

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 resur-

faced 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 volu-

metric 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 inputted 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 Leshner 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 & Todorovič 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 & Todorovič 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 im-

age 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 in-

formation 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 cir-

cumscribe 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 phenomenologi-

cally 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 dis-




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.

tinguish differences of perceived depth than differences of perceived height and width.

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 remain 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 argu-

ment 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 percip-

ient. 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 phenome-

nal 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.

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 sur-



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.

face. 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 solip-

sism, 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 transitional 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 com-

puter 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 dangle 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 surrealist paintings by Salvador 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-di-




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.

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.

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.

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 de-




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.

picted 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 shrunk 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 ex-

PLICIT 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

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

where v 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 $\rho = (\pi - v)$ and ρ ranges from 0 at $r = 0$ to π at $r = \text{infinity}$.

What does this kind of compression mean in an isomor-

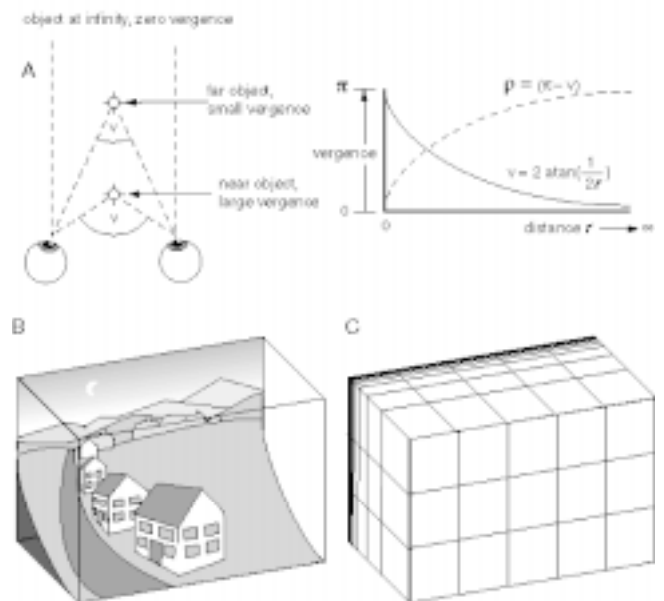


Figure 13. **A.** A vergence representation maps infinite distance into a finite range. **B.** This produces a mapping reminiscent of a museum diorama. **C.** The compressed reference grid in this compressed space defines intervals that are perceived to be of uniform size.




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 neuro-computation.

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 neuro-computation 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 neuro-computation 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 in-

clined 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 stim-

ulus, 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 am-

biguous 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) re-

ported 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.

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; Yarus 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

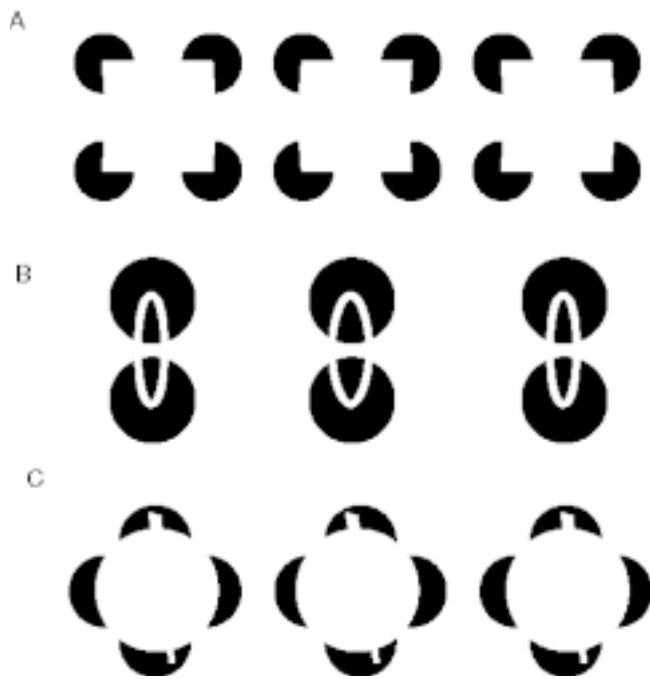


Figure 16. Perceptual interpolation in depth in illusory figure stereograms. Panel **A** was adapted from Carman et al. (1992), **B** from Kellman and Shipley (1991), and **C** from Idesawa (1991). Opposite disparity percepts are achieved by binocular fusion of either the first and second, or the second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

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 ab-

straction 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 $F_{\alpha r} = e^{-r^2} \sin(\alpha)$ 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 $F_{\alpha r} = e^{-r^2} \sin(\alpha)^P$, 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 $F_{\alpha r} = e^{-r^2} 2 \sin(\alpha)^P - 1$, 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

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

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 $G_{\alpha \beta r} = e^{-r^2} \sin(\alpha)^P$. 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 $\text{sgn}(\cdot)$) of a cosine on the orthogonal vector; that is, $G_{\alpha\beta r} = e^{-r^2} \sin(\alpha)^P \text{sgn}(\cos(\beta))$. 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

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

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 $H_{\alpha\beta r} = e^{-r^2} \sin(\alpha)^P \text{sgn}(\cos(\beta))$ – but then adding another similar plane from the orthogonal surface normal, producing the equation $H_{\alpha\beta r} = e^{-r^2} [\sin(\alpha)^P \text{sgn}(\cos(\beta)) + \sin(\beta)^P \text{sgn}(\cos(\alpha))]$. 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

$$H_{\alpha\beta r\theta} = e^{-r^2} [\sin(\alpha)^P \text{sgn}(\cos(\beta)) |\cos(\theta)^Q| + [\sin(\beta)^P \text{sgn}(\cos(\alpha)) |\cos(\theta)^Q|]. \quad (\text{Eq. 3})$$

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 $E_{\alpha\beta r} = e^{-r^2} [\sin(\alpha)^P \sin(\beta)^P]$. 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

$$E_{\alpha\beta r\theta\phi} = e^{-r^2} [\sin(\alpha)^p \sin(\beta)^p] \cos(\theta)^q \cos(\phi)^q, \quad (\text{Eq. 4})$$

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

$$I_{\alpha\beta r\theta\phi} = e^{-r^2} [\sin(\alpha)^p \sin(\beta)^p] \cos(\theta)^q \cos(\phi)^q \quad (\text{Eq. 5})$$

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

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

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

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

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 sug-

gested 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