The Dimensions of Conscious Experience: A Quantitative Phenomenology

Steven Lehar

slehar@vision.eri.harvard.edu

Submitted to Journal of Consciousness Studies June 2000

Rejected April 2001

Summary

Psychology was originally formulated as the science of the *psyche*, i.e. the subjective side of the mind / brain barrier. However time and again it has been diverted from this objective in the supposed interest of scientific rigor. The Behaviorists proposed to transform psychology to a science of behavior, and today the Neuroreductionists propose to transform it to a science of neurophysiology. In the process they attempt to deny the very existence of conscious experience as valid object of scientific scrutiny. However the subjective conscious experience is a primary source of evidence for the nature of the representation in the brain. I propose a quantitative phenomenolgy to express the dimensions of conscious experience in information theoretic terms. This approach leads to interesting observations of the properties of phenomenal perspective, that clearly reveal the phenomenal world as an internal rather than external entity.

Introduction

Psychology was originally formulated as the science of the *psyche*, i.e. the scientific investigation of the subjective side of the mind / brain barrier. The study of conscious experience therefore is the principal domain of psychology, and neurophysiology only enters into the picture to provide a physical substrate for mind. Over the years two distinct methodologies have emerged for the study of conscious experience, *phenomenology* and *psychophysics*. Each approach has its particular strengths and weaknesses, although in recent decades the phenomenological approach has fallen largely into disuse and even a certain disrepute. The reasons for this are complex, but ultimately unjustified (Varela & Shear 1999), for phenomenology is not only a perfectly valid and legitimate method of investigation, it is an indispensable component of the science of psychology. The recent neglect of this avenue of investigation has led psychology far astray from the original goal of a scientific study of conscious experience. It was this diversion of psychology from its original objective that has necessitated the creation of the new science of consciousness studies.>

There is a subtle distinction between the terms phenomenology and introspection, which are otherwise almost synonymous. Phenomenology is the older term, and involves a temporary suspension of belief in the objective world in order to concentrate on the phenomenon of experience in and of itself. Introspection involves contemplation of one's own thoughts, feelings, and sensations. However the latter term has become associated with the Introspectionist, or Structuralist movement (Wundt 1897, Titchener 1928) whose stated objective was to study the structure of conscious experience. This more specific variation of Introspection (capitalized here to distinguish it from introspection in general) marked the birth of psychology as a science distinct from philosophy and biology. However from the outset certain methodological problems emerged which have contributed to its current demise. There has always been a problem of objectivity with phenomenological observation, for the observer cannot help but focus on aspects of the conscious experience which are consistent with their preconceived notions or theoretical expectations. This does not invalidate the practice when exercised with the proper caution, but it is a constant danger to which the phenomenological observer must remain eternally vigilant. The classical Introspectionists began with the assumption that all levels of perceptual processing should be available to introspective observation, from the raw sensory experience through the higher cognitive understanding of that experience. They expected therefore that the most basic or elemental components of consciousness should be similar dimensionally to the raw sensory input. In order to begin the investigation from the bottom up, the Introspectionists placed great emphasis on training the Introspective observer to focus specifically on the elements of the conscious experience, and to ignore the higher order inferences by which they were supposedly bound. For example the retinal image is known to be two-dimensional, and therefore the Introspective observer was trained to ignore the three-dimensional component of the visual experience and to focus on its two-dimensional projection. However this kind of Intospection is not a true introspection, because the results depend on those initial assumptions. For even a casual inspection reveals the world of perception to be fundamentally three-dimensional, and it is only by an effort of self-delusion that it can possibly be seen as primarily a two- dimensional projection.

The Gestalt movement introduced a new approach to phenomenology that relied heavily on the visual illusions observed when viewing certain graphical figures. The principal finding of Gestalt theory was that the perceptual experience due to a figure is not a simple sum of the effects of its individual components, but that global configural factors are involved in perception that defy the kind of atomistic or elemental analysis proposed by the Introspectionists. In particular, Gestalt theory demonstrated that the world of perception is fundamentally three-dimensional, and that therefore the third dimension was not due to any kind of cognitive inference based on two-dimensional elements, but the three-dimensional nature of the percept is primary, and is manifest even at the lowest or most immediate levels of conscious experience. It takes considerable practice for an artist to learn to extract the two-dimensional projection from a three-dimensional scene, as required to produce a perspective sketch, because it is very difficult to ignore the three- dimensional component of the visual world. It is the two-dimensional projection therefore that must be inferred from the three-dimensional percept, rather than the other way round. The great disparity between the conscious experience and the sensory input on which it is based suggests that there is a considerable amount of visual processing that occurs below the level of consciousness, where the two-dimensional sensory input is transformed into the three-dimensional percept. Despite these significant achievements of Gestalt phenomenology, the disparity between the phenomenal observations of these different schools of psychology eventually led to a general indictment of phenomenology and introspection as hopelessly subjective and impossible to verify objectively.

The methodology of psychophysics was devised specifically to avoid the problem of observer subjectivity by using naïve subjects who are not informed of the theoretical significance of the experiment, with perceptual tasks limited to simple discriminations or judgements that can be reduced to a simple keypress response. Theoretical techniques have been borrowed from signal detection theory in order to eliminate any residual bias due to experimental protocol, and statistical techniques have been refined for rigorous analysis of the results. However the psychophysical method suffers from a major limitation, which is that the rich conscious experience simply cannot be communicated through the bottleneck of the simple keypress response. As a result, an impoverished view of the nature of perceptual processing has emerged in contemporary psychology that ignores some of the most prominent aspects of perception which are plainly manifest in the subjective conscious experience.

Despite these problems a great deal of useful knowledge has been gained both from the phenomenological and the psychophysical approaches. The two methodologies are somewhat complementary, for phenomenology offers initial insights into the general nature of conscious processes, while psychophysics offers a more objective means to settle issues of contention raised by the phenomenological observations. The one is more qualitative, and offers a big-picture overview, while the other is more quantitative, and measures the particulars with greater precision and reliability. The strong bias seen in recent decades in favor of the quantitative approach is a reflection of a general `physics envy' in psychology that attempts to avoid the fuzzy subjective aspect of the science in order to place psychology on a more rigorous scientific footing.

It was this same pursuit of scientific rigor that motivated the Behaviorist movement, for there is more certainty and objectivity in behavior observed from the outside than in subjective reports of the conscious experience from the inside. But in the pursuit of scientific rigor, the Behaviorists threw out the baby with the bath water, and the most extreme proponents of Behaviorism even denied the very existence of a subjective conscious experience. However a psychology that fails to account for the conscious experience is a psychology that essentially explains nothing, for it has lost sight of the original objectives of the science. Fortunately the excesses of the Behaviorist movement are now generally recognized, and the cognitive revolution has begun to swing the pendulum back towards a recognition of the validity of a science of conscious experience. The lesson from the Behaviorists' error is that the troublesome philosophical issues underlying the study of conscious experience cannot be avoided by simply ignoring them. Rather, these issues must be clearly identified and confronted head-on.

But we are not yet out of the woods, for once again the science of psychology is being diverted from its primary objective, again in the interests of supposed scientific rigor. This time the lure has been advances in neurophysiological recording technology, especially the single-cell electrode. This again seems to promise a more objective third-person approach to the investigation of the elusive psyche. However the picture of sensory processing revealed by neurophysiology is so divergent from the dimensions of conscious experience as observed phenomenologically, that once again we are exhorted by modern Neuroreductionists to deny the very existence of the conscious experience as being some kind of illusion which has no objective reality (Dennett 1991). However the real motivation behind this denial of consciousness is ultimately neurophysiological, for the properties of consciousness are a constant embarrassment to our neurophysiological theories of sensory processing. It would be very convenient for neural network theorists if consciousness were simply excluded from the set of data to be explained by psychology.

The Epistemological Divide

There is a central philosophical confusion underlying discussions of phenomenal experience, which explains why the very existence of conscious experience is so often denied as an objective reality. 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 *naive realism*, as opposed to *indirect realism*, or *representationalism*. Although this issue is not much discussed in contemporary psychology, it is an old debate that has resurfaced several times, but the continued 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.

Problems with Direct Realism

The direct realist view (Gibson 1972) 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 which 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 it stimulates 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. The difficulty with the concept of direct perception is most clearly seen when considering 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 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 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.

Problems with Indirect Realism

The indirect realist view is also incredible, for it suggests that the solid stable structure of the world that we perceive to surround us is merely a pattern of energy in the physical brain, i.e. that 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 it out beyond your direct experience. 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 concept of perception is so disorienting that only a handful of researchers have seriously entertained this notion or pursued its implications to its logical conclusion. (Kant 1781/1991, Koffka 1935, Köhler 1971 p. 125, Russell 1927 pp 137-143, Smythies 1989, 1994, Harrison 1989, Hoffman 1998)

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

Problems with Projection Theory

There is a third alternative besides the direct and indirect realist views, and that is a *projection theory*, whereby the brain does indeed process sensory input, but that the results of that processing get somehow projected back out of the brain to be superimposed back on the external world (Ruch 1950 quoted in Smythies 1954, O'Shaughnessy 1980 pp 168- 192, Velmans 1990, Baldwin 1992). According to this view, the world around us is part real, and part perceptual construction, and the two are spatially superimposed. However no physical mechanism has ever been proposed to account for this external projection. The problem with this notion becomes clear when considering how an artificial intelligence could possibly be endowed with this kind of external projection. Although a sensor may record an external quantity in an internal register or variable in a computer, there is no sense in which that internal value can be considered to be external to that register or to the physical machine itself, whether detected externally with an electrical probe, or examined internally by software data access. Unless the principle of external projection can be demonstrated in a simple artificial sensory system, this explanation too remains as mysterious as the property of consciousness it is supposed to explain.

Selection from Incredible Alternatives

We are left therefore with a choice between 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, while 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 reflect the truth of visual processing. For 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.

The problem with the direct realist view is of an epistemological nature, and is therefore a more fundamental objection, for direct realism 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 that neural processes in our brain are somehow also out in the world. Both of these paradigms have difficulty with phenomena of dreams and hallucinations (Revonsuo 1995), which present the same kind of phenomenal experience as spatial vision, except independently 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 that has led many to conclude that consciousness is deeply mysterious and forever beyond human comprehension. For example Searle (1992) contends 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; that there is no distinction between the observation and the thing observed.

The problem with the indirect realist view on the other hand is more of a technological or computational limitation, for we cannot imagine how contemporary concepts of neurocomputation, or even artificial computation for that matter, can account for the properties of perception as observed in visual consciousness. It is clear however that we have yet to discover the most fundamental principles of neural computation and representation, 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. For in that case 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 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 that either we accept a magical mysterious account of perception and consciousness that seems impossible in principle to implement in any artificial vision system, or we have to face the seemingly incredible truth that the world we perceive around us is indeed an internal data structure within our physical brain. The principal focus of neurophysiology should now be to identify the operational principles behind the three-dimensional volumetric imaging mechanism in the brain, that is responsible for the generation of the solid stable world of visual experience that we observe to surround us in conscious experience.

A Perceptual Modeling Approach

The divergence between the neurophysiological and phenomenal descriptions of visual processing is so great, that it is hard to know where to even begin in the attempt to model visual experience. The elements of neurophysiology appear to be neurons that generate pulses of electrical activity that are transmitted across the chemical synapse. The elements of conscious experience appear to be solid volumetric objects bounded by colored surfaces, embedded in a three-dimensional void of perceived space. There is a dimensional mismatch between these bottom-up and top-down descriptions of visual representation, that makes it impossible to model conscious experience in neurophysiological terms with any confidence that the model has any validity in the brain. For until the mapping between subjective experience and the corresponding neurophysiological state has been identified beyond question, it is impossible to verify whether the neural model has correctly replicated any particular phenomenal experience.

Nagle (1974) suggests that we set aside temporarily the relation between mind and brain and devise a new method of *objective phenomenology,* i.e. to quantify the structural features of the subjective experience in objective terms, without committing to any particular neurophysiological theory of perceptual representation. Chalmers (1995) extends this line of reasoning with the observation that the subjective experience and its corresponding neurophysiological state carry the same information content. Chalmers therefore proposes 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 cognition do not float free of one another, but cohere in an intimate way. 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 sufficiently specified as to be measurable in any physical system given that the coding scheme is known. A similar argument is made by Clark (1993, p. 50).

I propose therefore a *perceptual modeling* approach, i.e. a quantitative model of the content of conscious experience, expressed in the subjective variables of perceived color, shape, and motion, rather than in the neurophysiological variables of neural activations or spiking frequencies etc. The variables encoded in the perceptual model correspond to 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 that 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 the 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 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) presents 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 side-stepped, 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.

While this is of course only an interim solution, for eventually the neurophysiological basis of conscious experience must also be identified, the perceptual model does offer objective information about the informational content encoded in the physical mechanism of the brain. This is a necessary prerequisite to a search for the neurophysiological basis of conscious experience, for we must clearly circumscribe the *explanandum* before we can attempt an *explanans*. 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 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), and all of the visible points on a perceived object like 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 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. Since perceived color is expressed in the three dimensions of hue, intensity, and saturation, the perceived world can be expressed as a six-dimensional manifold (Clark 1993), with three spatial and three color dimensions. 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 apparent in perception, as a necessary prerequisite to the search for a neurophysiological mechanism that can perform that same transformation.

Once we recognize the world of experience for what it really is, it becomes clearly evident that the representational strategy used by the brain is an *analogical* one. In other words, objects and surfaces are represented in the brain not by an abstract symbolic code, as suggested in the propositional paradigm, nor are they encoded by the activation of individual cells or groups of cells representing particular features detected in the scene, as suggested in the neural network or feature detection paradigm. Instead, objects are represented in the brain by constructing full spatial effigies of them that appear to us for all the world like the objects themselves- or at least so it seems to us only because we have never seen those objects in their raw form, but only through our perceptual representations of them. Indeed the only reason why this very obvious fact of perception has been so often overlooked is because the illusion is so compelling that we tend to mistake the world of perception for the real world of which it is merely a copy. This is a classic case of not seeing the forest for the trees, for the evidence for the nature of perceptual representation in the brain has been right before us all along, cleverly disguised as objects and surfaces in a virtual world that we take to be reality. So for example when I stand before a table, the light reflected from that table into my eye produces an image on my retina, but my conscious experience of that table is not of a flat two-dimensional image, but rather my brain fabricates a three-dimensional replica of that table carefully tailored to exactly match the retinal image, and presents that replica in an internal perceptual space that includes a model of my environment around me, and a miniature copy of my own body at the center of that environment. The model table is located in the same relation to the model of my body as the real table is to my real body in external space. The perception or consciousness of the table therefore is identically equal to the appearance of the effigy of the table in my perceptual representation, and the experience of that internal effigy is the closest I can ever come to having the experience of the physical table itself.

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. However this argument is invalid, for there is no need for an internal observer of the scene, since the internal representation is simply a data structure like any other data in a computer, except that this data is expressed in spatial form. If the existence of a spatial data structure required a homunculus to view it, the same objection would also apply to symbolic or verbal information in the brain, which would also require a homunculus to read or interpret that data. 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. To deny the spatial nature of the perceptual representation 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 properties as I experience them to have, i.e. location, spatial extension, color, and shape. I think them, therefore they exist. 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 in my internal perceptual world is beyond question if I experience them, even if only as a hallucination.

The Neuroreductionist Objection

A number of theorists have proposed (Dennett 1991, 1992, O'Regan 1992, Pessoa et al. 1998) that consciousness is an illusion, and that in fact the conscious experience is considerably more impoverished than it appears subjectively. For example the loss of resolution in peripheral vision is not immediately apparent to the naïve observer. However the objective of perceptual modeling is not to quantify the casual experience of the naïve observer, but the careful observation of the critical observer. For the loss of acuity in peripheral vision is plainly evident under phenomenological observation, and can be easily verified psychophysically, and therefore this should also be reflected in the perceptual model. Dennett argues that visual information need not be encoded explicitly in the brain, but merely implicitly in some kind of compressed representation. For example the percept of a surface with uniform color could be abbreviated to a kind of edge image, with a single value to encode the color of the whole surface, as is the practice in image compression algorithms. This notion appears to be supported by neurophysiological studies of the retina which show that ganglion cells respond only to spatial or temporal discontinuities of the brightness profile, with no response within regions of uniform color or brightness. Dennett argues that the experience of a filled-in field of color in uniform fields, and in the blind spot, does not suggest an explicit filling-in mechanism in the brain, but that the color experience is encoded by "ignoring an absence" (Dennett 1991,1992). However an absence can only be ignored from a representation that already contains something in the place of the ignored item, otherwise one would experience nothing at all, rather than a spatially continuous field of color. In fact the experience of the retinal blind spot, or a uniformly colored surface, produces a distinct colored experience at every point throughout the colored region to a particular spatial resolution as a spatial continuum, and the informational content of that experience is greater than that in a compressed representation. If it is true that the retinal image encodes only brightness transitions at visual boundaries, then some other mechanism higher up in the processing stream must perform an explicit filling-in to account for the subjective experience of the filled-in surface. In fact the many illusory filling-in phenomena such as the Kanizsa illusion implicate exactly this kind of mechanism in perception. If it were sufficient for the brain to encode visual information only implicitly in some kind of compressed code, then there would be no need to posit any perceptual processing beyond the retina, because the retina already contains an implicit representation of all of the information in the visual scene. If visual information were indeed expressed in a compressed neurophysiological code, then our subjective experience of that information would have to also be correspondingly compressed or abstracted, as is the case for example with an experience of a remembered or imagined scene. The fact that our phenomenal experience is of a filled-in volumetric world is direct and concrete evidence for a volumetric filling-in mechanism in the brain.

This debate highlights the indispensable contribution of phenomenology, for in the absence of the subjective experience of vision there would be no check on theories of visual processing based exclusively on neurophysiological or psychophysical data. In fact the bottom-up approach that works upwards from the properties of the single cell, and the top-down approach that works downwards from the conscious experience, are equally valid and complementary approaches to the investigation of visual processing. Both are essential because each affords a view of the problem from its own unique perspective. If we follow Dennett's lead and ignore phenomenology whenever it is in conflict with contemporary neurophysiology, we end up like the Behaviorists, with a denial of the existence of the very object that psychology set out to explain in the first place. In fact Dennett's contention that consciousness is somehow less than it appears is tantamount to asking us to ignore the evidence of our own eyes. Although there are certain aspects of consciousness which may not be immediately apparent to the casual observer and therefore require more critical phenomenological observation, nevertheless consciousness is by definition exactly as rich and complex as it appears phenomenologically, and no less, even if that appearance happens to be inconvenient for neurophysiological theories of visual representation.

The Function of Conscious Experience

There is much discussion in philosophy about the possible function of conscious experience, and whether it is an epiphenomenon that has no direct functional value. The issue is highlighted with the notion of the hypothetical `zombie' whose behavior as observed externally is identical to that of normal people, except that this zombie supposedly lacks all conscious experience. This notion sounds very peculiar from the indirect realist perspective. For once we accept that the world which appears to be external to our bodies is in fact an internal data structure in our physical brain, the notion of the zombie as proposed becomes a contradiction in terms. For a zombie that does not possess an internal picture of the world around it, could not possibly walk about in the world avoiding obstacles as we do. Without a conscious memory of where it had just been, and a conscious intent of where it would like to go next, the zombie would behave much as we do when we are in an unconscious state, i.e. it would lie inert and immobile, with neither the incentive nor the capacity for action.

The notion of this kind of zombie presupposes a distinction between the structural aspects of the perceived world, which are supposedly a reflection of the objective spatial properties of the world, and the subjective qualia with which those perceptual structures are somehow painted or clothed. This harkens back to an old distinction in psychology between the primary and secondary qualities of perception. Immanuel Kant argued however that the perception of space and time are themselves *a priori* intuitions, i.e. they are a kind of qualium used by the mind to express the structure of external reality. Therefore the fact that the world of experience appears as a volumetric spatial structure is itself an aspect of conscious experience, rather than a veridical manifestation of the true nature of the external world. The phenomenon of hemi-neglect demonstrates that portions of perceived space can completely disappear from consciousness, making it impossible to form either mental or perceptual imagery in that portion of space. It is not just the objects in that space that become invisible, but the very space itself as a potential holder of objects that ceases to exist. This condition clearly indicates the reality of an explicit spatial representation in the brain.

The notion of the hypothetical zombie therefore is impossible in principle, because it is impossible to have any perceptual experience in the absence of some subjective qualia by which that experience is expressed. For qualia are the carriers of the information experienced in perception, just as electromagnetic waves are the carriers of radio and television signals. Again, the notion of information can help clarify the central role of qualia in perceptual representation. For information is defined independent of the physical medium by which it is carried, whether it be electromagnetic radiation, electrical voltages on a wire, or characters on a printed page, etc. However in every case there must be some physical medium to carry that information, for it is impossible for information to exist without a physical carrier of some kind. A similar principle holds on the subjective side of the mind / brain barrier, where the information encoded in perceptual experience is carried by modulations of some subjective qualium, whether it be variations of hue, brightness, saturation, pitch, heat or cold, pleasure or pain, etc. The notion of experience without qualia to support it is as absurd as the notion of information without any physical medium or mechanism to carry that information. Similarly, it is impossible in principle to have a spatial experience in the absence of a spatial representational medium to encode it, as demonstrated most clearly by the phenomenon of hemi-neglect, just as it is impossible for us to experience a four-dimensional space with our three-dimensional perceptual and imagery system. The zombie argument therefore is circular, for it presupposes the possibility of behavior in the absence of experience to demonstrate that behavior and experience are theoretically separable.

The functional purpose of conscious experience therefore is to provide an internal replica of the external world in order to guide our behavior through the world, for otherwise we would have no knowledge of the structure of the world, or of our location within it.

Bounded Nature of the Phenomenal World

The idea of perception as a literal volumetric replica of the world inside your head raises another question of boundedness, i.e. how an explicit spatial representation can 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. This can be seen most clearly in the night sky, where the distant stars produce a dome-like 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, for example from an airplane or a hot air balloon. The dome of the sky above, and the bowl of the earth below therefore define a finite approximately spherical space (Heelan 1983) that encodes distances out to infinity within a representational structure that is both finite and bounded. While 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.

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 therefore in some sense be perceived as being bowed, and yet while bowed, 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, i.e. 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 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, i.e. objects beyond a certain distance lose all perceptual depth. This is seen for example where 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 1 A, which depicts the perceptual representation of a man walking down a road. The phenomenon of perspective is by definition a transformation from a three-dimensional space through a focal point to a two- dimensional surface. The appearance of perspective on the retinal surface therefore is no mystery, and is similar in principle to the image formed by the lens in a camera. What is remarkable in perception is the perspective that is observed not on a two-dimensional surface, but 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.

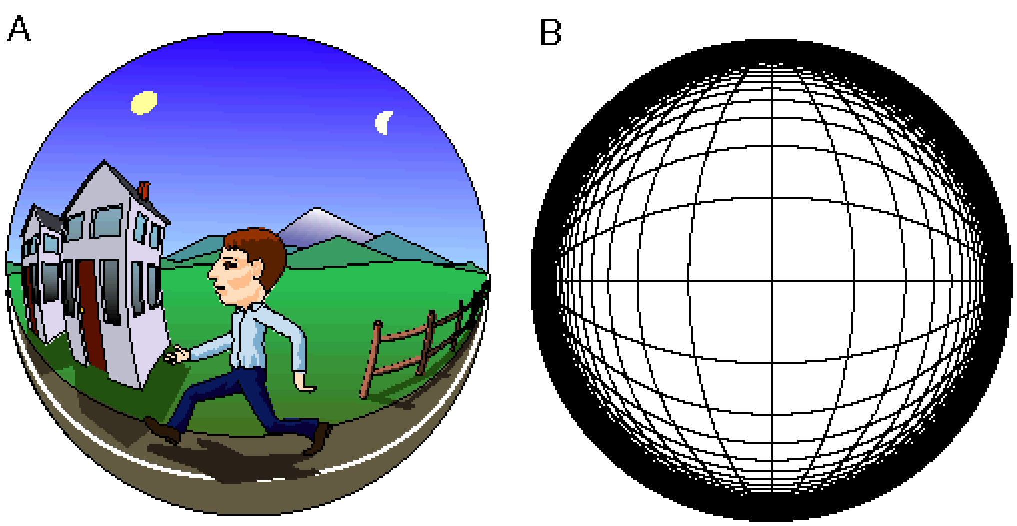

Figure 1

A: 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 yet at the same time they are perceived to converge to a point both up ahead and behind, and that point is perceived to be at a distance which is less than

infinite. B: The deformation of the infinite Cartesian grid caused by the perspective transformation of the perceptual representation. This distorted reference grid is used to make judgements about objective size and straightness in the perceived world.

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 yet at the same time to be undiminished in size. This corresponds to the difference in subjects' reports depending on whether they are given objective v.s. projective instructions (Coren et al., 1994 p. 500) in how to report their observations, showing that 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 objects are represented, so that objects that are warped coherently along with the space in which they are embedded appear undistorted perceptually. This concept is illustrated in figure 1 B that shows how a uniformly spaced Cartesian coordinate system is warped by the three-dimensional perspective transformation that maps points $P(\alpha,\beta,r)$ in Euclidean space, expressed in polar coordinates, to points $Q(\alpha,\beta,\pi-\nu)$ in perceptual space. The variables α and β represent azimuth and elevation angles in a polar coordinate system centered on the percipient, r represents radial distance from the center, and ν represents a vergence measure defined as

> [*Note for HTML version:* If your browser does not load the "Symbol" font, the greek letters will not appear correctly in the text, Pi appears as π, theta appears as θ, sigma appears as σ etc. If you see proper greek letters here, this problem does not apply to you.]

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

In other words, azimuth and elevation angles are preserved through this transformation, while radial distance is compressed by a nonlinear vergence measure that maps the infinity of Euclidean distance into a finite bounded range. This geometrical transformation of the infinite Cartesian grid actually represents a unique kind of perspective transformation on the Cartesian grid, because the transformed space looks like a perspective view of a Cartesian grid when viewed from the inside, with all parallel lines converging to a point in opposite directions, although this perspective is defined in three dimensions rather than two. 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. If the distorted reference grid of figure 1 B is used to measure lines and distances in figure 1 A, the bowed line of the road on which the man is walking is aligned with the bowed reference grid and is therefore perceived to be straight. Therefore 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. Similarly, the walls of the houses depicted in figure 1 A which bow outwards from the observer, conform to the distortion of the reference grid in figure 1 B, and are therefore perceived to be straight and vertical. Likewise, the nearer and farther houses are perceived to be of approximately equal height and depth in objective size, because they span the same number of reference lines in the perspective distorted grid, and yet at the same time the farther house is perceived to be smaller in projective size, as observed also in perspective. This model therefore offers a quantitative description of the otherwise paradoxical phenomenon of the simultaneous experience of objective and projective size in perception. It should be emphasized that figure 1 A does not depict the subjective experience of spatial vision, for we do not see the world of perception as distorted like a fish-eye lens view. Rather, this model represents the informational content of the kind of visual representation that embodies the logical contradictions observed in phenomenal experience, where parallel lines can be perceived to meet at two points in opposite directions which are less than an infinite distance from the observer, while passing to either side of the observer, and at the same time those lines are perceived as straight and parallel throughout their length even where they meet. This data structure is a literal embodiment of the *antinomy*, or logical contradiction built into perception observed by Kant, i.e. that perceived space is at the same time finite and infinite. For the representation truly encodes all of space out to infinity, although in a truly finite representation, using the trick of reducing the resolution all the way to zero resolution in the depth dimension beyond a certain perceived distance. 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. This representational trick is only invisible to us under normal circumstances because we have never seen space the way it really is. In the words of William Blake, "If the doors of perception were cleansed, everything would appear to man as it is, *infinite*".

The Embodied Percipient

This model of spatial representation emphasizes another aspect of perception that is often ignored in models of vision, that 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, for a view of the world around us would be useless if it were 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. As we walk down a street we experience a percept of our feet stepping beneath us, but our bubble of perceptual experience seems to remain centered on our head even as we move through the world. It seems almost as if we are pushing the street backwards with our feet, and this causes the image of the street up ahead to expand outwards from the vanishing point, and arc around us before converging back to a point behind us, all the while following the warped grid lines of our distorted perceptual world. Objects in the perceived world are observed to morph as we walk, swelling up like a balloon as they pass by, only to shrink back down again behind us. And yet since this morphing conforms to the distorted reference grid of our perceived world, we experience those objects as undistorted, as if they were embedded in a Euclidean world. However the distorted nature of perceived space is clearly evident when standing in a long corridor or hallway, whose four corners between the walls, floor, and ceiling expand outwards from a point up ahead, and converge inwards to a point behind, but appear to be straight and parallel throughout their length. This bizarre behavior of the world of experience is clearly not a property of the world itself, but a direct manifestation of the pecular representational scheme used by the brain to present the world to us in conscious experience.

Conclusion

Time and again throughout the history of the investigation of our own sentient nature, movements have arisen to acknowledge conscious experience as an objective reality in its own right, independent of the world of which it is an imperfect replica. The most significant such movement was psychology itself, which took on the challenge of the scientific investigation of the psyche. Time and again these movements have been followed by dark ages that questioned the very existence of consciousness as an objective reality, and sought to redefine psychology as the science of behavior, or neurophysiology, or anything other than the elusive psyche. This reaction is understandable given the enormous complexity and sophistication of conscious experience. We are all born naive realists, and it seems that only a few in each generation ever come to see through the illusion and recognize the world of experience for what it really is. The tendency towards naive realism has always been most urgent and passionate in those circles that undertake the daunting task of providing a neurophysiological or computational explanation of the mechanism of perceptual processing (Gibson 1972, Dennett 1991, Velmans 1990), for the naive realist perspective greatly simplifies our concept of the functional role of perception. However no real progress is made by modeling a simplified or reduced version of perception, especially one that leaves consciousness as something magical and mysterious that seems forever beyond human comprehension. The only way we will ever discover the neurophysiological principles behind perception is by first taking stock of the full measure of the phenomenon of conscious experience, even if that experience appears to violate everything we think we know about neurophysiology. If conscious experience is inconsistent with our notions of biological computation, then it is our neurophysiological theories that are in desperate need of fundamental revision, for the evidence of phenomenology is primary, and offers a direct and reliable view of the informational content of the representational mechanism of the brain. The time has come to restore the original charter of psychology as the science of the psyche, and to begin the search for a neurophysiological theory that is consistent with the observed properties of conscious experience.

References

Baldwin T. (1992) The Projective Theory of Sensory Content. In: T. Crane (Ed.) The Contents of Experience: Essays on Perception. Cambridge UK: Cambridge University Press, 177-195.

Chalmers, D. J. (1995) Facing Up to the Problems of Consciousness. Journal of Consciousness Studies 2 (3) 200-219. Reprinted in "Toward a Science of Consciousness II, The Second Tucson Discussions and Debates". (1996) S. R. Hameroff, A. W. Kaszniak, & A. C. Scott (Eds.) 5-28.

Clark A. (1993) Sensory Qualities. Oxford: Clarendon Press.

Coren S., Ward L. M., & Enns J. J. (1994) Sensation and Perception. Ft Worth TX, Harcourt Brace.

Dennett D. (1991) Consciousness Explained. Boston, Little Brown & Co.

Dennett D. (1992) `Filling In' Versus Finding Out: a ubiquitous confusion in cognitive science. In Cognition: Conceptual and Methodological Issues, Eds. H. L. Pick, Jr., P. van den Broek, & D. C. Knill. Washington DC.: American Psychological Association.

Gibson J. J. (1972) A Theory of Direct Visual Perception. In: The Psychology of Knowing. (J. R. Royce & W. W. Rozeboom (Eds.), Gordon & Breach.

Harrison S. (1989) A New Visualization on the Mind-Brain Problem: Naive Realism Transcended. In J. Smythies & J. Beloff (Eds.) The Case for Dualism. Charlottesville: University of Virginia.

Heelan P. A. (1983) Space Perception and the Philosophy of Science. Berkeley. University of California Press.

Hoffman D. D. (1998) Visual Intelligence: How We Create What We See. New York: W. W. Norton.

Kant I. (1781 / 1991) Critique of Pure Reason. Vasilis Politis (Ed.) London: Dent.

Koffka K, (1935) Principles of Gestalt Psychology. New York, Harcourt Brace & Co.

Köhler W. (1971) A Task For Philosophers. In: The Selected Papers of Wolfgang Koehler, Mary Henle (Ed.) Liveright, New York. pp 83-107.

Kuhn T. S. (1970) The Structure of Scientific Revolutions. Chicago: Chicago University Press.

Nagle T. (1974) What Is It Like to Be a Bat? Philosophical Review 83 435-450

O'Regan K. J., (1992) Solving the `Real' Mysteries of Visual Perception: The World as an Outside Memory. Canadian Journal of Psychology 46 461-488.

O'Shaughnessy B. (1980) The Will: A Dual Aspect Theory. (2 volumes) Cambridge UK: Cambridge University Press.

Pessoa L., Thompson E., & Noë A. (1998) Finding Out About Filling-In: A guide to perceptual completion for visual science and the philosophy of perception. Behavioral and Brain Sciences 21, 723-802.

Revonsuo A. (1995) Consciousness, Dreams, and Virtual Realities. Philosophical Psychology 8 (1) 35-58.

Ruch (1950) In J. F. Fulton (Ed.) Textbook of Physiology, 16th Ed. Philadelphia, p. 311. Pertinent passage quoted in Smythies (1954).

Russell B. (1927) Philosophy. New York: W. W. Norton.

Searle J. R. (1992) The Rediscovery of Mind. Cambridge MA: The MIT Press.

Shannon C. E. (1948) A Mathematical Theory of Communication. Bell Systems Technical Journal 27: 379-423.

Smythies J. R. (1954) Analysis of Projection. British Journal for the Philosophy of Science 5, 120-133.

Smythies J. R. (1989) The Mind-Brain Problem. In: J. R. Smythies & J. Beloff (Eds) The Case For Dualism. Charlottesville: University of Virginia Press.

Smythies J. R. (1994) The Walls of Plato's Cave: the science and philosophy of brain, consciousness, , and perception. Aldershot UK: Avebury.

Titchener E. B. (1928) A Text-Book of Psychology. New York: MacMillan.

Varela F. J. & Shear J. (1999), First-Person Methodologies: What, Why, How? Journal of Consciousness Studies, 6 (2-3), pp. 1-14.

Velmans M. (1990) Consciousness, Brain and the Physical World. Philosophical Psychology 3 (1) 77-99.

Wundt W. (1897) Outlines of Psychology. C. H. Judd (transl.) Leipzig: W. Engleman, New York: G. E. Stechert.