THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Prospects for Naturalizing Color

Citation for published version:

Isaac, A 2009, 'Prospects for Naturalizing Color', *Philosophy of Science*, vol. 76, no. 5, pp. 902-914.
<<http://www.jstor.org/stable/10.1086/605801>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Published In:

Philosophy of Science

Publisher Rights Statement:

© Isaac, A. (2009). Prospects for Naturalizing Color. *Philosophy of Science*, 76(5), 902-914.

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Prospects for Naturalizing Color, or “What’s blue and yellow and green all over?”

Alistair Isaac

June 2, 2008

Abstract

In [Chu07], Paul Churchland gives an argument for the “objective reality” of color; the strategy he deploys to make this argument is an instance of a more general research program, which he calls “domain-portrayal semantics.” In section 1, I point out some features of color vision which complicate Churchland’s conclusion, in particular, the context sensitive and inferential nature of color perception. In section 2, I examine and defend the general research program, concluding that it lies at the intersection of strategies to naturalize representational content. Such a minimally naturalistic, or operationalist, program involves two components: first, a mapping between the target domain of represented structures (colors, shapes, middle-sized dry goods, whatever) and the target range of representational cognitive structures; second, a detailed account of the causal processes which induce this mapping. I conclude with some conjectures concerning the future of such an operationalist program, in particular, that other perceptual domains may exhibit the same context sensitivity and inferential character as color.

1 Problems for a Metaphysically Robust Account of Color

I present Churchland’s argument in brief, emphasizing the problem he sees as the crux of the matter and his solution to it. Next, I briefly discuss the context sensitivity and inherently inferential nature of color perception. These consid-

erations suggest that Churchland’s program will not succeed as an argument for realism about color.

1.1 Churchland’s argument for the “objective reality” of color

In order to argue for the “objective reality” of color, Churchland provides himself with a specific explanatory task:

“How,” it may be asked, “does the peculiar and well-defined three-dimensional structure of the human phenomenological color space . . . map onto the objective space of possible electromagnetic reflectance profiles displayed by material objects? What is the internal structure of that objective target feature-domain in virtue of which the internal structure of our phenomenological color spindle constitutes an accurate map of that target domain?” ([Chu07], 123-4)

In providing such a mapping, Churchland focuses on one particular problem: how can the diverse group of distinct reflectance profiles¹ which appear subjectively identical be characterized as a single kind in a non-circular fashion? The naive strategy of characterizing “red” as just the equivalence class of reflectance profiles which induce a sensation of red (for a “normal” subject in standard viewing conditions, *etc.*) does not give Churchland the *objectivity* he desires. In order to defend his conclusion, Churchland must define the equivalence classes of reflectance profiles correlated with phenomenological color space in a manner which does not include reference to properties of the observer. In making this argument, Churchland helps himself to the idea that the theory of metamers, or classes of physically distinct spectral power distributions of light which match onto the same point in phenomenological color space, can be recast in terms of reflectance profiles. *Pace* Churchland, I will reserve the term “metamer”

¹A surface reflectance profile is a function which assigns a value (in the range 0-100) to each possible wavelength of electromagnetic radiation. This number represents the percent efficiency with which the surface reflects electromagnetic radiation at that wavelength.

for lights with different spectral power distributions which appear subjectively identical; I will use the term r-metamer to stand for different reflectance profiles which appear subjectively identical under some specified lighting conditions.² Recognizing the difficulties in specifying such lighting conditions, let us set that issue aside and examine Churchland’s solution to the problem of identifying the appropriate r-metamers in a non-circular fashion.

Although his argument is far from rigorous, Churchland’s intuitive account is quite ingenious. First, consider only that region of an object’s reflectance profile which encompasses the range of wavelengths visible to the human eye (roughly, from 400 nm to 700 nm). Next, consider this section of the reflectance profile rolled into a cylinder; this turns the “rectangular” space of possible reflectance profiles into a “cylindrical” space, with reflective efficiency values for different wavelengths falling on the surface of the cylinder. Churchland claims that any reflectance profile will have a unique ellipse, induced by intersecting the cylinder with a plane, which most closely approximates it. He proposes to identify equivalence classes of r-metamers with classes of reflectance profiles approximated by the same ellipse. In Churchland’s words:

This trick turns the original reflectance profile itself, whatever its idiosyncratic ups and downs, into a wraparound configuration that admits of an optimal approximation by a suitable planar cut through the now-cylindrical space. . . The peculiar ellipse produced by a specific cut will be said to be an optimal—or, as I shall say henceforth, canonical—approximation of the original or target reflectance profile when it meets the following two defining conditions:

1. The altitude of the ellipse must be such that the total area A above the canonical ellipse, but below the several upper reaches of the target reflectance profile, is equal to the total area B beneath the

²Since the experiments which produced the modern theory of subjective color space, and likewise of metamers, were performed with lights, I am worried that some difficulties will arise if we try to apply the theory of metamers *mutatis mutandis* to surface reflectance profiles. At the very least, we need an account of standard lighting conditions to make this work because the spectral power distribution of light entering the eye after bouncing off an object is a function of that object’s reflectance profile *and* the spectral power distribution of the source light.

canonical ellipse, but above the several lower reaches of the target reflectance profile. (This condition guarantees that the total area under the target reflectance profile equals the total area under the approximating ellipse.)

2. The angle by which the ellipse is tilted away from the horizontal plane, and the rotational or compass heading positions of its upper extreme, must be such as to minimize the magnitude of the two areas A and B . This condition guarantees that the approximating ellipse follows the gross shape of the target reflectance profile, at least to the degree possible. ([Chu07], 128)

The space of possible canonical approximation ellipses (“CA-ellipses”) can be characterized by three dimensions of variance: 1) altitude of the CA-ellipse’s center; 2) degree of tilt from the horizontal; and 3) rotational position within the cylinder. Churchland proposes to identify these three dimensions with the three dimensions of the spindle shaped space of phenomenological color perception. Roughly, CA-ellipse altitude corresponds to brightness, tilt corresponds to saturation, and rotational position corresponds to hue. Intuitively, the story looks plausible, and the space of canonical reflectance profiles indeed looks quite similar to that of phenomenological color perception. Churchland calls this similarity “the answer to a color realist’s prayer” ([Chu07], 133) and focuses in his discussion primarily on objections concerning the slight metrical deformation of the phenomenological color space with respect to the space of CA-ellipses under the proposed mapping.

1.2 Critique of the Churchland account

Compelling though it may be, and setting aside the problem of standard lighting conditions, Churchland’s argument fails to exhibit sensitivity to the role of context in color perception. The most obvious strategy for adapting Churchland’s argument involves reintroducing the subject into the characterization of viewing conditions. The manner in which this would have to be done, however, would

leave some objects with non-unique color assignments, drawing into question the objectivity of such.

1.2.1 The context-sensitivity of color perception.

When performing experiments on color perception, the usual setup³ involves a test light and three primary lights. The subject's visual field is divided, with the test light on one side and the primaries mixing on the other; the subject manipulates the intensities of the primary lights until their mixture appears subjectively identical to the test light. Symmetric matching experiments such as this can be used to identify the class of metamers for a given subject. Demonstrably, this class of metamers is invariant with respect to context, so long as the context remains symmetric. In other words, suppose A and B are two metameric spectral power distributions, β_1 and β_2 are two contexts (*i.e.* colored environments within which the test lights occur), and $\beta_i(X)$ is the value in phenomenological color space induced by spectral power distribution X in context β_i . Then, since A and B are metamers, it will be the case that $\beta_1(A) = \beta_1(B)$ and $\beta_2(A) = \beta_2(B)$, but it will not in general be the case that $\beta_1(A) = \beta_2(B)$, or vice versa (or even that $\beta_1(A) = \beta_2(A)$). The crucial point is that context affects the phenomenological color value assigned to a particular spectral power distribution.

For example, Jameson and Hurvich ([Jam59]) compare two equations for describing the coding of spectral power distributions in three variables. One such equation includes an additional factor (I , I' , *etc.*) for capturing the influence of context on neural processing while the other does not. They remark:

It seems clear that so long as we are concerned only with the identity or nonidentity of colors produced by stimuli of different wavelength distributions for any of a series of different luminance levels, or for

³For a more detailed discussion see [Wan95], 80ff.

any of a series of *symmetrical* adaptation and surround conditions, there is no way of discriminating between the alternative equations, and the introduction in these relations of the expression I or I' to express the induced neural activity remains arbitrary.

If, on the other hand, we are also interested in the subjective appearance of the color C , then the alternative relations are by no means equivalent, and the disparity between them will depend on the contribution of the induced response (I or I') relative to the total response. Moreover, when we come to compare color perceptions, color matches, or psychophysical functions for different, or nonsymmetrical conditions of adaptation and/or contrast, we find that the expression I (or I') is no longer a matter of indifference to the result, and its importance can be subjected to an empirical check. ([Jam59], 892)

A similar point is made by Suppes, *et al.* ([Sup89]):

We have emphasized that in a metamer color match $a \sim b$, the lights a , b are viewed under identical conditions. If the viewing conditions change, the appearance of the metamer color match may change drastically though it remains a match. It follows that asymmetry of the viewing conditions often produces a color mismatch between metamers or, for that matter, isomers (identical) distributions. In fact, the perceptual effect of viewing conditions is demonstrated... by showing that physically identical lights no longer match in color. ([Sup89], 252)

So, the subjective appearance of a colored light depends on the context, or the spectral power distribution of nearby colored light. Now, since Churchland has taken the theory of metamers for spectral power distributions of light and applied it, *mutatis mutandis*, to surface reflectance properties of objects, it seems fair to bring along as well this context sensitivity for subjective color assessments. In other words, even if we presume some standard lighting conditions, stipulated in terms of “broad daylight,” say, the same reflectance profile will produce different subjective color assessments depending upon the reflectance profile of nearby surfaces. In fact, this is empirically the case as we can see from the Munker illusion, the watercolor illusion, and the neon color spreading

illusion, to name just a few.⁴

Of course, color illusions involving reflectance properties of surfaces usually involve distinct surface reflectance profiles occurring in very close proximity. The realist may therefore attempt to argue that such contextual effects are purely “local,” *i.e.* occur only when the two surfaces subtend a sufficiently small angle in the visual field. One could therefore avoid the phenomenon by stating that the (homogenous) surface in question must subtend some substantially large angle in the visual field. No matter how recalcitrant, any contextually induced effects on subjective color assessment should disappear if the contextual factors are removed to a sufficient distance from the focal viewing area. Nevertheless, if one adds such a specification to the requirements of standard viewing conditions, one reintroduces the subject into the equation. In other words, one can only specify the visual angle the stimulus will occupy in terms of the subject’s position with respect to the stimulus, thus her distance from the stimulus will become a factor in the characterization of standard viewing conditions. At first blush, introducing such a fix to the context sensitivity problem removes the objectivity of the characterization of reflectance profiles.

There is a further wrinkle, however: the assumption that we can fix the standard viewing conditions by specifying the subject’s position with respect to the stimulus depends upon the homogeneity of surface reflectance properties. Yet, such surface properties are not in general homogenous. A consequence is that some objects will fit the standard viewing conditions multiple ways, and will thus be multiple colors simultaneously. Consider, for example, a cloth of finely woven blue and yellow threads. If viewed from a very close distance, such that individual threads subtend a large enough angle of visual space to conform to the standard viewing conditions, the object will be striped blue and yellow. If

⁴Each of these illusions involves creating the appearance that two patches of surface with the exact same reflectance profile are “different colors” by giving the surfaces surrounding them dramatically different reflectance profiles.

viewed from a further distance, however, such that the cloth as a whole subtends the appropriate angle, then it will be a homogenous shade of green. Depending, in other words, on the size of the patch of heterogenous surface one chooses to average over, the spectral reflectance profile, and thus the “objective” color, may be significantly different.

1.2.2 The inferential character of color perception

Ironically, it is the context-sensitivity of color vision which produces our subjective sense of color constancy. It is only because “[c]olor appearance depends more on the local contrast of the cone absorptions than on the absolute level of cone absorptions” ([Wan95], 289) that an object (say, an apple) will appear the same color in vastly different viewing conditions (say, high noon and twilight). Since Helmholtz, many vision scientists have found it useful to interpret the computations of the visual system which result in judgments such as color constancy as *inferential* in character. Although the *process* by which conclusions about color are derived is unconscious, and thus may be qualitatively different from that of conscious inferences, the epistemological status of these “unconscious conclusions” is the same.

[W]hile it is true that there has been, and probably always will be, a measure of doubt as to the similarity of the psychic activity in the two cases, there can be no doubt as to the similarity between the results of such unconscious conclusions and those of conscious conclusions. These unconscious conclusions derived from sensation are equivalent in their consequences to the so-called *conclusions from analogy*. ([Hel66], 174)

If we accept (as many perceptual scientists do) Helmholtz’s analysis, then phenomenological color space is a kind of theory about the world. The problem of determining the reality of colors, then, may be analogous to the problem of determining the reality of scientific theories. In particular, even the most robust

scientific theories (say, Newton’s gravitational theory) bear only a statistical relationship to the data which support them, leaving perpetually open the possibility that a better theory may come along and prove them “false.” Likewise, color judgments bear only a statistical relationship to surface reflectance properties. The realist strategy is to eliminate the probabilistic character of this relationship by specifying standard viewing conditions. As discussed above, however, there may exist objects for which even the best notion of standard viewing conditions makes no sense.

Perhaps it would be useful here to consider an analogy with a perceptual domain which is commonly taken to be objective, *size*. Size judgments exhibit the same context sensitivity as color judgments for precisely the same purpose, namely to induce a constancy over perceptions of size. The effects of context on size judgment can easily be seen in the Ames room illusion. Consider another optical illusion, however, the changing size of the moon. Subjectively, the moon appears significantly larger when it is close to the horizon; yet the moon always subtends the same angle of visual space, as its distance from any observer remains (approximately) constant. *Something* about the contextual cues when the moon is close to the horizon induces the visual system to infer that the moon is larger than it infers it to be when it is higher in the sky (although several explanations for the phenomenon have been proposed, there is no consensus on the correct explanation). Now, since it is quite easy to operationalize our notion of size in a manner that is invariant with respect to viewing conditions (say, by measuring with a ruler), we have an objective notion of the size of the moon. But suppose we had no recourse but to specify the size of the moon with respect to standard viewing conditions; what would the standard viewing conditions be such that we could make an appropriate judgment about the size of the moon? The problem here is that any of the candidates for standard viewing conditions

which might work for everyday objects (middle sized dry goods, *etc.*) cannot be adapted to viewing conditions for planetary bodies; planets are simply too big for the perceptual system's techniques for inferring size to apply.

Returning to color, I can now state the problem for Churchland more succinctly: so long as the only test available for the "objective" partition of the space of surface reflectance properties is that it matches phenomenological color space, the very properties of color perception which produce color constancy for everyday objects like oranges and apples will defeat the uniqueness of color values for atypical objects. Yet, if color is objective, then these atypical objects should have unique, determinate colors associated with them.

2 Minimal Naturalistic Semantics for Representational Content

As Wagner and Warner have noted, "eliminativists and proponents of folk theory tend to communicate through short dismissals of each other's positions" (13). Likewise, it is tempting for those of us who reject naturalism in its robust metaphysical form to dismiss Churchland's program on the basis simply of his reductivist and (naive?) realist conclusions. However, the research program proposed by Churchland should be of interest to any philosopher who countenances a notion of representational content and respects the methodology and results of the sciences. If we reinterpret Churchland's program in this light, rejecting strong reductivist conclusions in favor of a more tempered, or "minimal," reading, we can appreciate its explanatory value for any developed philosophical account of representational content.

2.1 A program for naturalizing representational content

Model theoretic semantics assigns primitive terms and predicates “meaning” by providing them an interpretation in a set-theoretical structure. Can we make sense of an analogous project for non-linguistic domains such as, say, the space of percepts for a given sense modality? Since perceptions can vary continuously with respect to certain features (*e.g.* “sourness,” “heat,” “loudness,” “redness,” *etc.*), an account of the meaning of these perceptual “qualia” might, analogously, be given via an interpretation of the axes which define the space of variation with respect to the axes of variation of some model (potentially, though not necessarily, characterized in set theoretic terms). If, furthermore, the “model” of a given perceptual space is a structure external to the brain and well defined by natural science, we might claim to have provided a “naturalistic semantics” for the perceptual space. Churchland motivates a project along these lines with a Just-So story about the evolutionary development of contentful cognitive structures:

A promising general approach to understanding how the brain—or any of its various subsystems—represents the external world posits the brain’s development, through learning, of a variety of (often high-dimensional) maps of the objective similarity-structure of this, that, or the other objective feature-domain. Through extended experience with the relevant objective feature-domain, the relevant part of the brain can construct an internal map of that domain—of the range of possible faces, the range of possible voices, the range of possible reaching options, the range of possible colors, and so forth. ([Chu07], 121)

This account suggests a bipartite program for providing a naturalistic semantics for representational mental states. One goal is to provide a mapping between a cognitive representational space (such as, say, the phenomenological color space) and some well defined domain of variance in the environment (whether it be surface reflectance profiles, spectral power distributions, or what have you).

The second goal is to provide a detailed causal story for how such a mapping is induced; such a causal story will, necessarily, incorporate the causal structure of the perceptual system’s “inferential” computations.

Churchland only explicitly acknowledges the first of these goals, claiming that by providing a mapping between cognitive domain and external feature domain, one has provided a semantics for the cognitive domain.

Such homomorphisms or second-order resemblances, on this view, are the essence of the brain’s representational achievements. One might call this account “domain-portrayal semantics,” to contrast it with such familiar doctrines as indicator semantics or causal covariation semantics. ([Chu07], 123)⁵

However, he seems to explicitly acknowledge the importance of the second goal when referring to “the nature and ground of our internal map” ([Chu07], 125). If I interpret him correctly here, by “nature” he means the well developed (axiomatizable, *etc.*) theory of phenomenological color space, and by “ground” he means the physiological story about how this space is neurally coded (partially understood, though many questions remain concerning the exact connection between the physiological and the phenomenological here).

Churchland himself has attempted to engage in the second goal in [Chu05], where he argues for a correspondence between idealized models of the neural wiring for opponent color vision (as offered in, for example, [Hur81], Ch. 11) and phenomenological color space. Yet neural network models such as Churchland considers (those in which the three types of cone cells are wired into three types of opponent color cells: red-green, blue-yellow, and white-black) are hand-tuned, as it were, to correspond to the phenomenological color spindle. A perhaps more critical step for the program outlined here is a detailed account of the relationship between such idealized models of neural wiring and the actual wiring of the visual system as discovered via experimentation. For example, Derrington, *et*

⁵*c.f.* [Gar00], which presents a similar idea in terms of “conceptual spaces”

al. ([Der84]) found neurons in the macaque LGN which respond as if receiving a red-green opponent color signal yet which “do not meet the specification of the canonical R-G opponent cell” (264). As Wandell discusses:

The spectral responses of these neural populations suggest that there is only a loose connection between the signals coded by the neurons and the perceptual coding into opponent hues; it is unlikely that excitation and inhibition causes our perception of red-green and blue-yellow as opponent colors. One difficulty is the imperfect correspondence between the neural responses and the hue-cancellation measurements. The second difficulty is that there is no substantial population of neurons representing a white-black signal. This is a very important perceptual dimension which must be carried in the signals of the LGN. Yet, no clearly identified group of neurons can be assigned this role. ([Wan95], 324)

2.2 The role of “reductions”

Churchland concludes [Chu05] with some remarks on the import of the “reduction” he has provided:

[T]he principal intellectual motive for embracing the systematic identities between color qualia and coding vectors proposed is simply the extent and quality of the predictive and explanatory unity that the relevant reduction provides. ([Chu05], 554)

I hope the above discussion has demonstrated, however, that we are still far from providing a detailed enough account of the relationship between phenomenological color space and reflectance properties (or even between idealized neural network models and empirical measurements of neural activity) for there to be any serious candidate for a *reduction* of color qualia on the table. In fact, if such a “reduction” is ultimately provided, it may be so convoluted and qualified that our intuitions concerning its ontological import could well be defeated.

Controversies arise when claims about reduction are ideological rather than scientific in character. It is not usually appreciated how involved and technical the actual reduction of one part of science—even a near neighbor—is to another. ([Sup02], 52)

These general considerations aside, if we do manage to succeed in providing a convincing “naturalistic semantics” for a particular perceptual domain along the lines outlined above, will we have then provided a “reduction” of that domain?

Here is precisely where the non-naturalist’s eyes begin to glaze over, her resistance to “reduction” in any form will make her skeptical of the entire project. Yet, the strategy for providing “naturalistic semantics” outlined above is minimal in just this sense: *it is neutral with respect to the ontological import of such an account*. Consider, for example, Mary’s question, namely: “will the complete scientific account of color vision tell one (ideally one who has never seen it) what it’s like to experience the color red?” To oversimplify, here, briefly, are three potential answers:

1. [*metaphysical naturalist*] *Yes*, because we have reduced color “qualia” to X (= reflectance properties, whatever)
2. [*methodological naturalist*] *No*, because the question is ill-posed; science is not in the business of explaining what subjective experiences are “like.” Insofar as the question can be recast in operational terms, then *yes*.
3. [*non-naturalist*] *No*, there will always be an additional phenomenon (the “qualia”) in need of (philosophical) investigation, which cannot be explained via the scientific account.

Yet, as stated above, the project for providing a naturalistic semantics for color is neutral between these answers. Even if it succeeds as an account of “meaning” for color percepts, the implications of this semantics for ontological questions, or questions of explanatory power, will remain open. In fact, even the non-naturalist should benefit from such a proposed semantics: surely one can better critique the view that color qualia can be reduced to physical phenomena if one is informed by the exact proposal? As we’ve seen in the quote from Wandell which concludes the previous section, ammunition for the anti-reductive view may come from science itself.

A minimal naturalistic semantics will provide the best available candidate for a reduction of any perceptual domain to physical phenomena, yet an argument that such a reduction has been provided will be needed *in addition to* the minimal account. Furthermore, even philosophers resolutely opposed to reductive programs may benefit from such an account. Finally, the resolute naturalist himself would do well to clearly distinguish between those aspects of his argument which carry ontological weight and those of a more minimal character, the better to communicate with those whose too hasty dismissals depend upon the former and blind them to the general value of the latter.

3 Future Prospects for Naturalistic Semantics

Color is uniquely suited to the research program discussed above. This is because psychology has provided us with a robust and detailed theory of phenomenological color space, physics has provided us with a detailed theory of the properties of light, and neurophysiology has uncovered many of the neural mechanisms for coding color information. Ironically, on all three accounts, scientific theories pertinent to color are far better developed than those for many “primary” properties such as shape. Scientific theories of the phenomenological space of space perception, for example, are woefully underdeveloped, and theories of the neural substrate of shape perception are similarly poor. Contrary to *a priori* analysis, then, it looks as if a naturalistic account of phenomenological color space may actually be much *easier* to provide than one of phenomenological shape space (and, *a fortiori*, the corresponding spaces of middle-sized dry goods or linguistic abstracta). Given that these domains are much more complex than that of color, we may expect a fully detailed naturalistic semantics to be that much more difficult to develop, potentially encountering problems far more intricate than those posed by the context sensitivity and inferential character of

color perception. Perhaps, then, rather than struggling to defend realism about colors, the naturalistic semanticist may wish to question realist assumptions about the objects of everyday experience. For it is by applying lessons learned in simple domains to progressively more complex ones that science proceeds, and so also with the cautious naturalist.

References

- [Chu07] Churchland, Paul (2007) “On the Reality (and Diversity) of Objective Colors: How Color-Qualia Space is a Map of Reflectance-Profile Space,” in *Philosophy of Science*, 74; 119 - 149.
- [Chu05] Churchland, Paul (2005) “Chimerical Colors: Some Phenomenological Predictions from Cognitive Neuroscience,” in *Philosophical Psychology*, 18:5; 527 - 560.
- [Der84] Derrington, A. M., J. Krauskopf, and P. Lennie (1984) “Chromatic Mechanisms in Lateral Geniculate Nucleus of Macaque,” in *Journal of Physiology*, 357; 241 - 265.
- [Bay87] Baylor, D. A., B. J. Nunn, and J. L. Schnapf (1987) “Spectral Sensitivity of Cones of the Monkey *Macaca Fascicularis*,” in *Journal of Physiology*, 390; 145 - 160.
- [Gar00] Gardenfors, Peter (2000) *Conceptual Spaces: The Geometry of Thought*, MIT Press; Cambridge, MA.
- [Hel66] Helmholtz, Hermann (1866 / 1968) “Concerning the Perceptions in General,” in *Helmholtz on Perception: Its Physiology and Development*, ed. Richard M. Warren and Roslyn P. Warren, John Wiley & Sons, Inc.; New York, NY.

- [Hur81] Hurvich, Leo M. (1981) *Color Vision*, Sinauer Associates Inc.; Sunderland, MA.
- [Jam59] Jameson, D. and L. M. Hurvich (1959) “Perceived Color and Its Dependence on Focal, Surrounding, and Preceding Stimulus Variables,” in *Journal of the Optical Society of America*, 49:9; 890 - 898.
- [Sup89] Suppes, Patrick, David H. Krantz, R. Duncan Luce, and Amos Tversky (1989, 2007) *Foundations of Measurement, vol. II: Geometrical, Threshold, and Probabilistic Representations*, Dover reprint; Mineola, NY.
- [Sup02] Suppes, Patrick (2002) *Representation and Invariance of Scientific Structures*, CSLI Publications; Stanford, CA.
- [Wag93] Wagner, Steven J. and Richard Warner (1993) “Introduction” to *Naturalism: A Critical Appraisal*, University of Notre Dame Press; Notre Dame, Indiana.
- [Wan95] Wandell, Brian A. (1995) *Foundations of Vision*, Sinauer Associates Inc.; Sunderland, MA.