

The Protophenomenal Structure of Consciousness,  
With Especial Application to the  
Experience of Color

Extended Version

Technical Report UT-CS-99-418

Bruce J. MacLennan\*

Computer Science Department  
University of Tennessee, Knoxville  
maclennan@cs.utk.edu

March 1999

Revised April 18, 2002

**Abstract**

This report begins with an overview of the protophenomenal approach to consciousness, which relates the smallest units of subjective experience to electrochemical processes in the brain. The approach is applied to the long-standing philosophical problem of inverted qualia (“Could you tell if your experience of the color spectrum were opposite to mine?”). This problem is critical for an adequate scientific theory of consciousness. The protophenomenal approach is used to show that a spectral inversion is impossible, since abnormal color vision would result in predictable anomalies in the conscious experience of color.

---

\*This report is based on a presentation at the *International Conference on Consciousness in Science and Philosophy '98*, Charleston, IL, November 6–7, 1998. It is in the public domain and may be used for any non-profit purpose provided that the source is credited.

# 1 Introduction

This paper addresses the *principal problem of consciousness*, which is to reconcile our experience of subjective awareness with the scientific world view; it is essentially the same as Chalmers’s “Hard Problem” (Chalmers (1995, 1996)). This problem arises because subjective experience has a special epistemological status, since it is the personal (and private) substratum of all observation, whereas empirical science is traditionally based on common (nonpersonal, public) specific observations (MacLennan 1995).

Although, because of this special status, direct reduction of subjective experience to physical principles is impossible, we can use another sort of reduction, for the essence of reduction is an explanation of the more complex in terms of the simpler. This is a quantitative reduction within the domain of phenomenal experience.

It is apparent that in a certain sense our conscious experience is extended and therefore has “parts” or “regions.” For example, visual awareness comprises spatially distributed components or units of visual experience. Similarly one’s sense of touch is experienced as spatially distributed over the body. In general we may understand our experience as constituted of smaller, simpler elements of experience. This approach may sound like the long-discredited idea of elementary sense data (e.g., “red-here-now”), but it is not. The story is certainly much more complicated than this, but we can begin to see how subjective experience might be decomposed into elementary units of experience.

In itself this approach does little to reconcile subjective experience with the scientific worldview (the “principal problem of consciousness”). However, as we begin to identify these elementary units of experience, we find that they correspond quite closely to physical processes at the neural level. Further, we find the relations between the elements of experience, which integrate them into a phenomenal world, correspond to the connections between the neurons. There is thus a close correspondence, perhaps even an isomorphism, between the phenomenological and neurological analyses. Yet this is not a reduction of the phenomenological to the neurological, for that is fundamentally impossible. Rather, it is a parallel reduction of subjective processes into elementary subjective processes, and of neurological processes into simpler neurological processes. In this way the neurological and phenomenological analyses may each illuminate the other.

The phenomenological analysis depends on the observation of one’s own subjective states. Certainly naive introspectionism is treacherous, and psychology did well to abandon it, but phenomenological training permits unbiased (or less biased) analysis of the structure of consciousness. Through phenomenologically trained observers we may acquire unbiased (public) data about the structure of consciousness (MacLennan 1995, 1996a). (We use “phenomenology” and related terms in the sense of Husserl and Heidegger, that is, to refer to the analysis of the *phenomena*, the givens (*data*) of conscious experience.)




Figure 1: Representation of a Hypothesized Protophenomenon of Visual Sensation. The real part of a Gabor elementary function corresponding to the receptive field profile of a cell in primary visual cortex.

## 2 Protophenomena

The foregoing considerations have led me to postulate *protophenomena* as the elementary constituents of phenomena (MacLennan 1995, 1996a). Each protophenomenon has the property of elementary (irreducible) subjectivity. A few examples of protophenomena may prove useful before we turn to a more careful description. Very simple protophenomena include the experience of a spot of color at a particular location in the visual field and the feeling of pressure at a particular location on the skin. To take a slightly more realistic example, there is considerable evidence (reviewed in Pribram 1991, MacLennan 1991) that cells in primary visual cortex respond optimally to “oriented grating patches” (more precisely, to Gabor elementary functions) at particular retinal locations (Fig. 1). Therefore, to a first approximation, the protophenomenon associated with such a cell can be taken as the experience of such a bit of spatial frequency at a particular location in the visual field.

However, protophenomena are more than simply “sense data,” and some of the more complex and subtle protophenomena include elementary components of recognitions, judgments, expectations, intentions, moods and so forth. Further, protophenomena are very “small,” in the sense that changes in the intensity of individual protophenomena will not typically affect the macroscopic phenomenal state; nevertheless the state of consciousness is no more than the sum total of the states of all the protophenomena. (The total conscious state of an individual comprises perhaps 10 to 100 billion protophenomena.)

Protophenomena are postulated to be associated with *activity sites* in the brain,

the “intensity” (degree of presence in consciousness) of a protophenomenon corresponding to some physical variable at that site. For example, just as the activity of a neuron in visual cortex may reflect the presence of the pattern in Fig. 1 on the retina, so the intensity of the corresponding protophenomenon reflects the degree of presence of that pattern in consciousness. There are a number of candidates for the activity sites, but their identity remains an open question. Some of the possibilities include synapses, neural somata, ion channels and dendritic microtubules, but their exact identity is not crucial for the theory of protophenomena. If they are sufficiently small then the activity sites would be subject to the laws of quantum mechanics. (Protophenomena and their activity sites need not be discrete, but that seems the most likely possibility at this time. Protophenomena would be continuous if, for example, the entire dendritic membrane acted as a continuum of activity sites.)

What is the ontological status of protophenomena; do they exist? It is best for now to treat protophenomena as “theoretical entities,” analogous to atoms when they were first hypothesized. Theoretical entities are validated by the role they hold in the theory and by their fruitfulness for scientific progress. Ultimately, we may find that protophenomena exist individually (in the same way that atoms were found to exist), for example as properties of individual activity sites. Alternately, we may find that protophenomena exist only in the context of large numbers of activity sites, and thus that they are emergent properties, analogous to emergent physical properties. For example, according to acoustical theory a sound wave has an objective pressure at each point in space, but it doesn’t make much sense to talk about the pressure of a single molecule or even of a small group of them. For now this is an open question. However, it is a question with empirical content, since it is possible, at least in principle, for phenomenologically trained observers to report changes in their conscious experience resulting from altering physical quantities at potential activity sites.

Causal dependencies among activity sites suggest how protophenomena are integrated into a phenomenal world. In particular, there is no reason to suppose that a neuron in auditory cortex is different in any essential way from one in visual cortex, therefore it must be the connections that cause one to correspond to the experience of sound and the other to the experience of color. I expect that protophenomena acquire their phenomenal quality in the same way. That is, an isolated protophenomenon has no qualitative character, just a subjective intensity. It is the dependencies between protophenomena that give them their character, and make one the experience of a sound and the other the experience of a color.<sup>1</sup> The remainder of this paper will, I hope, make it more plausible that the qualitative aspect of a protophenomenon is exhausted by its dependencies.

The role of connections in establishing a phenomenal world is illustrated also by the many topographic maps in the brain. In these maps the arrangement of neu-

---

<sup>1</sup>This structuralist account of qualia is not a necessary consequence of the theory of protophenomena, but it seems the simplest account in the absence of contradictory evidence.

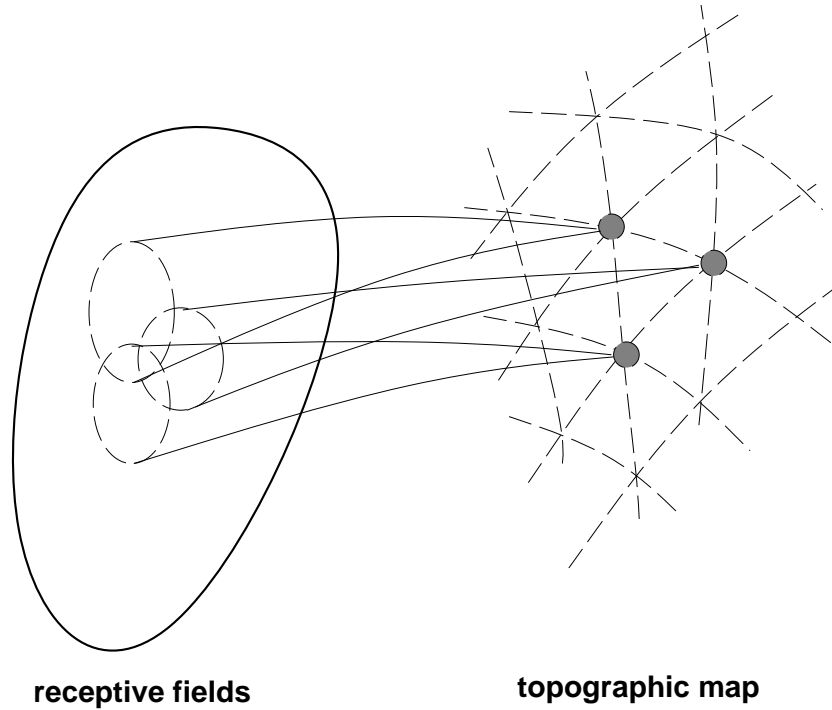


Figure 2: Continuity Through Topographic Maps.

rons corresponds to significant (micro-)features of the stimulus (e.g. spatial location, spatial or auditory frequency, orientation). Neurons represent specific microfeatures by virtue of their connections to other neurons in the same map or in other maps. Figure 2 illustrates how the overlapping receptive fields of the neurons in a map can define a continuous representational space.

The connections between neurons are paralleled by dependencies between protophenomena, which give them their characters in a web of protophenomena. Therefore we can begin to understand the dependencies among protophenomena by investigating the parallel dependencies between neurons.

The activity of a neuron typically depends directly on the activities of thousands of other neurons (several hundreds of thousands in many cases), and distributes its output in turn to thousands of destinations. Figure 3 is a simplified diagram of a neuron and its dendritic tree. Signals are transmitted from the other neurons to the dendritic tree and propagate in a complex way to the cell body. If, for simplicity, we take the membrane voltage in the cell body to be the activity site, then the physical parameters of the dendritic tree determine how this activity depends in the activities of other neurons. Therefore, we must consider the dendritic tree in more detail.

Figure 4 shows an equivalent circuit for the electrical dependencies in a small part of a neuron's dendritic tree, which describes how the activity in the neuron's soma




Figure 3: Simplified Diagram of Typical Neuron. The neuron's cell body is in the center, surrounded by its dendritic tree (which in reality has many more branches). In the upper right is another neuron, which makes contact in several places with the first neuron's dendritic tree.




Figure 4: Equivalent Circuit of Dendritic Tree. In the upper-left an axon forms a (bidirectional) synapse on a dendritic spine. In the lower-right the dendrites converge on the soma (cell body), where they are integrated into a potential, which may trigger impulses leaving on the right. Triangles represent voltage-controlled voltage sources, which are found at the synapses and axon hillock. For simplicity the diagram omits the conductance of the membrane.

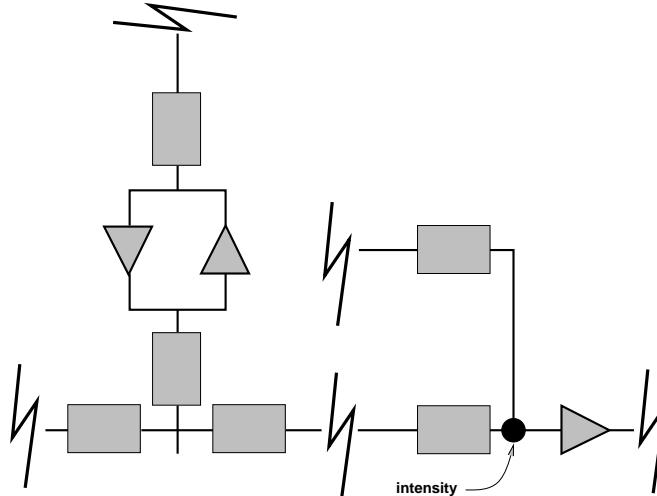


Figure 5: Structure of a Protophenomenon. For concreteness, protophenomenal intensity has been identified with the membrane potential of the soma. The grey rectangles and triangles represent the transfer functions or impulse responses that characterize how intensities are transferred from one protophenomenon to another. The rectangles (passive impedances) correspond to the electrical behavior of the dendrites, the triangles (active amplifiers) to that of the synapses or axon hillocks. See MacLennan (1996b, Appendix) for more details.

depends on the activities in other neurons. (This circuit corresponds to the passive electrical properties of the dendrite, to which we restrict our attention, for the sake of simplicity.) Figure 5 shows the corresponding dependencies in the protophenomenal realm. Electrical properties (such as resistance and capacitance) have been replaced by abstract mathematical relationships between protophenomenal activities.

Of course, the nervous system is not closed; the activities of sensory neurons depend on physical processes outside of the nervous system. In a parallel way the intensities of the corresponding protophenomena depend on extrinsic variables (i.e. variables whose values do not depend directly on the intensities of other protophenomena). Figure 6 illustrates a rod cell in the retina and some of its connections together with the corresponding protophenomenal dependencies.

The dynamics of protophenomenal activity can be described by differential equations, in particular, by the same differential equations that define the electrochemical behavior of the dendrites (MacLennan 1996b). In many cases the dependencies (the equations) are approximately linear, and protophenomenal activity can be described in terms of a *characteristic function* (often known as an “impulse response”). The characteristic function represents the spatiotemporal input signal to which the dendritic tree is optimally tuned. In this sense a protophenomenon can be considered to represent a spatiotemporal pattern in the intensities of the other protophenomena and extrinsic variables on which it depends.

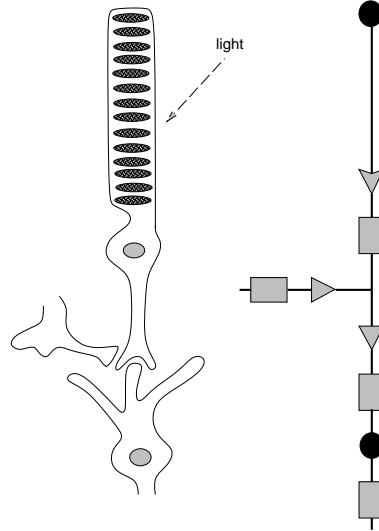


Figure 6: Sensory Input. The activity of a rod in the retina depends on light energy as well as on other neurons. The intensity of the corresponding protophenomenon depends on an extrinsic independent variable (black circle at top) as well as on other protophenomena.

Protophenomenal dependencies establish connections among protophenomena and thereby assemble them into a phenomenal world. One way they do this is by establishing continuity through expectations. Another way is by means of conjunctive dependencies and by more complex temporal dependencies. As a result a phenomenal world may be described by a set of possible trajectories in protophenomenal state space (MacLennan 1995, 1996a, 1996b).

In summary, the fact of phenomenal experience corresponds to a protophenomenon's intensity, since that intensity represents its degree of presence in conscious experience; the quality of conscious experience corresponds to the protophenomenon's dependencies, which relate it to other protophenomena.

### 3 Color and Spectral Inversions

As an example of the protophenomenal approach, we can consider the well-known problem of *inverted qualia*. We begin with a simple case, a pitch inversion, to show the technique, before turning to the more complex, but more interesting color spectrum inversion. We must distinguish physical sound frequencies, which may be high or low, from the experience of high and low pitches, which we may term *phenomenological high pitch* ( $\Phi$ -High) and *phenomenological low pitch* ( $\Phi$ -Low). The question then may be put: Is it possible that you experience low frequencies as high pitches ( $\Phi$ -High) and vice versa?



The possibility seems plausible, but careful observation of our experience of sound reveals that the inversion is impossible because high and low pitches have a different phenomenological topology. Suppose you listen to a tone of gradually decreasing frequency; as it does so, the tone will gradually become a rhythm, and the rhythm will eventually become an amplitude variation. This phenomenology is exactly mirrored in the neurology, for tones are represented spatially (via a tonotopic map) in auditory cortex, but as the frequency decreases the nerve impulses begin to synchronize with the sound waves; at still lower frequencies the individual waves are represented by bursts of impulses. Thus we may say that our experience of low pitches is entangled with our experience of rhythm and amplitude, and this distinguishes  $\Phi$ -Low from  $\Phi$ -High. Indeed, the phenomenology of pitches really can't be much different than it is. See MacLennan (1995, 1996a) for a fuller discussion of the impossibility of a pitch inversion.

Now we can consider the well-known problem of a spectral inversion, which dates back to Locke's *Essay Concerning Human Understanding* (1690) (e.g., Hardin 1988, Nida-Rümelin 1996, Palmer 1999, MacLennan 1999). In brief, the problem is as follows: Although we agree on the names for various wavelengths, is it possible that you experience red wavelengths the same way I experience blue wavelengths, and vice versa? Before we can solve this problem we need a more accurate phenomenology of color. The plausibility of a spectral inversion derives in part from an oversimplified phenomenology of color, since we have imagined that color can be reduced to a single dimension (wavelength) but a phenomenological analysis shows it to be much more complex. For example, the ancient Greek word *chlôros* is often translated "green," but then we are surprised to find in the literature that dew, tears and blood can sometimes be *chlôros*. Clearly, this category refers to more than a set of wavelengths. We can improve our understanding by observing that *chlôros* often applies to fresh, living or moist things. This is not so different from English, in which we can refer to a green twig or a green rider without implying that either one is green in color. The conclusion to be drawn is that it may be inaccurate to assume that an apparent color term denotes no more than a range of wavelengths, and that the rest is unessential connotation. Phenomenology demands that we take the entire semantic field evoked by the word as our data to be analyzed, and not be misled by preconceived ideas about supposedly simple sensory properties. See MacLennan (1998) and the references cited therein for a more detailed discussion.

Setting aside many of the higher-level complexities of color (e.g. its emotional and cultural connotations), yet avoiding the trap of a one-dimensional view, we can observe that it has long been known that we can identify four pure hues, which are termed the "unique hues," an observation that has led to the double-opponent theory of color vision. In this theory the three color receptors (short, medium and long wavelength, henceforth  $S$ ,  $M$  and  $L$ ) are combined in various ways to yield three orthogonal axes (Figs. 7, 8). The light-dark axis is formed by  $+S + M + L$  and its opposite; the yellow-blue axis is formed by  $-S + M + L$  and its opposite; the red-




Figure 7: Approximate Response Curves of Three Color Receptors. S, M and L represent the response of short, medium and long wavelength receptors. Wavelength in measured in nanometers.




Figure 8: Double Opponent Theory of Color Vision. Capital letters represent phenomenal brightness and color (L/D = light/dark, Y/B = yellow/blue, R/G = red/green).




Figure 9: Normal Response Curves of Double Opponent Cells. Capital letters represent phenomenal colors corresponding to peaks in the response curves. Letters with a subscript  $u$  represent the three spectral “unique” (unmixed) hues.




Figure 10: Normal Locus of Monochrome Light in Color Plane. Notice that some purples are *nonspectral*, that is, these experiences cannot result from seeing monochromatic light, although they may result from polychromatic light. The curve is generated from an approximate model of cell responses; the lower end of the curve corresponds to short wavelengths, the upper to long.

green axis is formed by  $+S - M + L$  and its opposite (Fig. 9).<sup>2</sup> The two zeroes on each of the two chromic axes (yellow-blue and red-green) define the four unique hues (Fig. 8). It is interesting to note that unique-red is a *nonspectral hue*, which means that it does not correspond to any wavelength of monochromatic light. This is because the red-green curve is bimodal and has two zeroes, whereas the yellow-blue curve is unimodal and has only one. Nevertheless, the sensation of unique-red can be created by mixing red light with a combination of yellow and blue that balances response on the yellow-blue axis. Figure 10 shows the locus of monochromatic light in the color plane, and we can see that certain reds and purples (defined by activity in the red-blue quadrant) are not generated by monochromatic light.

Before proceeding, I must answer a possible objection to this reconstruction of the topology of normal color vision. The reconstruction is based in part, but critically, on the overlap of the response curves of the three color receptors, but this overlap depends

---

<sup>2</sup>Here we use a common form of the theory; for recent accounts see De Valois & De Valois (1988, 1993) and Kaiser & Boynton (1996).

on optochemical relationships that are external to the nervous system proper, which I have presented as the basis for protophenomenal dependencies. Indeed, the inputs from the color receptors might be considered three independent extrinsic variables. However, they are not independent. We know from our understanding of the physical processes (represented by the response curves) that the receptors have correlated activities, and therefore that the corresponding protophenomena will have correlated intensities. Furthermore, we know that “Hebbian” (or correlational) learning in the brain will strengthen excitatory connections between interdependent activity sites with correlated activity. (And similarly such adaptation will lead to the strengthening of inhibitory connections in center-surround structures.) Therefore, through learning, which takes place in the phenomenal world in parallel with the brain, external correlations will become encoded as neural connections, which correspond to protophenomenal dependencies. Therefore, the structure of the phenomenal world becomes an image mirroring the structure of the external world. In this way conscious experience adapts to the structure of the environment. To return to the case of color experience: overlaps between response curves become mirrored in protophenomenal dependencies; among the consequences is our subjective experience of chronic similarity.

The problem of a spectral inversion can be recast in terms of inversions between the poles on one or more of these axes or in terms of exchanges between two or more of the axes. However, we will show by phenomenological analysis that such spectral inversions are impossible, that is, that abnormal neurological connections would lead to abnormalities in conscious experiences that could be detected by the subject. Here the arguments will be summarized briefly.

First, it is fairly obvious that dark and light have phenomenologically distinct characters, and hence are noninterchangeable: in the dark, forms and hues are indistinguishable, but not in the light. As Francis Bacon said, “All colours will agree in the dark.”

Second, phenomenological analyses of color from ancient times to our own have observed that yellow is intrinsically brighter than blue (and red and green), the so-called “yellow anomaly”; the neurophysiological reason is the large overlap between  $+S + M + L$  and  $-S + M + L$ . Conversely, blue is the intrinsically darkest hue. Hence, blue and yellow are phenomenologically similar to dark and light, and hence noninterchangeable. Therefore, in a case of abnormal vision, whatever receptor combination has the largest overlap with  $+S + M + L$  will be experienced as phenomenal-yellow ( $\Phi$ -Yellow), and if this does not correspond to spectral-yellow then the anomaly will be detectable. Further, we may define  $\Phi$ -Yellow as the experienced unique hue most similar to  $\Phi$ -Light, and its opposite,  $\Phi$ -Blue to be that experienced unique hue most similar to  $\Phi$ -Dark.

The case of a red-green inversion is more subtle, but phenomenological analysis again exposes a difference. For example, Goethe (1840) observed that green is a phenomenological mixture of yellow and blue, whereas red (*Purpur*) results from




Figure 11: Topology of Normal Color Sensation. Capital letters represent phenomenal colors.

an “augmentation” of yellow and blue. Further, the experience of “unique red” is nonspectral; that is, it cannot be created by monochromatic light, whereas experience of the other three unique hues (including green) can. The well-known studies of Berlin and Kay also support the phenomenal differences between red and green (Berlin & Kay 1969, Kay & McDaniel 1978, Saunders & van Brakel 1997).

Finally, the red-green axis cannot be exchanged with the yellow-blue, because the former is less similar to light-dark than the latter. This phenomenological fact, which has been recognized since ancient times, is consequence of  $+S - M + L$  (“red”) having a smaller overlap with  $+S + M + L$  (“light”) than does  $-S + M + L$  (“yellow”).

As a result of this neurophenomenological analysis, we can begin to understand the topology of color (Fig. 11). First we have the three axes, which define three polar oppositions: light-dark, yellow-blue, red-green. Superimposed on this structure are relations of similarity: yellow is most similar to light, and blue is most similar to dark.




Figure 12: Response of Opponent Cells with Abnormal Unimodal Channel. Response is measured in arbitrary units, wavelength in nanometers.

Green is most similar to yellow and blue and is intermediate in its similarity to light and dark. Red is similar to yellow, but not to blue. These conclusions are objective in that they result from observations made independently by many phenomenologists over the centuries.

Finally, we will consider several more examples of abnormal or nonhuman color perception. For example, if we have  $+S + M - L$  instead of  $-S + M + L$  in the unimodal (yellow-blue) channel, then spectral blues will be experienced as yellows, while spectrally green light will still be experienced as green (Fig. 12); a blue light (e.g. 460 nm) will generate the unique-yellow experience and a reddish-orange light (e.g. 605 nm) the unique-blue experience. To understand why, observe that if the  $+S + M - L$  (spectral blue) channel is most similar to the  $+S + M + L$  (light) channel, then  $+S + M - L$  will be experienced as  $\Phi$ -Yellow, and its opposite  $-S - M + L$  (spectral orange-red) as  $\Phi$ -Blue. In the bimodal channel ( $-S + M - L$ ), the mean between  $\Phi$ -Yellow and  $\Phi$ -Blue, experienced as  $\Phi$ -Green, will be spectrally green; its opposite ( $+S - M + L$ ), experienced as  $\Phi$ -Red, will be nonspectral red-purple. The result is the topology shown in Fig. 13; the differing correspondence between the phenomenological topology and spectral color would allow the abnormality to be detected.

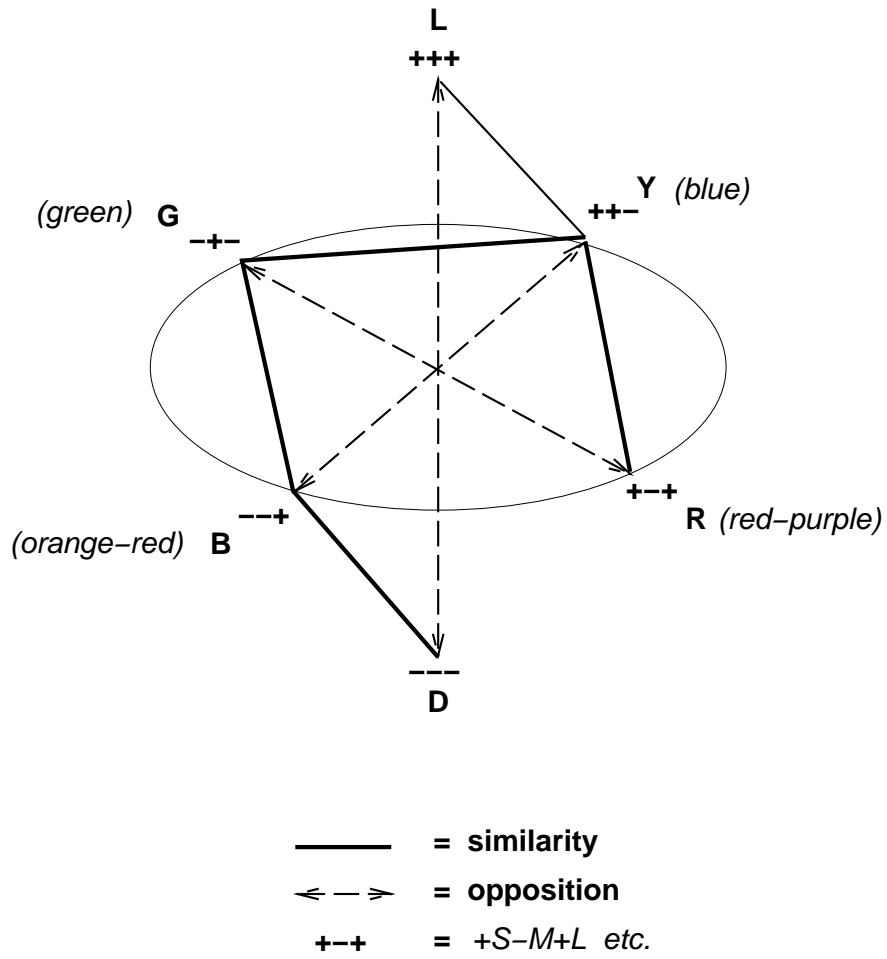


Figure 13: Topology with an Abnormal Unimodal Channel. Boldface capital letters represent phenomenal colors, lower-case words denote spectral colors.



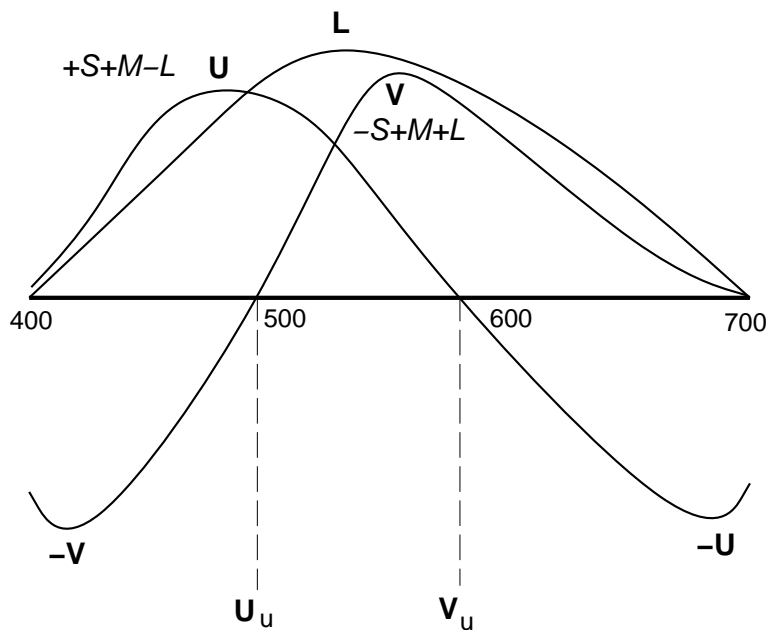


Figure 14: Response Curves with Bimodal Channel Replaced by Unimodal Channel. The corresponding phenomenal colors are labeled arbitrarily  $U$  and  $V$ .

On the other hand, if we have  $+S + M - L$  and  $-S + M + L$  (two unimodal channels) for the chromic channels, then color phenomenology will have several detectable anomalies: there will be two spectral unique hues as opposed to three (Fig. 14), and one whole phenomenal color quadrant ( $-U, -V$ ) will be nonspectral, that is, unexperiencable with monochromatic light (Fig. 15). This is because the bimodal channel has two zeros, whereas the unimodal channels have only one each. This leaves open the question of how a person with such an abnormality would experience color. The question is difficult because the two unimodal channels have approximately an equal overlap with the light ( $+S + M + L$ ) channel, and therefore an equal claim to be experienced as  $\Phi$ -Yellow. Indeed, for the sake of argument, we can assume that they have *exactly* the same overlap, giving a topology such as that in Fig. 16. The difference between the two unimodal channels might not be perceived as one of color (since, in effect, differences of experienced color correspond to differential activity compared with  $L = +S + M + L$ ).

Many other neural anomalies can be hypothesized. However, if a sensory system is too different from our own, we may be neurologically unable to imagine the experience, although we can describe its topology. Since imaginal areas in the brain have parallel structures to perceptual areas, we have limited ability to imagine qualia that are radically different from what we can perceive.




Figure 15: Locus of Monochrome Light with Bimodal Channel Replaced by Unimodal Channel. Lower part of curve corresponds to shorter wavelengths, upper to longer.




Figure 16: Topology of Color with Bimodal Channel Replaced by Unimodal Channel. Thick lines represent strong similarities, thin lines weak; dashed arrows represent oppositions.

## 4 Conclusions

The protophenomenal perspective has several benefits. First, it allows the fact of conscious experience to be integrated into scientific theory without denying or distorting the nature of that experience. Second, it permits a form of reduction of the more complex to the simpler while acknowledging the complexity of phenomena and avoiding naive introspectionism. Third, it permits a detailed account of the structure of conscious experience.

Of course, many open questions remain. For example: What are the activity sites and what sorts of physical systems can be activity sites? (This has implications for nonbiological consciousness.) What distinguishes conscious from nonconscious neural activity? Are protophenomena emergent? (This has implications for degrees of consciousness.) Are protophenomena qualitatively exhausted by their mutual dependencies (structuralism)? What can we say about the boundaries and unity of consciousness? Finally, much careful neurophenomenological work remains to be done before we will understand the detailed structure of consciousness.

## References

- Berlin, B. & Kay, P. (1969). *Basic Color Terms: Their Universality and Evolution*, University of California Press.
- Chalmers, D. J. (1995). Facing up to the problem of consciousness, *Journal of Consciousness Studies* **2**: 200–219.
- Chalmers, D. J. (1996). *The Conscious Mind*, Oxford University Press, Oxford.
- De Valois, R. L. & De Valois, K. K. (1988). *Spatial Vision*, Oxford University Press, New York.
- De Valois, R. L. & De Valois, K. K. (1993). A multi-stage color model, *Vision Research* **33**: 1053–1065.
- Goethe, J. W. v. (1840). *Goethe's Theory of Colours*, John Murray. Transl. C. L. Eastlake.
- Hardin, C. L. (1988). *Color for Philosophers: Unweaving the Rainbow*, Hackett Publishing Company, Indianapolis/Cambridge.
- Kaiser, P. K. & Boynton, R. M. (1996). *Human Color Vision*, second edn, Optical Society of America, Washington, DC.
- Kay, P. & McDaniel, C. K. (1978). The linguistic significance of the meanings of basic color terms, *Language* **54**: 610–646.

- MacLennan, B. J. (1991). Gabor representations of spatiotemporal visual images, *Technical Report CS-91-144*, Computer Science Department, University of Tennessee, Knoxville. Accessible via URL <http://www.cs.utk.edu/~mclennan>.
- MacLennan, B. J. (1995). The investigation of consciousness through phenomenology and neuroscience, in J. King & K. H. Pribram (eds), *Scale in Conscious Experience: Is the Brain Too Important to be Left to Specialists to Study?*, Lawrence Erlbaum, Hillsdale.
- MacLennan, B. J. (1996a). The elements of consciousness and their neurodynamical correlates, *Journal of Consciousness Studies* **3**: 409–24. Reprinted in *Explaining Consciousness: The Hard Problem*, ed. by J. Shear, MIT Press, Cambridge.
- MacLennan, B. J. (1996b). Protophenomena and their neurodynamical correlates, *Technical Report CS-96-331*, Computer Science Department, University of Tennessee, Knoxville. Accessible via URL <http://www.cs.utk.edu/~mclennan>.
- MacLennan, B. J. (1998). Finding order in our world: The primacy of the concrete in neural representations and the role of invariance in substance reidentification, *Behavioral and Brain Sciences* **21**: 78–79.
- MacLennan, B. J. (1999). Neurophenomenological constraints and pushing back the subjectivity barrier, *Behavioral and Brain Sciences* **22**: 961–963. Commentary on Palmer (1999).
- Nida-Rümelin, M. (1996). Pseudonormal vision: An actual case of qualia inversion?, *Philosophical Studies* **82**: 145–157.
- Palmer, S. E. (1999). Color, consciousness, and the isomorphism constraint, *Behavioral and Brain Sciences* **22**: 923–989. Includes peer commentary.
- Pribram, K. H. (1991). *Brain and Perception: Holonomy and Structure in Figural Processing*, Lawrence Erlbaum, Hillsdale.
- Saunders, B. A. C. & van Brakel, J. (1997). Are there non-trivial constraints on colour categorization?, *Behavioral and Brain Sciences* **20**: 167–228.