I. Background: The Structure(s) of Color Space

Philosophical and psychological theorizing about color has traditionally made much of the idea that colors (or perhaps color appearances) are essentially, i.e. by their very own nature, ordered in ways that are phenomenally evident to a perceiver in virtue merely of her experience of them.
In 1921, for example, the Cambridge logician, W. E. Johnson, intriguingly suggested that our grasp of "adjectives" such as number, color, and shape is governed by a certain characteristic structure. Johnson started by noting that there is no one character that is common to the members of such classes: there is, for instance, no "adjectival character" shared by red and blue, or by red things and blue things. The colors fall into a single class, then, not because there share something, but because of a "special kind of difference which distinguishes one color from another; whereas no such difference exists between a color and a shape." This "special kind of difference" is an ordering, Johnson thought:
within such classes, we are able to generate certain sub-classes simply by our grasp of the relevant ordering relations.
Number is his favored example: (1) It is ordered by the relation more than (and its converse, less than).
(2) This ordering relation is constitutive both of number itself, and of our grasp of it: we could not properly understand number without grasping more than as it applies to the numbers.
Finally, (3) Our grasp of number-ordering enables us to construct subclasses such as {x: x is a number and x is less than 4}: we know that the extension of this class is {1,2,3} directly as a consequence of possessing an adequate grasp of the number concept.
Johnson implies that we do not learn about the numbers less than 4 separately from learning about number.
By way of contrast, consider what he calls "substantives": the grasp of an inclusive concept such as mammal, gives us no specific knowledge about subclasses such as dog -- or even knowledge that such a subclass exists.
Johnson's claim, then, is that ordering relations are a constitutive part of certain concepts, and that they afford us the means by which to construct certain subclasses of these concepts.
He called these concepts "determinables," and the sub-classes so generated "determinates."

Johnson thought that color is ordered in this manner; like number, it is a "determinable".
Here, he was drawing on an ancient tradition.
The idea that colors are structured by ordering relations has always struck philosophers as quite intuitive.
However, there was no consensus, either in philosophy or in psychology, of what the ordering relations might be.

(See Kuehni, section I on various alternatives.)
Darker/lighter is a partial ordering relation that caught Aristotle's attention, and it has stayed in the running through the centuries; saturated/desaturated entered the discourse at some point -- but even taken together these are not sufficient to generate all the determinate colors.
It is likely that philosophers in the middle of the twentieth century thought that the colours were additionally ordered roughly as in the rainbow --- or as in the Newtonian color wheel, perhaps --- with brightness and saturation as additional dimensions, but the literature contains little that is explicit about these matters.
(However, Wittgenstein had some interesting musings about the phenomenology of color in his posthumous Remarks on Colour.)
More recently, psychologists and philosophers have been particularly interested in the ordering relations, 'is perceptually similar to' and 'is perceptually distinguishable from'.
These can be defined (and psychophysically operationalized) over the domain of colors.
Moreover, one can use these relations together with mathematical tools such as multidimensional scaling to define a space of colors with a "distance" between any two colors.
That is, plotting out the colors with the appropriate distances between them yields a multidimensional arrangement.
The claim is that our perceptual grasp of colors is essentially ordered by this metrically organized space.
This gives us a structure that conforms to Johnson's analysis.
Suppose that you are shown a sample of red -- crimson, say.
From this you can mentally "generate" brighter shades of red.
You do this simply by your grasp of color.

By the nineteen fifties, however, Leo Hurvich and Dorothea Jameson were urging the adoption of an alternative to the Newtonian color wheel. Using ideas foreshadowed in the work of Thomas Young and Ewald Hering, Hurvich and Jameson urged that hue consisted of two dimensions --- one running from blue to yellow, and the other from red to green.
Adding darkness/lightness to hue (darkness/brightness in the case of luminous sources), we then get three dimensions of color.
This proposal is stronger, one should note, than the "similarity space of color" idea just mooted, because in addition to the idea that colors are ordered in three dimensions, the Hurvich-Jameson model identifies three privileged and simple dimensions.
In this model the primary colors are opposite poles of a continuous measure --- for instance blue is the maximum value in one direction of its axis, and yellow is the opposed extreme value.
Every hue (including all the saturated hues) is a combination of some value along each of these privileged dimensions: for instance, orange

is a combination of more or less equal amounts of red and yellow; various shades of orange have more or less of red and orange, and they are also arrayed along the bright-dark axis (a peachy orange being lighter than the color of Seville marmalade).
Moreover, since these axes are simple, red is not a combination of purple and orange (even though it is in-between them on the hue-circle).
Note also that in Hurvich-Jameson hue-space, no color can be both reddish and greenish, or both bluish and yellowish, since these are opposite ends of a single dimension; this follows from the fact that each color occupies only one position in color space.
(See Byrne and Hilbert for more discussion.)

The Hurvich-Jameson model is closer to Johnson's notion of a determinable, because it posits ordering relations that are separately known, not just an ordering in a three-dimensional space, but a space constructed from these separate dimensions.
The Hurvich-Jameson hue circle is quite different from the ordering of colors in the rainbow (or the Newton color wheel), since red would be one right angle away from blue and yellow and two right angles away from green --- whereas in the rainbow, green and yellow are relatively close together.
(Incidentally, the talk of mixture should not be confused with that about the mixing of pigments, which has rules of its own: we are talking about color as it is experienced.
The claim is that color is experienced as a combining of intensive values from scales just mentioned.)

The work of Hurvich and Jameson was relatively little known in philosophical circles until C. L. Hardin's very influential book Color for Philosophers: Unweaving the Rainbow was published in 1988.
But it rapidly became standard to talk of color ordering in this way.
Johnson's ordering relations came via these authors' model of color space to possess a standard content, to the point where the "traditional color space", as we shall call it --- or TCS for short --- is now standardly taken to include the Hurvich-Jameson hue circle.
TCS is simultaneously thought of as generating the colors, and also as a graphical representation of their similarity relations --- the closer two colors are in TCS, the more similar they are.
This is the view now adopted by many philosophers of perception.

One point of disagreement that is obscured by the foregoing is whether the elements of TCS that are ordered by the various ordering relations

under consideration are color appearances (construed as psychological states elicited by perceptual interaction with objects in the extramental world), or, on the other hand, colors themselves (construed as properties of extra-mental surfaces, lights, films, and the like).
(Reinhard Niederee takes them in the first way, for instance, in the bulk of his paper, though he clearly recognizes the distinction; Churchland takes them the second.)
It is certainly possible to take subjects' reactions to stimuli, and therefore the relations based on these, as revealing facts about those subjects' psychological states --- color appearances --- rather than about the distal stimuli.
But since the relevant psychological states are typically produced by distal stimuli, or are representations of them, the very same reactions can be treated as inducing ordering (and other) relations on the distal stimuli --- the colors.
Of course, depending on what explanatory work one wants TCS to perform, extending this organizational scheme to distal properties may or may not meet one's needs, and therefore not everyone is interested in thinking about TCS as organizing the space of colors.
Treating TCS as a space of color appearances is a weaker, more conservative option since it is compatible with, but does not require, extending the organizational scheme to colors.
On the other hand, this way of looking at things implies that perceivers possess awareness of similarity relations on their own perceptual experiences -- and many philosophers want to avoid this kind of reflexivity.
They prefer to say that in color vision we experience colors themselves as ordered in certain similarity relations.

TCS is a source of several further scientific and philosophical questions.

First, what is the full and correct description of TCS?
Does the above characterization exhaust the dimensionality of color experience?
What is the relationship of color dimensions to one another?
Are there privileged axes in TCS such that all the hues are in some non-conventional sense composed of values along these axes?
Scientifically, we want to understand the neural and psychological processes that underlie TCS in the visual system.
Physical objects such as lights and colored surfaces send light to our eyes.
The light that arrives at our eyes is processed by the brain and emerges somehow as experienced color.
Somewhere in this process it is stamped with the qualities that we

experience --- red/ green, blue/yellow, etc.
What are the relevant details of this process?
What features of the world does it reliably track or preserve?
Second, what is the nature of our knowledge of TCS and what it represents?
Is our grasp of TCS innate or learned?
Under what conditions ought one to say that it reveals reality?
What aspect of reality does it reveal?
What is the significance of the dimensions that order color-experience?
Hardin, in particular, has been vocal in arguing that since TCS does not reflect a physical ordering of colours (in the way that even Newton's circle did, since it preserved the rainbow-ordering, albeit with the two ends joined together).
Thus, he says, colour vision and color experience do not correspond to any objective feature of the world.
According to Hardin, TCS is an ordering imposed by sensory processing, and should simply not be taken to have any objective significance.
(Note, however, that as we explained above, saying that color-similarity is experiential is not to imply that the substrate on which it is imposed is also experiential.)

The papers collected together in this volume are all, in one way or another, about TCS and its significance.
Interestingly, it emerges from these papers that the traditional picture is itself under some stress.
The papers themselves fall under three broad (if somewhat overlapping) categories which we outline below.

II. Color and Structure: Current Views

The first category is about the ordering itself.
Is it a reasonable framework for describing experienced color?
If it is, what is the significance of the ordering with respect to the physical counterparts of color out in the world?
Some of the authors are content to operate within the traditional notion of a color-ordering: though they are ready enough to allow that it might be mistaken in some aspects, these authors are nonetheless confident enough of the broad outlines that they are willing to use it as a scaffold for constructing their ideas.
To take one example, Paul Churchland argues against the idea that all naturally experienced colors fall into TCS --- he shows how to generate certain experiences outside of its limits.
But he has no doubt that blue-yellow and red-green are the fundamental axes of TCS.

Others are skeptical: they argue that in one way or another the traditional conception of the color-ordering is misleading in fundamental ways.
Here are two examples.
Don MacLeod is worried by our lack of understanding of the neural basis of color phenomenology.
At one time, it was thought that TCS was a fairly straightforward outcome of "opponent processing" --- a hypothetical neural process that transformed the outputs of the three kinds of retinal cone cells, which underlie color vision, so as to maximize the capture of the information they provide.
MacLeod, however, does not think that things are that simple.
Kimberly Jameson is another sceptic.
Not so much in her contribution to this volume, but in her work at large, Jameson argues for a much greater role for culture and learning than traditional opponent processing theory is willing to concede.
Her emphasis on the acquisition and construction of color is considerably at odds with TCS, which could be taken as committed to the claim that color is the product of innate processes, and that every experienced color has the same component structure, regardless of cultural influences.

In his contribution, Rolf Kuehni gives a magisterial summary of color ordering conceptions, though he notes emphatically that these are all, in one way or another, mathematical abstractions, and that perceptual order is at best (but, significantly, often) on the ordinal level.
In other words: don't take TCS models too literally, but do be guided by them in a qualitative way.
The central problem concerning the sensory representation of color can be seen in this way.
The amount of information contained in light incident upon the eye is enormously greater than that which any organic color-vision system can analyse and represent.
These systems re-configure the information borne by light in a more compact form.
One way of looking at this reduction is in terms of dimensionality.
Any given light signal can be thought of as a spectral power distribution: an intensity of light in each of the wavelength intervals in the visible spectrum.
If, for example, we sample a signal at 10 nanometer intervals, we get 31 values between 400 and 700 nm, each independent of the others --- a 31 dimensional representation.
Color vision systems need to reduce this dimensionality: they possess neither a receptoral array sufficiently sensitive to small differences in such a representation, nor the processing capacity to deal with such quantities of information, nor indeed the need for such

exquisitely fine sensory response.

If we put the matter in this way, we can isolate two problem areas for a psychophysicist trying to describe sensory color-ordering.
First, how do you measure the dimensionality of color sensation?
How, in other words, do you map out a color-ordering given empirical observations of color-discrimination performance?
Second, what is the nature of the dimension-reduction between signal and percept?
The limited number of kinds of retinal cone cell is one obvious source of dimensional reduction.
Humans possess only three kinds, and each samples intensity over a broad interval (with varying sensitivity at different wavelengths).
This effects a huge reduction of information.
Is such a reduction strategic or circumstantial?
That is, is it a well-designed sampling that records a large amount of relevant information, throwing away what is not needed?
Or is it just leakage --- information simply being lost as if by a careless housekeeper sweeping the mess off his employer's desk?

Kuehni implicitly suggests that to a very great extent, the reduction is simply a loss: "Experiments have shown that the ability to 'measure' absolute luminance is very poor and relative luminance is, at best, estimated."
However, he seems to concede that color vision might be quite effective in normal ecological situations: "The full process used by the brain to assign a specific color experience to a given stimulus in a given visual field is as yet unknown, but it seems to be the result of relatively flexible interpretation of the [spectral power distributions] arriving at the eyes."

We noted before that the principal hue-axes of TCS induce an ordering quite different from that of the physical rainbow.
In addition to this, the dimension reduction just discussed means that signals that are distinguishable in a large number of dimensions might be indistinguishable in an informationally impoverished sensation.
This is the problem of metamerism
(metamers being physically distinct signals that are perceptually indiscriminable).
Paul Churchland says that "the apparently unprincipled diversity of metamers poses a genuine problem for a reductive account of objective colors."
Metamerism is a direct challenge to
Churchland's desire that such an account provide a one-to-one correlation between our internal (phenomenal) representation

of color and the set of physically distinct color signals.
He wants to argue, nonetheless, that "there is a way to construe the initially opaque space of possible reflectance profiles so that its structural homomorphism with human phenomenological space becomes immediately apparent ... thus the argument against color realism evaporates."

Churchland's strategy consists in finding a mathematical transformation of the 31 (or whatever) dimensional signal vector, the "CA-ellipse".
He claims that color sensation "tracks CA-ellipses".
In short, his claim is that there is a geometric construct in many-dimensional input space, such that color sensation corresponds to its smoothed-out approximation in a few-dimensional TCS.
Metamers do not constitute an "unprincipled diversity" according to Churchland; they are united by the fact that they share a CA-ellipse.
Color, then, is indeed a physically based ordering for Churchland: experienced color is the human color-vision system's best-shot approximation to a very rich reality, using CA-ellipses as an approximation tool.

How does TCS function in cognition?
Mohan Matthen takes TCS to be a semantic, or representational, vehicle by which visual sense orders real individual objects by color.
His argument turns on the kinds of knowledge we have of the colors as opposed to the kinds of knowledge we have concerning what color belongs to what object.
It is only recently that philosophers have distinguished the question 'What color is that object?' from questions of the type 'What is the nature of orange? is it, for instance, mixed or pure?'
Matthen argues that answers to the second kind of question --- questions about the colors themselves --- are known with "Cartesian" rather than merely "empirical" certainty, and that this can only be explained by supposing that we possess innate knowledge of the representational system that our perceptual systems use to denote the color-properties of environmental objects.
He takes it to be characteristic of innate concepts that they are acquired developmentally, not by learning.
The occasions that lead to the concept being activated do not epistemically or rationally constrain the result: a great variety of triggering experiences will lead to the same color ordering.
He also suggests that the notion of a dimensional ordering opens up an approach to solving a traditional philosophical problem: the color sorites.

Churchland and Matthen are optimists about TCS; they are not committed

to the details, but they do lean on the analysability of color experience into independent dimensions that are not spontaneously chosen by the perceiver.
Kuehni too grants the concept some validity and importance, though he demonstrates the multiplicity of theoretical approaches to the phenomenal structure of color.
Others, such as Reinhard Niederee, are less sanguine.

To understand the nature of Niederee's challenge, we need first to distinguish two ways in which the dimensionality of color can be taken.
In the description of TCS given above, we were assuming that TCS provides an answer to the following question: how shall we describe our phenomenal experience of the color of a thing?
In order to answer this question, we can ask people to match and discriminate things by color, and record how they respond: such experiments lead to a graphical representation of similarity relations something like TCS.
But we can also ask another question: how does color experience vary with changes in objects in our environment?
As it turns out, color experience depends not only on properties of the object to which color is attached, but also on the surroundings.
For example, the color you attribute (on the basis of perception) to a fruit you are looking at depends not just on the fruit, but also on its background and surroundings.
The fruit might, for instance, look luminously orange if it is placed in a dark surround, and somewhat flatter and yellower when it is placed in a light-colored surround.
Nieder,ee refers to this contextuality of color when he says, "there is a plethora of well-known spatial and temporal context effects (labeled as color contrast, assimilation, adaptation, and the like), which often even lead to 'new' colors not observed in isolated patches, such as (most shades of) brown or gray."
When you try to account for variations of color-in-context, Niederee shows, TCS is too weak an instrument.
Here is a simple example of what he has in mind.
The color brown marks a region in TCS.
Yet, it is a "contrast color": a color that is visible only when a contrast with other colors is available.
(White and black are contrast colors, available only when there are brightness differences in a scene.
Since brown is a blackish orange, yellow, or red, it too is dependent on brightness contrast.)
Nieder,ee analyses the change of color appearance of a sample light, as it is placed in surrounds differing in color and intensity.

He shows that a three-dimensional ordering cannot account for the change of color appearance in such a situation, even conceding trichromacy.
As he says, "this dimensionality result does not imply that color vision is based on four basic classes of retinal receptor types rather than three.
Instead, higher dimensionality of the kind considered here obviously is the consequence of context effects, . . . [which] turn out to be more complex than commonly assumed."

Rainer Mausfeld approaches TCS from a radically skeptical point of view.
He starts by asking, skeptically, how it could be possible that a series of physical processes starting from an object in the physical world and ending up with a conscious percept could, in Russell's words, "suddenly [jump] back to the starting point, like a stretched rope when it snaps" --- how in other words, can we coherently suppose that the culminating percept represents the original object, when everything in between is a transformation?
Mausfeld's own approach is to propose that the "Perceptual System" has a coding and conceptual system that is used to tie the organism to its environment. (Here, he is in broad agreement with the "Sensory Classification Thesis" of Matthen 2005, chapter 1 and passim.)
And he maintains that there is no such thing as "color per se".
One of the things Mausfeld means by this insistence is the denial of a certain strand of philosophical realism that takes the perceived world to be pre-sorted into types that are somehow recreated by the perceptual system --- the kind of realism that we find in Churchland, for example.
He is also out of sympathy with the opinion, dominant among contemporary philosophers of perception, that sensation has only "non-conceptual content" --- an opinion powerfully advanced by Gareth Evans (1982).
Mausfeld posits a radical plurality of color-type parameters.
Aperture-colors do not signify the same kind of feature or situation as surface colors; lustrous colors "cannot be reduced to or understood from physical considerations of physical material properties", and so on.
Additionally, in claiming that there is no such thing as "color per se", Mausfeld means to assert that the very construct of color is defined implicitly by the range of transformations carried out and uses within and outside the system to which it is put, and that it is not out in the world in this sense (just as is plausibly the case of, say, the syntactic property of being a control verb).

III. Color Spaces and Explanatory Spaces

A second broad class of questions concern the relationships between colors and color appearances, as putatively conceptualized by TCS (or some inheritor conception) and various other aspects of color perception.
These relationships matter because the significance of TCS lies largely in the explanatory projects relating to color perception in which it plays a role.
The papers in this section of the anthology are devoted to asking whether and how TCS can after all serve various (psychological, neural, metaphysical) explanatory projects for which it has been enlisted.

For example, Kimberly Jameson is interested in the ways in which colors are partitioned into categories by perception and cognition.
While it is hard to imagine that this sort of color categorization is completely independent of the organization imposed on colors by TCS, there is plenty of room for debate about just how large a role this factor plays in color categorization.
The flames of this debate were fanned by the celebrated results of Berlin and Kay (1969), which purported to demonstrate color categorization to be cross-culturally universal.
Many have thought that the best explanation of these findings (and subsequent studies in the same tradition, such as the 200x World Color Survey) lay in taking TCS to underlie color categorization.
For example, it is at least somewhat suggestive that the color categories corresponding to all four poles of the chromatic axes distinguished by TCS (viz., red, green, blue, and yellow) were found to be fundamental rather than derived.
And if, as many have held, TCS is itself more or less universal in adult human beings --- perhaps even fixed innately somehow in the human biological/cognitive/perceptual endowment, then this would explain the cross-cultural convergence in color categorization reported by Berlin and Kay and others.
In her contribution to the present volume, however, Jameson disputes the extent to which TCS can explain data about color categorization.
While she does not doubt that there is an ordering of color appearances in human beings, she does doubt its universality, its independence from cultural influence, and its ability to explain color categorization by itself.

Austen Clark is concerned with another role that TCS has traditionally been assigned in connection with phenomenal consciousness.
A view that finds broad acceptance amongst philosophers of perception and perceptual psychologists is that one cannot undergo the appearances that are organized by TCS without being in a conscious mental state.

That is, the standard view has it that there cannot be phenomenal appearances without perceptual consciousness.
However, Clark argues that there are explanatory benefits to be gained by rejecting the standard view, and instead holding that states found in early visual processing can have qualitative characters without being conscious.
This raises a further questions about what ingredients above and beyond qualitative character are necessary for a state to count as conscious; Clark's answer is that the state must be a locus of selective attention.

Just as with Jameson, it is not an aim of Clark's contribution to challenge the existence of TCS per se, though he implicitly challenges its connection to consciously experienced similarity.
Instead, Clark is questioning some natural and widespread assumptions about what it means for a state of color (or other) appearance to be phenomenal, and what else this implies about the perceptual psychology of the subject in whom the state occurs.
Thereby, he is challenging at least one of the explanatory roles traditionally assigned to TCS.

TCS finds another, quite different home in its use by some theorists to understand the metaphysics of color properties.
This is the topic Jonathan Cohen confronts.
Of course, there has been (increasingly in the last fifteen years) a large range of different proposals about how to understand the metaphysics of colors.
One important class of such proposals that Cohen is concerned to defend, and that has its roots in great modern thinkers such as Locke and Boyle, has it that colors are fundamentally constituted in terms of relations to subjects.
More specifically, relationalists typically hold that colors are what they are in virtue of the color appearances that they induce in subjects and that are organized by TCS.
The traditional (so-called 'physicalist') rival of views of this sort are those according to which colors are constituted independently of the appearances they induce in subjects, and instead in terms of their physical makeups.
The clash between relationalist and physicalist theories of color is not over whether there is such a thing as TCS.
Rather, it concerns whether TCS should be used as a basis for the ontology of color properties: relationalists think the relations colors bear to states organized in TCS are constitutive of colors, while physicalists think of these relations as interesting, and even important, but ultimately contingent features of colors --- viz., features that are not part of the essences of those properties.

(Recall that, as noted above, it is possible to construe TCS as ordering either just proximal color appearances or both proximal color appearances and distal color properties.
Because they construe the distal properties in terms of their relations to proximal appearances, relationalists get the extension of TCS to distal properties for free, as it were.)

While Cohen spends some time in his paper arguing for relationalism over physicalism, he devotes most of the paper to the project of defending relationalism from important criticisms put forward by Hardin and others and that are independent of physicalism.
These criticisms have it that, just because relationalism grounds colors in TCS, it results in a conception of color that is insufficiently objective and insufficiently unified.
Cohen responds to the worry about objectivity by arguing that we lack reasons to demand the objectivity of colors in any of the forms that color relationalism is unable to secure.
And he argues that, while there is a good sense in which relationalism does not provide for the unity of colors by itself, this is a reflection only of relationalism's schematic character: any more specific form of relationalism that would be counted an adequate theory of color will provide the ingredients needed to answer the unity demand.
(He does not address, however, the epistemic objections to color relationalism put forward by Matthen.)
Yet again, what hangs on Cohen's defense is not the existence or viability of TCS, but the viability of a certain use of TCS --- specifically, the use of TCS in order to understand the metaphysics of color properties.

In contrast, Don MacLeod raises a rather more fundamental challenge to TCS itself and its neural underpinnings.
MacLeod urges that our understanding of why TCS has the features it appears to are at least badly incomplete in several important ways.
For example, he argues that the standard three-receptor type explanation for trichromacy is empirically inadequate.
Moreover, he contends that there is no viable neurophysiological explanation for the psychological color primaries that form the poles of TCS (which is not to say that the textbooks fail to offer explanations!).
He urges that we are without an account of what neural mechanisms underlie the most basic of relations in terms of which TCS is constructed --- color discrimination.
MacLeod goes on to claim that the standard (receptor-based) understanding of the differences between color appearances in normal

and color-blind subjects make false predictions.
And he contends that we are almost completely ignorant about the neural basis for color constancy.

MacLeod's pessimism about TCS could be taken in different ways.
On the one hand, it could turn out that the rough structure of TCS is as we take it to be, even if its mechanistic underpinnings have so far eluded us; if so, then MacLeod's list of concerns might be read as a twenty-first century color science analog of David Hilbert's 1900 list of then-unresolved mathematical problems.
fn: Just for clarity: the David Hilbert who authored this list (the famous German mathematician) is distinct from the David Hilbert (the American philosopher) who is a co-author of chapter ? of the present anthology.
Alternatively, it could be that current views about the organization of TCS are themselves in need of serious revision.
Either way, MacLeod worries that there remain serious gaps in some of the most fundamental aspects of our understanding TCS.
And either way, MacLeod's unresolved questions can serve as warnings against premature triumphalism about our understanding of TCS, and also against overconfident appeals to TCS and its features in other explanatory projects such as those pursued in the remainder of this section of the anthology.

Jonathan Westphal takes up the problem of explaining how certain exclusionary relations that appear to structure the TCS (and that are at the heart of Wittgenstein's concerns in the Remarks on Colour) are true about afterimage colors.
As an example of one such exclusionary relation, consider the (now controversial --- see Crane and Piantanida 19xx) idea that, necessarily, nothing can be both red and green all over and at the same time.
Westphal proposes that the best explanation of why ordinary surfaces cannot be both red and green all over and at the same time lies in the idea that being red and being green require incompatible kinds of interactions with incident light.
Because the interactions in question are incompatible, no object can have both of them all over and at the same time; consequently, if being red and being red indeed require of things that they have such interactions, it is necessary that no object can be both red and green all over and at the same time.
Although Westphal is generally sympathetic to this explanation of the incompatibilities, he points out that it cannot be applied straightforwardly to afterimages, since there is no light incident on afterimages (the latter are not surfaces).
And yet, Westphal wants to insist, afterimages are colored, and the

exclusionary relations at issue are no less true of them: just as with surfaces, no afterimage can be both red and green all over and at the same time.
Westphal's question, then, is how to account for the relationships that appear to structure the TCS in afterimages.

Taken together, these articles underline the importance of the TCS by arguing about how (and in some cases, whether) it is related to a range of substantially disparate topics in color perception.
But they also show up tensions in and challenges to the traditional assumptions about the TCS.

IV. Color Blindness

A final pair of papers in the present volume debate the nature and extent of the differences between the color spaces of statistically normal trichromatic color perceivers, on the one hand, and those of perceivers with certain types of deficiencies --- specifically, red-green dichromats, on the other.
One way in which these matters interact with general questions about TCS is by addressing the way in which the general shape of TCS varies systematically with various kinds of statistically abnormal alterations.
These issues also shed light on questions about whether and to what extent our knowledge of TCS is mediated by certain kinds of experience that might be differentially available to different classes of observers --- and, therefore, on the putative innateness of this knowledge.

Byrne and Hilbert argue for a revised version of a more or less orthodox position that they call the "Reduction View," according to which the color space of red-green dichromats is a proper subset of TCS.
They contrast the Reduction View with the "Alien View," which has it that the color space of red-green dichromats and TCS have no non-trivial overlap --- rather, the elements of the former are just entirely different, or alien, with respect to TCS.
Byrne's and Hilbert's contention is that if dichromatic vision is a reduction of normal trichromatic vision, then the (Revised) Reduction View is correct.

Their main argument in support of the Reduction View depends on the Hurvich-Jameson-Hering-inspired idea that TCS is structured around just two chromatic opponent channels (red-green and yellow-blue).
They note that, on the standard understanding, the value of the red-green channel depends entirely on the difference between the

output of retinal L-cones and that of retinal M-cones, whereas the value of the yellow-blue channel depends on the output of three retinal cone-types (L-, M-, and S-cones).
But since red-green dichromats have no functioning M-cones (if deuteranope) or no functioning L-cones (if protanope), Byrne and Hilbert reason that they should have a functioning blue-yellow channel but no functioning red-green channel.
If that is true, then red-green dichromats should represent the chromatic properties of the world entirely through the contribution of their working yellow-blue channels, whereas normal trichromats represent the chromatic properties of the world through the contribution of both their yellow-blue channels and their red-green channels.
This leads directly to the Reduction View: red-green dichromats should represent a subset of the trichromat's TCS.

As noted, Byrne and Hilbert endorse (again, assuming dichromatic vision is a reduction of normal trichromatic vision) not the standard Reduction View, but a Revised version of it.
Their Revision comes from their contention that what the yellow-blue channel carries information about is not the determinable colors yellow and blue, but the super-determinable colors yellowishness and bluishness.
Adding this position to the mix turns Byrne's and Hilbert's Reduction View into an Alien View, because "normal trichromats never see this hue [yellowishness] without seeing more determinate hues like orange and yellow.
On the Revised Reduction View, a red-green dichromat sees yellowishness unaccompanied" (xx).

Justin Broackes thinks that this whole approach is wrong. In the first place, he throws considerable doubt on the experimental work that purportedly supports what Byrne and Hilbert call the Reduction View: most deuteranopes and protanopes show under laboratory conditions a substantial, but of course reduced, ability to discriminate red and green, especially in large fields of colour, and
many report seeing these colors, without finding them just a variety of either yellow or blue or grey. Secondly, he points out that the traditional view gives a poor account of how dichromats experience the colors they tend to confuse with one another. To put this in a somewhat technical way, the traditionalists construct "confusion lines" in trichromatic color spaces -- lines that connect the colors that dichromats supposedly cannot distinguish -- but do not notice that these confusion lines do not tell us much about how the confused colors look. Do all the colors that lie along a confusion line look to the dichromat the way that one of those colors look to the trichromat? (This would be Byrne and Hilbert's Reduction View.) Or do they look fundamentally different from any of these colors? (This would be the Alien View.) To Broackes, neither view is compelling. Firstly, though dichromats may be rather less reliable than trichromats, they simply do not in practice make the confusions indicated in these

diagrams when they are presented with large fields and with surface colours (as opposed to just small fields of light seen through the tube of a colorimeter). And secondly, though protanopes often find a red light at high intensity indistinguishable from a yellow light at much lower intensity, this does not rule out the possibility that the red at a reduced intensity might actually produce a sensation of a different hue, (e.g., red!). Thus, even if the confusion diagrams were correct (which they are only rather approximately, and for small stimuli), they would not guarantee that the range of sensation of a dichromat collapsed to two principal hues and the chromaticities lying between them.
Thus, the whole business of "confusion lines" is thus overplayed, Broackes argues.

Broackes is out of sympathy with the view that TCS is somehow constructed out of cone inputs, with the consequence that the three dimensional TCS of colour normals is replaced in cone-deficient perceivers by a two- or enhanced-two-dimensional system. From his point of view, TCS is specified (in humans at least) independently, one might say, of cone-cell inputs. As he sees the problem for color vision, then, it is not to construct a color space commensurate with the number of channels of input that cone cells provide. Rather, the problem is to make the discriminations needed to fill a color space given, to some extent, independently of receptors. Thus, dichromats, or at least many of them, do not end up with a fewer-dimensional color space: for instance lacking the red-green axis or with a radically different way of perceiving blue-yellow. Rather they start out with TCS (or some other three-dimensional color space), and fill it by making the requisite distinctions with less information, though perhaps with a coarser grain and less reliably.

How can dichromats fill up TCS if they lack information from a whole class of cone cells? Now, given the large overlap between the M- and the L-cones, it is not right away clear that this is even a good way of putting the question. Broackes, however, side-steps this point and attacks the question as posed. He invokes an interesting maxim -- that the information available to a reduced receptoral array may be enhanced by integrating information gathered over time and with movement of the perceiver. Consider the claim that distance information available to one eye is impoverished compared to that available to two. This is true instantaneously, but false over time, since all the parallax information available to two eyes at a moment is available to one eye over time, provided that this eye is permitted to move. Similarly, while the information gathered by a trichromat at a moment is a lot more color-sensitive than that which can be gathered by a dichromat, this deficit is mitigated over time. Consider the following: in the reddish light of evening, the sides of red objects that are turned toward the light are brighter than those of green objects of the same overall brightness. In general, the objects that gain most in lightness as the illuminant changes will be those that are closest in color to the new illuminant. Again, as objects turn towards or away from the light, differently colored objects will display different darkening patterns. In these and other ways, dichromats can gather almost all the information that trichromats gather, and they can use this information to array objects in a three dimensional color space.