, with **Tse**'s contention that in amodal completion there is no perceptual filling-in. It is true that there is no *modal* filling-in of explicit surfaces, as in the Kanizsa figure, which is what Tse probably intended. But there is filling-in nonetheless, although of a nonsensory, *amodal* manner. When we see an occluded object, like a horizontal branch, part of which is occluded by a nearer vertical tree trunk, we can reach back behind the tree trunk and grab exactly that point in space that we "know" to be occupied by the occluded branch, based exclusively on the configuration of its visible portions. Although it is a semantic issue whether such experience is really "seeing" at all, there is no question that a three-dimensional volumetric *experience* is involved, and that experience is produced by filling-in processes very much like those seen in the modal surfaces of the Gestalt illusions.

Wright supports the representational stance in the target article, and provides further arguments to defeat the alternative direct realist view. Wright objects, however, to the use of the term "subjective" when applied to the sensory field, saying the sensory representation of colored volumes embedded in perceived space is thoroughly nonconceptual, and therefore no kind of subjective judgment is involved at that low level of experience. If television is an apt analogy for the televisual function of vision, Wright suggests I acknowledge the nonconceptual nature of the pattern of glowing phosphor dots on the television screen.

The semantic distinction that **Wright** draws between nonconceptual sensory processes and subjective judgment

https://doi.org/10.1017/S0140525X03410091

based on that sensory data may serve him well for his own purposes, but it is at odds with a prominent theme in Gestalt theory: that there is no difference in essential principles between the lower-level functions of sensation and perception, and the higher-level functions of recognition and cognition, except for a difference in complexity (cf. Lehar 2003b, Ch. 6). The higher-level recognition of a table as a whole is not different in principle from the recognition of its edges and surfaces as edges and surfaces. In the television analogy, the individual pixels of a photosensor array can be seen as very simple "feature detectors" tuned to respond to their feature, the brightness and color of light from a narrow angle of the visual field. Similarly, the local spatial fields proposed in the Gestalt Bubble model can be seen as three-dimensional surface-element, edge-element, and corner-element "feature detectors" that, in cooperation and competition with their neighbors, make a collective "subjective judgment" about the presence or absence of edge or corner features in particular parts of space. What is missing in the Gestalt Bubble model is the strict input/output function normally ascribed to feature detectors, because the "output," or final state of a particular "detector," depends not only on the input from the retina, or only on local interactions in perceived space, but on the total configuration of all of the other local elements across the whole of phenomenal space simultaneously. Even the highest-level global recognition has an influence on the state of the lowest-level edges and surface brightnesses of a scene, as seen in the subjective reversals of Figure 11D of the target article. That perceived corners and surfaces are observed to change their configuration with the perceptual shift, clearly indicates the "subjective" nature of these low-level components of experience, which are not strictly invariant with the input, as Wright suggests.

References

Letters "a" and "r" appearing before authors' initials refer to target article and response, respectively.

Adelson, E. H. (1993) Perceptual organization and the judgment of brightness. *Science* 262:2042–44. [JAS]

(2000) Lightness perception and lightness illusions. In: *The new cognitive neurosciences,* 2nd edition, ed. M. S. Gazzaniga. MIT Press. [rSL, JH]

Aldrich, V. C. (1979) Objective sense-data. *The Personalist* 60:36–42. [EW]

Amit, D. J. (1994) The Hebbian paradigm reintegrated: Local reverberations as internal representations. *Behavioral and Brain Sciences* 18:617–26. [WD]

Andrews, D. P. (1964) Error-correcting perceptual mechanisms. *Quarterly Journal of Experimental Psychology* 16:102–15. [HER]

Angelucci, A., Levitt, J. B., Walton, E. J., Hupe, J. M., Bullier, J. & Lund, J. S. (2002) Circuits for local and global signal integration in primary visual cortex. *Journal of Neuroscience* 22:8633–46. [NPM]

Anstis, S. & Howard, I. (1978) A Craik-O'Brien-Cornsweet illusion for visual depth. *Vision Research* 18:213–17. [aSL]

Aristotle (1976) *De anima (On the soul).* Arno Press. [JAS]

Arnheim, R. (1969) *Art and visual perception: A psychology of the creative eye.* University of California Press. [aSL]

Attneave, F. (1971) Multistability in perception. *Scientific American* 225:142–51. [aSL]

(1982) Prägnanz and soap bubble systems: A theoretical exploration. In: *Organization and representation in perception,* ed. J. Beck. Erlbaum. [aSL, RL, SM]

Ayer, A. J. (1957) *The problem of knowledge.* Penguin Books. [EW]

Barlow, H., Blakemore, C. & Pettigrew, J. (1967) The neural mechanism of binocular depth discrimination. *Journal of Physiology* 193:327–42. [aSL]

Barrow, H. G. & Tenenbaum, J. M. (1981) Interpreting line drawings as three-dimensional surfaces. *Artificial Intelligence* 17:75–116. [aSL]

Bisiach, E., Capitani, E., Luzatti, C. & Perani, D. (1981) Brain and conscious representation of outside reality. *Neuropsychologia* 19:543–52. [aSL]

Bisiach, E. & Luzatti, C. (1978) Unilateral neglect of representational space. *Cortex* 14:129–33. [aSL]

Blakemore, C. & Tobin, E. A. (1972) Lateral inhibition between orientation detectors in the cat's visual cortex. *Experimental Brain Research* 15:439–40. [NPM]

Blank, A. A. (1958) Analysis of experiments in binocular space perception. *Journal of the Optical Society of America* 48:911–25. [aSL]

Blumenfeld, W. (1913) Untersuchungen Über die Scheinbare Grösse im Sehraume. *Zeitschrift für Psychologie* 65:241–404. [aSL]

Booth, D. A. (1978) Mind-brain puzzle versus mind-physical world identity. (Commentary on R. Puccetti & R. W. Dykes, Sensory cortex and the mind-brain problem.) *Behavioral and Brain Sciences* 3:348–49. [DAB]

Booth, D. A. & Freeman, R. P. J. (1993) Discriminative measurement of feature integration in object recognition. *Acta Psychologica* 84:1–16. [DAB]

Boring E. G. (1933) *The physical dimensions of consciousness.* Century. [aSL]

(1942) *Sensation and perception in the history of experimental psychology.* Appleton-Century. [RL]

(1950) *A history of experimental psychology.* Appleton-Century-Crofts. [RL]

Bosking, W. H., Zhang, Y., Schofield, B. & Fitzpatrick, D. (1997) Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *Journal of Neuroscience* 17:2112–27. [NPM]

Bressan, P. (1993) Neon colour spreading with and without its figural prerequisites. *Perception* 22:353–61. [aSL]

Broad C. D. (1925) *The mind and its place in nature.* Routledge & Kegan Paul. [aSL]

Brookes, A. & Stevens, K. (1989) The analogy between stereo depth and brightness. *Perception* 18:601–14. [aSL]

Bruce, V. & Green, P. (1987) *Visual perception: Physiology, psychology, and ecology.* Erlbaum. [aSL]

Carman, G. J. & Welch, L. (1992) Three-dimensional illusory contours and surfaces. *Nature* 360:585–87. [aSL]

Chalmers, D. J. (1995) Facing up to the problems of consciousness. *Journal of Consciousness Studies* 2:200–19. Reprinted in: *Toward a science of consciousness II: The second Tucson discussions and debates* (1996), ed. S. R. Hameroff, A. W. Kaszniak & A. C. Scott. MIT Press. [aSL]

Charnwood, J. R. B. (1951) *Essay on binocular vision.* Halton Press. [aSL]

Chemero, A. (2003). An outline of a theory of affordances. *Ecological Psychology.* [CRF]

Churchland, P. M. (1984) *Matter and consciousness: A contemporary introduction to the philosophy of mind.* MIT Press. [aSL, WD]

Clark, A. (1993) *Sensory qualities.* Oxford University Press. [aSL]

(1998) *Being there: Putting brain, body, and world together again.* MIT Press. [CRF]

(2000) *A theory of sentience.* Oxford University Press. [DLl]

Collett, T. (1985) Extrapolating and interpolating surfaces in depth. *Proceedings of the Royal Society of London, Series B* 224:43–56. [aSL]

Coren, S., Ward, L. M. & Enns, J. J. (1994) *Sensation and perception.* Harcourt Brace. [aSL]

Cornsweet, T. N. (1970) *Visual perception.* Academic Press. [aSL]

Crick, F. (1994) *The astonishing hypothesis: The scientific search for the soul.* Scribners. [aSL]

Crick, F. & Koch, C. (1990) Toward a neurobiological theory of consciousness. *Seminars in the Neurosciences* 2:263–75. [aSL,SM]

Cronly-Dillon, J. R. & Gregory, R. L., eds. (1991) *Vision and visual dysfunction, vol. 2: Evolution of the eye and visual system.* CRC Press. [DDH]

Cutting, J. E. & Vishton, P. M. (1995) Perceiving layout and knowing distances: The interaction, relative potency, and contextual use of different information about depth. In: *Perception of space and motion,* ed. W. Epstein & S. Rogers. Academic Press. [JH]

Davidson, D. (1970) *Mental events.* Oxford University Press. [aSL]

Dennett, D. C. (1981) Two approaches to mental images. In: *Imagery,* ed. N. Block. MIT Press. [arSL]

(1991) *Consciousness explained.* Little, Brown. [aSL]

(1992) "Filling in" versus finding out: A ubiquitous confusion in cognitive science. In: *Cognition: Conceptual and methodological issues,* ed. H. L. Pick Jr., P. van den Broek & D. C. Knill. American Psychological Association. [aSL]

Denny-Brown, D. & Chambers, R. A. (1958) The parietial lobe and behavior. *Annual Research in Nervous and Mental Disease* 36:35–117. [aSL]

de Renzi, E. (1982) *Disorders of space exploration and cognition.* Wiley. [aSL]

Deubel, H. et al. (1998) Immediate post-saccadic information mediates space constancy. *Vision Research* 38:3147–59. [PUT]

Dobbins, A. C., Jeo, R. M., Fiser, J. & Allman, J. M. (1998) Distance modulation of neural activity in the visual cortex. *Science* 281:552–55. [HER]

Drake, D. (1920) The approach to critical realism. In *Essays in critical realism: A co-operative study of the problem of knowledge,* pp. 3–32, ed. D. Drake, A. O.

Lovejoy, J. B., Pratt, A. K. Rogers, G. Santayana, R. W. Sellars & C. A. Strong. Gordian Press. [aSL]
(1925) *Mind and its place in nature.* Macmillan. [EW]
Drake, D., Lovejoy, A. O., Pratt, J. B., Rogers, A. K., Santayana, G., Sellars, R. W. & Strong C. A. (1920) *Essays in critical realism: A co-operative study of the problem of knowledge.* Gordian Press. [aSL]
Duch, W. (1997) Platonic model of mind as an approximation to neurodynamics. In: *Brain-like computing and intelligent information systems,* ed. S.-I. Amari & N. Kasabov. Springer. [WD]
Duffy, C. & Wurtz, R. (1995) Response of monkey MST neurons to optic flow stimuli with shifted centers of motion. *Journal of Neuroscience* 15:5192–208. [CRF]
(1997a) Medial superior temporal area neurons respond to speed patterns in optic flow. *Journal of Neuroscience* 17:2839–51. [CRF]
(1997b) Planar directional contributions to optic flow responses in MST neurons. *Journal of Neurophysiology* 77:782–96. [CRF]
Earle, D. C. (1998) On the roles of consciousness and representations in visual science. Commentary on Pessoa et al. 1998. *Behavioral and Brain Sciences* 21:757–58. [aSL]
Ebenholtz, S. M. (2001) *Oculomotor systems and perception.* Cambridge University Press. [CRF]
Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M. & Reitboeck, J. (1988) Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biological Cybernetics* 60:121–30. [aSL, SM]
Einstein, A. (1950) *Essays in physics.* Philosophical Library. [DDH]
Evans, G. (1982) *The varieties of reference.* Clarendon Press. [EW]
Farah, M. J. (1990) *Visual agnosia.* MIT Press. [NPM]
Fechner, G. T. (1860/1966) *Elemente der Psychophysik.* Breitkopf and Härtel. English edition, 1966: *Elements of psychophysics, vol. 1,* trans. H. E. Adler. Holt, Rinehart and Winston. [DL]
Feigl, H. (1958) *The "mental" and the "physical."* University of Minnesota Press. [arSL]
Felisberti, F. & Derrington, A. M. (2001) Long-range interactions in the lateral geniculate nucleus of the New-World monkey, *Callithrix jacchus. Visual Neuroscience* 18:209–18. [NPM]
Foley, J. M. (1978) Primary distance perception. In: *Handbook of sensory physiology, vol. 7: Perception,* ed. R. Held, H. W. Leibowitz & H.-L. Tauber. Springer Verlag. [aSL]
Forsee, A. (1963) *Albert Einstein: Theoretical physicist.* Macmillan. [DDH]
Foster, J. (2000) *The nature of perception.* Oxford University Press. [EW]
Fox, C. R. (1999) Special senses 3: The visual system. In: *Neuroscience for rehabilitation,* 2nd edition, ed. H. Cohen. Lippincott Williams & Williams. [CRF]
Gallese V., Craighero L., Fadiga L.& Fogassi, L. (1999) *Perception through action.* PSYCHE 5, 1999; Online journal: http://psyche.cs.monash.edu.au/v5/psyche-5-21-gallese.html [CRF]
Gibson, J. J. (1950) *The perception of the visual world.* Houghton Mifflin. [VR]
(1966) *The senses considered as perceptual systems.* Houghton Mifflin. [RL, CRF]
(1972) A theory of direct visual perception. In: *The psychology of knowing,* ed. J. R. Royce & W. W. Rozeboom. Gordon & Breach. [aSL]
(1977) The theory of affordances. In: *Perceiving, acting and knowing: Toward an ecological psychology,* ed. R. Shaw & J. Bransford. Erlbaum. [EW]
(1979) *The ecological approach to visual perception.* Houghton Mifflin. [CRF, aSL, RL, SM]
Gibson, J. J. & Crooks, L. E. (1938) A theoretical field-analysis of automobile driving. *The American Journal of Psychology* 51:453–71. [aSL]
Gilchrist, A. L. (1977) Perceived lightness depends on perceived spatial arrangement. *Science* 195:185–87. [JAS]
Gillam, B. (1971) A depth processing theory of the Poggendorf illusion. *Perception and Psychophysics* 10:211–16. [aSL]
(1980) Geometrical illusions. *Scientific American* 242:102–11. [aSL]
(1998) Illusions at century's end. In: *Handbook of perception and cognition,* 2nd edition, ed. J. Hochberg. Academic Press. [HER]
Graham, C. H. (1965) Visual space perception. In: *Vision and visual perception,* ed. C. H. Graham. Wiley. [aSL]
Grealy, M. (2002) Closing gaps: Can the generalized intrinsic T-guide model provide a unified account of brain and behavior? Paper presented at the Seventh European Workshop on Ecological Psychology, Bendor Island, France, July 4–6, 2002. [CRF]
Green, M. & Odum, V. J. (1986) Correspondence matching in apparent motion: Evidence for three dimensional spatial representation. *Science* 233:1427–29. [aSL]
Gregory, R. L. (1963) Distortion of visual space as inappropriate constancy scaling. *Nature* 199:678–79. [aSL]
(1971) *The intelligent eye.* Weidenfeld & Nicolson. [SM]
Greve, D., Grossberg, S., Guenther, F. & Bullock, D. (1993) Neural representations for sensory-motor control, I: Head-centered 3-D target positions from opponent eye commands. *Acta Psychologica* 82:115–38. [SG]
Grossberg, S. (1987) Cortical dynamics of three-dimensional form, color, and brightness perception: II. Binocular theory. *Perception and Psychophysics* 41:117–58. [SG, aSL]
(1990) Neural FAÇADE: Visual representations of static and moving form-and-color-and-depth. *Mind and Language* 5 (Special Issue on Understanding Vision):411–56. [aSL]
(1994) 3-D vision and figure-ground separation by visual cortex. *Perception and Psychophysics* 55:48–120. [SG, aSL, NPM, VR]
(1995) The attentive brain. *American Scientist* 83:483–49. [WD]
(1997) Cortical dynamics of three-dimensional figure-ground perception of two-dimensional pictures. *Psychological Review* 104:618–58. [SG]
(1999a) How does the cerebral cortex work? Learning, attention, and grouping by the laminar circuits of visual cortex. *Spatial Vision* 12:163–86. [SG]
(1999b) The link between brain learning, attention, and consciousness. *Consciousness and Cognition* 8:1–44. [SG]
(2000) How hallucinations may arise from brain mechanisms of learning, attention, and volition. *Journal of the International Neuropsychological Society* 6:583–92. [SG]
Grossberg, S. & Howe, P. D. L. (2003) A laminar cortical model of stereopsis and three-dimensional surface perception. *Vision Research.* (in press). [SG]
Grossberg, S. & Kelly, F. (1999) Neural dynamics of binocular brightness perception. *Vision Research* 39:3796–816. [SG]
Grossberg, S. & McLoughlin, N. (1997) Cortical dynamics of 3-D surface perception: Binocular and half-occluded scenic images. *Neural Networks* 10:1583–605. [SG]
Grossberg, S. & Mingolla, E. (1985) Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review* 92:173–211. [aSL]
Grossberg, S., Mingolla, E. & Pack, C. (1999) A neural model of motion processing and visual navigation by cortical area MST. *Cerebral Cortex* 9:878–95. [SG]
Grossberg, S., Mingolla, E. & Ross, W. D. (1997) Visual brain and visual perception: How does the cortex do perceptual grouping? *Trends in Neurosciences* 20:106–11. [SG]
Grossberg, S. & Pessoa, L. (1998) Texture segregation, surface representation, and figure-ground separation. *Vision Research* 38:2657–84. [SG]
Grossberg, S. & Raizada, R. (2000) Contrast-sensitive perceptual grouping and object-based attention in the laminar circuits of primary visual cortex. *Vision Research* 40:1413–32. [SG]
Grossberg, S. & Seitz, A. (2003) Laminar development of receptive fields, maps, and columns in visual cortex: The coordinating role of the subplate. *Cerebral Cortex* (in press). [SG]
Grossberg, S. & Swaminathan, G. (2003) A laminar cortical model for visual perception of slanted and curved 3D surfaces and 2D images: Development, attention, and bistability (submitted). [SG]
Grossberg, S & Todoroviç, D. (1988) Neural dynamics of 1-D and 2-D brightness perception: A unified model of classical and recent phenomena. *Perception and Psychophysics* 43:241–77. [aSL]
Grossberg, S. & Williamson, J. W. (2001) A neural model of how horizontal and interlaminar connections of visual cortex develop into adult circuits that carry out perceptual grouping and learning. *Cerebral Cortex* 11:37–58. [SG]
Guenther, F., Bullock, D., Greve, D. & Grossberg, S. (1994) Neural representations for sensory-motor control, III: Learning a body-centered representation of 3-D target position. *Journal of Cognitive Neuroscience* 6:341–58. [SG]
Gunderson, K. (1971) Asymmetries and mind-body perplexities. *Minnesota Studies in Philosophy of Science* 4:273–309. [KG]
(1984) Skinnerian privacy and Leibnizian privacy. (Invited commentary on aspects of B. F. Skinner's work with response from Skinner). *Behavioral and Brain Sciences* 7(4): Reprinted in: *Canonical papers of B. F. Skinner,* ed. Stephen Harnad, Cambridge University Press, 1988. [KG]
Harrison, S. (1989) A new visualization on the mind-brain problem: Naive realism transcended. In: *The case for dualism,* ed. J. Smythies & J. Beloff. University of Virginia Press. [aSL]
Hatfield, G. & Epstein, W. (1985) The status of minimum principle in the theoretical analysis of visual perception. *Psychological Bulletin* 97:155–86. [SM]
Heckenmuller, E. G. (1965) Stabilization of the retinal image: A review of method, effects, and theory. *Psychological Bulletin* 63:157–69. [aSL]
Heelan, P. A. (1983) *Space perception and the philosophy of science.* University of California Press. [aSL]
Heilman, K. M. & Watson, R. T. (1977) The neglect syndrome – A unilateral defect of the orienting response. In: *Lateralization in the nervous system,* ed. S. Harnad, R. W. Doty, L. Goldstein, J. Jaynes & G. Krauthamer. Academic Press. [aSL]

Heilman, K. M., Watson, R. T. & Valenstein, E. (1985) Neglect and related disorders. In: *Clinical neuropsychology,* ed. K. M. Heilman & E. Valenstein. Oxford University Press. [aSL]

Helmholtz, H. (1925) Physiological optics. *Optical Society of America* 3:318. [aSL]

Hillebrand, F. (1902) Theorie der Scheinbaren Grösse bei Binocularem Sehen. *Denkschrift von acadamesche Wissenschaft, Wien* 72:255–307. [aSL]

Hochberg, J. (1970) Attention, organization and consciousness. In: *Attention: Contemporary theory and analysis,* ed. D. I. Mostofsky. Appleton-Century-Crofts. [JH]

(1978) *Perception.* Prentice Hall. [SM]

(1998) Gestalt theory and its legacy: Organization in eye and brain, in attention and mental representation. In: *Perception and cognition at century's end,* ed. J. Hochberg. Academic Press. [JH]

(in press) Acts of perceptual inquiry: Parsing objects by diagnostic coupling and consequences. *Acta Psychologica.* [JH]

Hochberg, J. & Brooks, V. (1960) The psychophysics of form: Reversible perspective drawings of spatial objects. *The American Journal of Psychology* 73:337–54. [aSL]

Hochberg, J. & McAlister, E. (1953) A quantitative approach to "figural goodness." *Journal of Experimental Psychology* 46:361–64. [JH]

Hochberg, J. & Peterson, M. A. (1987) Piecemeal organization and cognitive components in object perception: Perceptually coupled responses to moving objects. *Journal of Experimental Psychology: General* 116:370–80. [JH]

Hoffman, D. D. (1998) *Visual intelligence: How we create what we see.* W. W. Norton. [aSL]

Hoffman, W. C. (1989) The visual cortex is a contact bundle. *Applied Mathematics and Computation* 32:137–67. [VR]

Howe, P. D. L. & Grossberg, S. (2001) Laminar cortical circuits for stereopsis and surface depth perception. *Society for Neuroscience Abstracts* Vol. 164, No. 17. [SG]

Hubel, D. H. & Wiesel, T. N. (1959) Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology* 148:574–91. [aSL]

Humphrey, G. K. & Goodale, M. A. (1998) Probing unconscious visual processing with the McCollough Effect. *Consciousness and Cognition* 7:494–519. [NPM]

Hurvich, L. & Jameson, D. (1957) An opponent-colors theory of color vision. *Psychological Review* 64:384–404. [JH]

Husserl, E. (1907/1997) *Thing and space: Lectures of 1907.* Kluwer Academic. [VR]

(1966) *Zur Phänomenologie des inneren Zeitbewusstseins (Phenomenology of inner time consciousness).* Martinus Nijhoff. [DLl]

(1974) *Ding und Raum (Thing and space): Lectures of 1907.* Martinus Nijhoff. [DLl]

Idesawa, M. (1991) Perception of 3-D illusory surface with binocular viewing. *Japanese Journal of Applied Physics* 30(4B):L751–L754. [aSL]

Indow, T. (1991) A critical review of Luneberg's model with regard to global structure of visual space. *Psychological Review* 98:430–53. [aSL]

Intraub, H. (1997) The representation of visual scenes. *Trends in Cognitive Sciences* 1:217–21. [JH]

Intraub, H., Hoffman, J. E., Wetherhold, C. J. & Stoehs, S. (2001) Does direction of a planned eye movement affect boundary extension? Paper presented at the Forty-first Annual Meeting of the Psychonomic Society, Orlando, FL, November, 2001. [JH]

James, W. (1904) Does "consciousness" exist? In: *Body and mind: Readings in philosophy,* ed. G. N. A. Vesey. George Allen & Unwin. [MV]

Jameson, D. & Hurvich, L. (1967) The science of color appearance. *Color Engineering* 5:29. [JH]

Jones, H. E., Andolina, I. M., Oakely, N. M., Murphy, P. C. & Sillito, A. M. (2000) Spatial summation in lateral geniculate nucleus and visual cortex. *Experimental Brain Research* 135:279–84. [NPM]

Jones, H. E., Grieve, K. L., Wang, W. & Sillito, A. M. (2001) Surround suppression in primate V1. *Journal of Neurophysiology* 86:2011–28. [NPM]

Julesz, B. (1971) *Foundations of Cyclopean perception.* University of Chicago Press. [aSL]

Kanizsa, G. (1979) *Organization in vision.* Praeger. [aSL,RL]

(1994) Gestalt theory has been misinterpreted, but also has had some real conceptual difficulties. *Philosophical Psychology* 7:149–62. [RL]

Kanizsa, G. & Luccio, R. (1986) Die Doppeldeutigkeiten der Prägnanz. *Gestalt Theory* 8:99–135. [RL]

(1990) The phenomenology of autonomous order formation in perception. In: *Synergetics of cognition,* ed. H. Haken & M. Stadler. Springer. [RL]

Kant, I. (1781/1991) *Critique of pure reason,* ed. Vasilis Politis. Dent. [arSL]

Kapadia, M. K., Westheimer, G. & Gilbert, C. D. (2000) Spatial distribution of contextual interactions in primary visual cortex and in visual perception. *Journal of Neurophysiology* 84:2048–62. [NPM]

Kassubek, J., Otte, M., Wolter, T., Greenlee, M. W., Mergner, T. & Lücking, C. H. (1999) Brain imaging in a patient with hemimicropsia. *Neuropsychologia* 37:1327–34. [HER]

Kaufman, L. (1974) *Sight and mind.* Oxford University Press. [aSL]

Kellman, P. J., Machado, L. J., Shipley, T. F. & Li, C. C. (1996) Three-dimensional determinants of object completion. *Annual Review of Vision and Ophthalmology (ARVO) Abstracts* 3133 37 (3):S685. [aSL]

Kellman, P. J. & Shipley, T. F. (1991) A theory of visual interpolation in object perception. *Cognitive Psychology* 23:141–221. [aSL]

Kelly, F. & Grossberg, S. (2000) Neural dynamics of 3-D surface perception: Figure-ground separation and lightness perception. *Perception and Psychophysics* 62:1596–618. [SG]

Kim, J. (1998) *Mind in a physical world: An essay on the mind-body problem and mental causation.* MIT Press. [aSL]

Kinsbourne, M. (1987) Mechanisms of unilateral neglect. In: *Neurophysiological and neuropsychological aspects of spatial neglect,* ed. M. Jeannerod. North-Holland. [aSL]

(1993) Orientational bias model of unilateral neglect: Evidence from attentional gradients within hemispace. In: *Unilateral neglect: Clinical and experimental studies,* ed. I. H. Robertson & J. C. Marshall. Erlbaum. [aSL]

Kirschfeld, K. (1999) Afterimages: A tool for defining the neural correlate of visual consciousness. *Consciousness and Cognition* 8:462–83. [HER]

Knight, G. (2001) Idealism, intentionality, and nonexistent objects. *Journal of Philosophical Research* 26:43–52. [AR]

Koenderinck, J. J. (1990) The brain as a geometry engine. *Psychological Research* 52:122–27. [VR]

Koffka, K. (1935) *Principles of Gestalt psychology.* Harcourt Brace/ Kegan, Paul, Trench & Trubner. [CRF, arSL, RL, SM, VR]

(1969) *The task of Gestalt psychology.* Princeton University Press. [CRF]

Kohler, I. (1961) Interne und externe Organisation in der Wahrnehmung. *Psychologische Beiträge* (Festschrift für W. Köhler) 6:426–38. [BD]

Köhler, W. (1920/1938) Die physische Gestalten in Ruhe und stationären Zustand: Eine naturphilosophische Untersuchung. (Physical Gestalten). In: *A source book of Gestalt psychology,* ed. W. D. Ellis (1938). Routledge & Kegan Paul. (Reprinted from Brownschweig: Vieweg & Son.) [SM]

(1924) *Die physischen Gestalten in Ruhe und im stationären Zustand: Eine naturphilosophische Untersuchung.* Verlag der Philosophischen Akadamie. [aSL]

(1927/1971) Zum Problem der Regulation (On the problem of regulation). In: *The selected papers of Wolfgang Köhler,* ed. M. Henle. Liveright. (English edition, 1971). [SM]

(1940) *Dynamics in psychology.* Liveright. [RL, CRF]

(1947) *Gestalt psychology.* Liveright. [SM, CRF]

(1960) The mind-body problem. In: *Dimension of mind,* ed. S. Hook. New York University Press. [RL]

(1961) Gestalt psychology today. In: *Documents of Gestalt Psychology,* ed. M. Henle. University of California Press. [CFR]

(1969) *The task of Gestalt psychology.* Princeton University Press. [aSL, RL]

(1971) A task for philosophers. In: *The selected papers of Wolfgang Köhler,* ed. Mary Henle. Liveright. [arSL]

Köhler, W. & Held, R. (1947) The cortical correlate of pattern vision. *Science* 110:414–19. [aSL]

Kolb, B. & Whishaw, I. Q. (1996) *Fundamentals of human neuropsychology.* W. H. Freeman. [aSL]

Kopfermann, H. (1930) Psychologische Untersuchungen über die Wirkung Zweidimensionaler Darstellunger körperliche Gebilde. *Psychologische Forschung* 13:293–364. [JH]

Kosslyn, S. M. (1975) Information representation in visual images. *Cognitive Psychology* 7:341–70. [aSL]

(1980) *Image and mind.* Harvard University Press. [aSL]

(1994) *Image and brain: The resolution of the imagery debate.* MIT Press. [aSL]

Kuhn, T. S. (1970) *The structure of scientific revolutions.* Chicago University Press. [aSL]

Làdavas, E., Berti, A., Ruozzi, E. & Barboni, F. (1997) Neglect as a deficit determined by an imbalance between spatial representations. *Experimental Brain Research* 116:493–500. [aSL]

Laming, D. (1986) *Sensory analysis.* Academic Press. [DL]

(1997) *The Measurement of sensation.* Oxford University Press. [DL]

(2004) *Understanding human motivation: What makes people tick?* Blackwell. [DL]

Lee, D. N. & Aronson, E. (1974) Visual proprioceptive control of standing in human infants. *Perception and Psychophysics* 15:529–32. [PUT]

Lee, D., Georgopoulos, A., Pepping, G.-J. & Lee, T. M. (2002) Information for movement guidance in the nervous system. Paper presented at the Seventh European Workshop on Ecological Psychology, Bendor Island, France, July 4–6, 2002. [CRF]

Lehar, S. (2003a) Directional harmonic theory: A computational Gestalt model to account for illusory contour and vertex formation. *Perception* 32:423–448. Also available at: http://cns-alumni.bu.edu/~slehar/webstuff/dirhr1/dirhr1.html. [aSL]

(2003b) *The world in your head: A Gestalt view of the mechanism of conscious experience.* Erlbaum. [arSL]

Lesher, G. W. (1995) Illusory contours: Toward a neurally based perceptual theory. *Psychonomic Bulletin and Review* 2:279–321. [aSL]

Levitt, J. B. & Lund, J. S. (1997) Contrast dependence of contextual effects in primate visual cortex. *Nature* 387:73–76. [NPM]

Livingstone, M. & Hubel, D. (1988) Segregation of form, color, movement, and depth: Anatomy, physiology, and perception. *Science* 240:740–49. [NPM]

Llinas, R. R., Ribary, U., Joliot, M. & Wang, X.-J. (1994) Content and context in temporal thalamocortical binding. In: *Temporal coding in the brain,* ed. G. Buzsaki, R. R. Llinas & W. Singer. Springer-Verlag. [aSL]

Lloyd, D. (2002) Functional MRI and the study of human consciousness. *Journal of Cognitive Neuroscience* 14:818–31. [DLl]

(2003) *Radiant cool: A novel theory of consciousness.* MIT Press. [DLl]

Logvinenko, A., Kane, J. & Ross, D. A. (2002) Is lightness induction a pictorial illusion? *Perception* 31:73–82. [JAS]

Löw, R. (1984) The metaphysical limits of evolutionary epistemology. In: *Concepts and approaches in evolutionary epistemology,* ed. F. Wuketits. Reidel. [AR]

Luchins, A. S. & Luchins, E. H. (1982) An introduction to the origins of Wertheimer's Gestalt psychology. *Gestalt Theory* 4:145–71. [RL]

(1999) Isomorphism in Gestalt theory: Comparison of Wertheimer's and Köhler's concepts. *Gestalt Theory* 21:208–34. [RL]

Luneburg, R. K. (1950) The metric of binocular visual space. *Journal of the Optical Society of America* 40:627–42. [aSL]

Lyons, W. (1986) *The disappearance of introspection.* MIT Press. [DAB]

Mack, A. & Rock, I. (1998) *Inattentional blindness.* MIT Press. [DLl]

Marković, S. & Gvozdenović, V. (2001) Symmetry, complexity and perceptual economy: Effects of minimum and maximum simplicity conditions. *Visual Cognition* 8:305–27. [SM]

Marr, D. (1982) *Vision.* W. H. Freeman. [SM]

Marr, D. & Poggio T. (1976) Cooperative computation of stereo disparity. *Science* 194:283–87. [aSL]

Marshall, J. C. & Halligan, P. W. (1995) Seeing the forest but only half the trees? *Nature* 373:521–23. [aSL]

Maund, B. (1975) The representative theory of perception. *Canadian Journal of Philosophy* 5:41–55. [EW]

(1995) *Colours: their nature and representation.* Cambridge University Press. [EW]

McCready, D. (1985) On size, distance, and visual angle perception. *Perception and Psychophysics* 37:323–34. [HER]

McFie, J. & Zangwill, O. L. (1960) Visual-constructive disabilities associated with lesions of the left cerebral hemisphere. *Brain* 83:243–60. [aSL]

McLoughlin, N. & Grossberg, S. (1998) Cortical computation of stereo disparity. *Vision Research* 38:91–99. [SG]

Metzger, W. (1936) *Gesetze des Sehens.* W. Kramer. [BD]

Michaels, C. & Carello, C. (1981) *Direct perception.* Prentice-Hall. [RL]

Millar, A. (1991) *Reasons and experience.* Clarendon Press. [EW]

Miller, G. A., Galanter, E. & Pribram, K. (1960) *Plans and the structure of behavior.* Holt, Rinehart and Winston. [JH]

Milner, A. D. & Goodale M. A. (1995) *The visual brain in action.* Oxford University Press. [CRF, aSL]

Mitchison, G. (1993) The neural representation of stereoscopic depth contrast. *Perception* 22:1415–26. [aSL]

Morel, J. M. & Solimini, S. (1995) *Variational methods in image segmentation.* Birckhaüser. [VR]

Moutoussis, K. & Zeki, S. (2002) The relationship between cortical activation and perception investigated with invisible stimuli. *Proceedings of the National Academy of Sciences USA* 99:9527–32. [WAM]

Movshon, J. A., Adelson, E. H., Gizzi, M. S. & Newsome, W. T. (1986) The analysis of moving patterns. In: *Pattern recognition mechanisms,* ed. C. Chagas, R. Gattass & C. Cross. Springer Verlag. [aSL]

Müller, G. E. (1896) Zur Psychophysik der Gesichtsempfindungen. *Zeitschrift für Psychologie* 10. [aSL]

Munk, M. H. J., Roelfsema, P. R., König, P., Engel, A. K. & Singer, W. (1996) Role of reticular activation in the modulation of intracortical synchronization. *Science* 272:271–74. [WAM]

Nagel, T. (1965) Physicalism. *Philosophical Review* 74:339–56. [KG]

(1974) What is it like to be a bat? *Philosophical Review* 83:435–50. [KG, aSL]

Olson, C. & Gettner, S. (1995) Object-centered direction selectivity in the macaque supplementary eye field. *Science* 269:985–88. [CRF]

Opie, J. (1999) Gestalt theories of cognitive representation and processing. *Psycoloquy* 10(021). [aSL]

O'Regan, J. K. (1992) Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology* 46:461–88. [arSL, NPM, PUT]

O'Regan, J. K. & Noë, A. (2001) A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences* 24:883–975. [WD, JH]

(2002) A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences* 24:793–810. [JH, EW]

Pace-Schott, E. F. & Hobson, J. A. (2002) The neurobiology of sleep: Genetics, cellular physiology and subcortical networks. *Nature Reviews Neuroscience* 3:591–605. [WAM]

Palmer, S. E. (1992) Modern theories of Gestalt perception. In: *Understanding vision,* ed. G. W. Humphreys. Blackwell. [aSL]

(1999) Color, consciousness, and the isomorphism constraint. *Behavioral and Brain Sciences* 22:1–21. [aSL]

Parks R. W., Levine D. S. & Long D., ed. (1998) *Fundamentals of neural network modeling.* MIT Press. [WD]

Penrose, L. & Penrose, R. (1958) Impossible objects: A special type of visual illusion. *British Journal of Psychology* 49:31–33. [JH]

Pessoa, L., Thompson, E. & Noë, A. (1998) Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences* 21:723–802. [aSL]

Peterson, M. A. (1994) Shape recognition can and does occur before figure-ground organization. *Current Directions in Psychological Science* 3:105–11. [JH]

Peterson, M. A. & Gibson, B. S. (1993) Shape recognition contributions to figure-ground organization in three-dimensional display. *Cognitive Psychology* 25:383–429. [JH]

Petitot, J. (1999) Morphological eidetics. In: *Naturalizing phenomenology. Issues in contemporary phenomenology and cognitive science,* ed. J. Petitot, F. Varela, B. Pachoud, & J. M. Roy. Stanford University Press. [VR]

Pinker, S. (1980) Mental imagery and the third dimension. *Journal of Experimental Psychology* 109:354–71. [aSL]

(1988) A computational theory of the mental imagery medium. In: *Cognitive and neuropsychological approaches to mental imagery,* ed. M. Denis, J. Engelkamp & J. T. E. Richardson. Martinus Nijhoff. [aSL]

Poincaré, H. (1905/2001) *The value of science: essential writings of Henri Poincaré.* Modern Library. [VR]

Previc, F. H. (1998) The neuropsychology of 3-D space. *Psychological Bulletin* 124:123–64. [HER]

Pribram, K. H. (1971) *Languages of the brain: Experimental paradoxes and principles in neuropsychology.* Brandon House. [MV]

Price, H. H. (1932) *Perception.* Methuen. [aSL]

Pritchard, R. M. (1961) Stabilized images on the retina. *Scientific American* 204:72–78. [BD]

Pritchard, R. M., Heron, W. & Hebb, D. O. (1961) Visual perception approached by the method of stabilized images. *Canadian Journal of Psychology* 14:67–77. [BD]

Raizada, R. D. S. & Grossberg, S. (2001) Context-sensitive binding by the laminar circuits of V1 and V2: A unified model of perceptual grouping, attention, and orientation contrast. *Visual Cognition* 8:431–66. [SG]

(2003) Towards a theory of the laminar architecture of cerebral cortex: Computational clues from the visual system. *Cerebral Cortex* 13:100–13. [SG]

Ramachandran, V. S. (1992) Filling in gaps in perception: Part 1. *Current Directions in Psychological Science* 1:199–205. [aSL]

Ramachandran, V. S. & Anstis, S. M. (1986) The perception of apparent motion. *Scientific American* 254:80–87. [aSL]

Randrup, A. (1997) An alternative to materialism. *Cybernetics and Human Knowing* 4:15–24. [AR]

(2002) *What is real? Conscious experience seen as basic to ontology. An overview* [On-line]. Available from http://cogprints.ecs.soton.ac.uk/archive/00002660/01/reality.html and http://www.mobilixnet.dk/~mob79301/reality.html. [AR]

(submitted) An idealist approach to the study of evolution and cognition. *Evolution and Cognition.* [AR]

Revonsuo, A. (1995) Consciousness, dreams, and virtual realities. *Philosophical Psychology* 8:35–58. [aSL, AR, MV]

(1998) Visual perception and subjective visual awareness. Commentary on Pessoa et al. (1998). *Behavioral and Brain Sciences* 21:769–70. [aSL, AR]

(2000) Prospects for a scientific research program on consciousness. In: *Neural correlates of consciousness,* ed. T. Metzinger. MIT Press. [AR]

(2001a) Can functional brain imaging discover consciousness in the brain? *Journal of Consciousness Studies* 8:3–23. [AR]

(2001b) Dreaming and the place of consciousness in nature. *Behavioral and Brain Sciences* 24:1000–1001. [AR]

Richards, W. (1977) Lessons in constancy from neurophysiology. In: *Stability and constancy in visual perception,* ed. W. Epstein. Wiley. [HER]

Riehle, A., Grammont, F., Diesmann, M. & Grün, S. (2000) Dynamical changes and temporal precision of synchronized spiking activity in monkey motor cortex during movement preparation. *Journal of Physiology (Paris)* 94:569–82. [WAM]

Rock, I. & Brosgole, L. (1964) Grouping based on phenomenal proximity. *Journal of Experimental Psychology* 67:531–38. [aSL]
(1983) *Logic of perception.* MIT Press. [SM]
Roland, P. E. (2002) Dynamic depolarization fields in the cerebral cortex. *Trends in Neurosciences* 25:183–90. [WAM]
Rosenberg, G. (2003) *A place for consciousness: The theory of natural individuals.* Oxford University Press (forthcoming). [aSL]
Rosenthal, V. & Visetti, Y. M. (2003) *Köhler.* Les Belles Lettres. [VR]
Ross, H. E. & Plug, C. (2002) *The mystery of the moon illusion: Exploring size perception.* Oxford University Press. [HER]
Ross, W., Grossberg, S. & Mingolla, E. (2000) Visual cortical mechanisms of perceptual grouping: Interacting layers, networks, columns, and maps. *Neural Networks* 13:571–88. [SG]
Rubin, E. (1921) *Visuell wahrgenommene Figuren.* Gyldendalska. [DL]
Rubin, N. (2001) Figure and ground in the brain. *Nature Neuroscience* 4:857–58. [WAM]
Ruse, M. (1986) *Taking Darwin seriously.* Basil Blackwell. [AR]
Russell, B. (1921) *The analysis of mind.* George Allen & Unwin. [aSL]
(1927) *Philosophy.* W. W. Norton. [arSL]
Ryle, G. (1949) *The concept of mind.* Hutchinson's University Library. [DL]
(1949/1966) *The concept of mind.* Penguin Books. [EW]
Sacks, O. (1985) *The man who mistook his wife for a hat.* Harper & Row. [aSL]
Schirillo, J., Reeves, A. & Arend, L. (1990) Perceived lightness, but not brightness, of achromatic surfaces depends on perceived depth information. *Perception and Psychophysics* 48:82–90. [JAS]
Schwartz, E. L. (1980) Computational anatomy and functional architecture of striate cortex: A spatial mapping approach to perceptual coding. *Vision Research* 20:645–69. [HER]
Searle, J. R. (1980) Minds, brains, and programs. *Behavioral and Brain Sciences* 3:417–24. [aSL]
(1992) *The rediscovery of mind.* The MIT Press. [aSL]
(1997) *The mystery of consciousness.* New York Review. [aSL]
Sellars, R. W. (1916) *Critical realism: A study of the nature and conditions of knowledge.* Rand McNally. [aSL, EW]
(1922) *Evolutionary naturalism.* Open Court. [EW]
(1970) *The principles, perspectives and problems of philosophy.* Pageant Press International. [EW]
Shannon, C. E. (1948) A mathematical theory of communication. *Bell Systems Technical Journal* 27:379–423. [aSL]
Shaw, R. E. & Bransford, J. (1977) Introduction: Psychological approaches to the problem of knowledge. In: *Perceiving, acting and knowing,* ed. R. Shaw & J. Bransford. Erlbaum. [SM]
Shaw, R. E. & Turvey, M. T. (1981) Coalitions as models for ecosystems: A realist perspective on perceptual organization. In: *Perceptual organization,* ed. M. Kubovy & J. R. Pomerantz. Erlbaum. [SM]
Shepard, R. N. (1981) Psychophysical complementarity. In: *Perceptual organization,* ed. M. Kubovy & J. Pomeranz. Erlbaum. [aSL]
(1987) Toward a universal law of generalization for psychological science. *Science* 237:1317–23. [WD]
(1994) Perceptual-cognitive universals as reflections of the world. *Psychonomic Bulletin and Review* 1:2–28 (1994). Reprinted in: *Behavioral and Brain Sciences* 24(4):581–601. Special Issue on the work of Roger Shepard (2001). [WD]
Shepard, R. N. & Chipman, S. (1970) Second-order isomorphism of internal representations: Shapes of states. *Cognitive Psychology* 1:1–17. [aSL]
Shepard, R. N. & Metzler, J. (1971) Mental rotation of three-dimensional objects. *Science* 171:701–703. [aSL]
Simons, D. J., ed. (2000) *Change blindness and visual memory: A special issue of the Journal of Visual Cognition.* Psychology Press. [DLl]
Singer, W. (1999) Neuronal synchrony: A versatile code for the definition of relations? *Neuron* 24:49–65. [aSL, SM]
Singer, W. & Gray, C. (1995) Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience* 18:555–86. [aSL]
Singh, M. & Hoffman, D. D. (1998) Active vision and the basketball problem. Commentary on Pessoa et al. (1998). *Behavioral and Brain Sciences* 21:772–73. [aSL]
Smythies, J. R. (1989) The mind-brain problem. In: *The case for dualism,* ed. J. R. Smythies & J. Beloff. University of Virginia Press. [aSL]
Solomon, S. G., White, A. J. & Martin, P. R. (2002) Extraclassical receptive field properties of parvocellular, magnocellular, and koniocellular cells in the primate lateral geniculate nucleus. *Journal of Neuroscience* 22:338–49. [NPM]
(1994) *The walls of Plato's Cave: The science and philosophy of brain, consciousness, and perception.* Avebury. [aSL]
Stettler, D. D., Das, A., Bennett, J. & Gilbert, C. D. (2002) Lateral connectivity and contextual interactions in macaque primary visual cortex. *Neuron* 36:739–50. [NPM]
Stuart, G. W., Bossomaier, T. R. J. & Johnson, S. (1993) Preattentive processing of object size: Implications for theories of size perception. *Perception* 22:1175–93. [HER]
Swaminathan, G. & Grossberg, S. (2001) Laminar cortical circuits for the perception of slanted and curved 3D surfaces. *Society for Neuroscience Abstracts* Vol. 619, No. 49. [SG]
Takeichi, H., Watanabe, T. & Shimojo, S. (1992) Illusory occluding contours and surface formation by depth propagation. *Perception* 21:177–84. [aSL]
Tausch, R. (1954) Optische Täyschungen als artifizielle Effekte der Gestaltungs-prozesse von Grössen und Formenkonstanz in der natürlichen Raumwahrnehmung. *Psychologische Forschung* 24:299–348. [aSL]
Taylor, J. G. (1999) *The race for consciousness.* MIT Press. [WD]
Thau, M. (2002) *Consciousness and cognition.* MIT Press. [EW]
Thiele, A. & Stoner, G. (2003) Neuronal synchrony does not correlate with motion coherence in cortical area MT. *Nature* 421:266–70. [WAM]
Thouless, R. H. (1931a) Phenomenal regression to the "real" object. I. *British Journal of Psychology* 21:339–59. [DL]
(1931b) Phenomenal regression to the "real" object. II. *British Journal of Psychology* 22:1–30. [DL]
Trehub, A. (1991) *The cognitive brain.* MIT Press. [HER]
Tse, P. U. (1999a) Illusory volumes from conformation. *Perception* 27:(8):977–94. [aSL]
(1999b) Volume completion. *Cognitive Psychology* 39:37–68. [aSL, PUT]
Tse, P. U. & Logothetis, N. K. (2002) The duration of 3-D form analysis in transformational apparent motion. *Perception and Psychophysics* 64:244–65. [PUT]
Vallar G. (1998) Spatial hemineglect in humans. *Trends in Cognitive Sciences* 2(3):87–96. [aSL]
Velmans, M. (1990) Consciousness, brain, and the physical world. *Philosophical Psychology* 3:77–99. Available from http://cogprints.ecs.soton.ac.uk/archive/00000238/. [aSL,MV]
(1991a) Consciousness from a first-person perspective. *Behavioral and Brain Sciences* 14:702–26. Available from http://cogprints.ecs.soton.ac.uk/archive/00000594/. [MV]
(1991b) Is human information processing conscious? *Behavioral and Brain Sciences* 14:651–69. Available from http://cogprints.ecs.soton.ac.uk/archive/00000593/. [MV]
(1993) A reflexive science of consciousness. In: *Experimental and theoretical studies of consciousness. CIBA Foundation Symposium 174.* Wiley. Available from http://cogprints.ecs.soton.ac.uk/archive/00000243/. [MV]
(1996) Consciousness and the "causal paradox." *Behavioral and Brain Sciences* 19:537–42. Available from http://cogprints.ecs.soton.ac.uk/archive/00000596/. [MV]
(2000) *Understanding consciousness.* Routledge/Psychology Press. [MV]
(2001) A natural account of phenomenal consciousness. *Communication and Cognition* 34:39–59. Available from http://cogprints.ecs.soton.ac.uk/archive/00001813/. [MV]
(2002a) How could conscious experiences affect brains? *Journal of Consciousness Studies* 9:3–29. http://cogprints.ecs.soton.ac.uk/archive/00002750/. [MV]
(2002b) Making sense of the causal interactions between consciousness and brain (a reply to commentaries). *Journal of Consciousness Studies* 9:69–95. http://cogprints.ecs.soton.ac.uk/archive/00002751/. [MV]
(2003) Is the world in the brain, or the brain in the world? (Unabridged version of this *Behavioral and Brain Sciences* paper) http://cogprints.ecs.soton.ac.uk/archive/00002798/commentary). [MV]
von Ehrenfels, C. (1890) Über Gestaltqualitäten. *Vierteljahresschrift für wissenschaftliche Philosophie* 14:249–92. [BD]
Ware C. & Kennedy, J. M. (1978) Perception of subjective lines, surfaces and volumes in 3-dimensional constructions. *Leonardo* 11:111–14. [aSL]
Westheimer, G. & Levi, D. M. (1987) Depth attraction and repulsion of disparate foveal stimuli. *Vision Research* 27:1361–68. [aSL]
Wittgenstein, L. (1953) *Philosophical investigations.* Blackwell. [DAB]
Wright, E. L. (1990) Two more proofs of present qualia. *Theoria* 60:3–22. [EW]
(1992) The entity fallacy in epistemology. *Philosophy* 67:33–50. [EW]
(1996) What it isn't like. *American Philosophical Quarterly* 33:23–42. [EW]
Wuketits, F. M. (1984) *Concepts and approaches in evolutionary epistemology.* Reidel. [AR]
Yarbus, A. L. (1967) *Eye movements and vision.* Plenum Press. [aSL]
Zeki, S. & Bartels, A. (1998) The asynchrony of consciousness. *Proceedings of the Royal Society (London) Series B* 265:1583–85. [WAM]
Zucker, S. W., David, C., Dobbins, A. & Iverson, L. (1988) The organization of curve detection: Coarse tangent fields and fine spline coverings. In: *Proceedings of the Second International Conference on Computer Vision, IEEE Computer Society, Tampa FL,* pp. 568–77. http://www.cs.yale.edu/homes/vision/zucker/publications_full.html [aSL]