Classes of Sensory Classification:
A commentary on Mohan Matthen, *Seeing, Doing, and Knowing* (Oxford: Clarendon Press, 2005)

Austen Clark
Department of Philosophy
103 Manchester Hall U-2054
University of Connecticut
Storrs, CT 06269-2054

Presented at an Author Meets Critics session during the meeting of the Pacific APA, at Portland Oregon on 24 March 2006.

Suppose that you, as the open-minded and tolerant person that you are, acknowledge that there exist animals besides persons who also have colour vision. Such tolerance is commendable, as it allows you to accept with aplomb the flood of findings about comparative color vision: the astoundingly good color vision of many birds and fish (including the tetrachromatic or perhaps pentachromatic color discriminations of pigeons), the rarity and weird distribution of trichromatism among mammals, its checkered history among primates, and so on. Unfortunately, your open-mindedness also leaves you vulnerable to a puzzle--a puzzle that seems to grow only deeper the more you puzzle over it. Do those other species see at least some of the same colours--the very same properties as--we see? If we humans, with our paltry trichromatic discriminations, sometimes perceive some colours veridically, then what are we to say of the chromatic sensory states of a pigeon, with its putatively pentachromatic prowess? Who gets the differences among the colours right: us, or them? To choose exactly one disjunct seems wrong; to say "both" seems wrong; to say "neither" seems wrong.

Mohan Matthen has thought long and hard over this and other problems in the philosophy of perception. The result is his wonderful book, *Seeing, Doing, and Knowing*. I think it gives the best available answer to the questions just posed. He calls it "Pluralistic Realism". It is complicated, and parts require some major, and counter-intuitive, readjustments in the way we think. But the exercize is well worth the effort.

My plan is to focus on one central and recurring theme: the classes used in sensory classification. I shall track this theme doggedly through the first four parts of Matthen's book. These four parts are: Classification, Similarity, Specialization, and Content. I will bark out problems as they are found. I'll say nothing about part five, sensory reference, because the disagreement Matthen and I have about that topic is already getting aired elsewhere, and it would be boring to repeat it all here.

I. Classification

The story starts with what Matthen calls "sensory classification". Part 1.a. of the "sensory classification thesis" reads:

> Sensory systems classify and categorize: they sort and assign distal stimuli (i.e. external sense objects) to classes. (Matthen 2005, 13)

It is very clear that, at least in the early parts of the book, these "classes" are classes of equivalent *appearance*. He says as much in

> **The Classificatory Equivalence Thesis**: Two stimuli present the same appearance in some respect if and only if they have been assigned to the same sensory class. (Matthen 2005, 31)

"The stimuli that are co-classified *appear* the same in some sensory respect" (Matthen 2005, 23). He also calls the appearances "sense features", defined as:

> a property a stimulus appears to have by virtue of an act of sensory classification ... (Matthen 2005, 14)

This is actually part II of the Sensory Classification thesis. The upshot is that when one is aware of two things looking the same, one is aware of one of the byproducts of sensory classification. As he puts it, "classificatory sameness and difference are expressed as phenomenal sameness and difference" (Matthen 2005, 90)

Now there is a very natural, plausible, initial, interpretation of the Sensory Classification Thesis which in the end I think is just wrong. I adopted it myself, at first, but eventually realized that it under-states Matthen's thesis. The idea, in a nutshell, is to take these two conditions as describing something like the CIE system of colorimetric specification. We read "membership in the same sensory class" to mean *nothing more than* "having equivalent effects on some determinative stage of sensory processing", in particular, for example, having equivalent effects on receptors. Talk of the sensory system "sorting, categorizing, and assigning" is understood as a somewhat colorful way of saying that different stimuli have different effects on those different stages.

What makes this potent is that we can use an understanding of those differences to predict when two stimuli, presented under controlled conditions, will look the same. In the retina we have a class of cones that have optimal sensitivity in the long wavelength portion of the spectrum--the "L cones", for short. Likewise, a second class optimally sensitive to short wavelengths (S), and a third, optimally sensitive in the middle (M). If two different stimuli both cause the same numbers of activations in each of the three cone systems of an observer, they will look the same to that observer. If they are to be discriminable, then one of them must produce some effect on one of the classes of cones that differs from that of the other. So these effects on the cones seem to be the basis in terms of which the sensory system does its sorting and classifying. And the classes in question are indeed classes of equivalent appearance: predicting whether the two stimuli look the same.

This is how the CIE system, or any system of colorimetric specification, works. The output is a prediction to the effect that, when presented under controlled viewing conditions, stimulus sigma and stimulus tau will match in appearance. To get there, we have to know the wavelength composition of each stimulus. Here a "stimulus" is a bundle of photons, entering the eye. We laboriously calculate effects on each of three classes of receptor, and then calculate sameness or difference with respect to "chromaticity coordinates". Results can (sometimes) yield a verdict on sameness of appearance. If that is sufficient for assignment to the same sensory class, then the verdict is of the form: stimuli sigma and tau are both assigned to class kappa. So is this what Matthen means by "sensory classification"?

Well, no. The crux: in this example, we the theoreticians choose what we mean by "having similar effects". The equivalence classes that are derived (sameness in chromaticity, etc) are products of our choice. We construct them because they provide a useful way for us to generate appearance standards. Whereas Matthen is claiming something much stronger than this. His claim is that the act of classification corresponds to something real--something real within the sensory system itself. The "act of classification" is some datable event that occurs someplace and sometime in early sensory processing; the sensory system itself "classifies", "categorizes", and "sorts" stimuli. So the system itself employs some representations of the classes into which the stimuli can be sorted. These are implications that Matthen endorses, and indeed, insists upon. He says

> Sensory awareness can be expressed in terms of a set of singular propositions, messages to the effect that a particular individual is assigned to a certain class, and is identified as exemplifying a certain property. (Matthen 2005, 14)

The generic form is: stimulus sigma is a member of class kappa. If sensory awareness can be "expressed" in such terms, then the capacities to identify individuals sigma and classes kappa are essential to the capacity of sensing itself.

And, in an ingenious twist, the label for the class is hypothesized to be the appearance that all the members of that class share. So the "look" of a color is the label for the class of stimuli that all look that way. Further, if one is aware of the appearance of a color, one is aware (per hypothesis) of a label identifying the class to which the stimulus has been assigned.

> the stimuli that are co-classified *appear* the same in some sensory respect. Suppose, for example, that two objects look crimson: this is the means by which the colour vision system delivers to us, or our epistemic organs, the news that it has co-classified them. ...The car in the parking lot looks red: this is how I know that my visual system has assigned it to the class of red things. (Matthen 2005, 23)

So here is an example. Look about the room at all your well-dressed colleagues in this wonderful discipline of ours of academic philosophy. You will notice patches of variously different colours and textures. Some parts of the panoply--if that's the right word--look the same colour; others differ to a greater or lesser degree. Pick two that look the same. According to Matthen, early on in sensory processing there were "acts of classification" in which those two patches were both assigned to the same class of stimuli. They were categorized as members of class kappa. So matching patches in the visual field betoken assignement to the same class. The visual field becomes a class act. Looks label the classes.

This is a bold and intriguing thesis, but it does face a prima facie difficulty. Call this problem number one. Not all chromatic perception seems to be categorical in the way that (say) perception of phonemes is categorical. Some prima facie counter-examples: perceiving the changes in the daylight over the course of a day; watching someone slowly blush; seeing colors of reflections, shadows, and glares; and in general, perceiving any optic array in which there are no sharp boundaries between patches of different color, or in which colors appear to characterize something other than material objects. Conveying these contents would take lots and lots of singular propositions of the form "stimulus sigma is a member of class kappa", and both the sigmas and the kappas would have to overlap.

II. Similarity

The need to employ overlapping classes of appearance leads us directly into the topic of part two: similarity. Here a significant disclaimer is attached to the Sensory Classification thesis. Namely: the thesis is an idealization, or an over-simplification, resting on a premise known to be false. The false premise is that the relations "same in appearance" (matches, indiscriminable, etc) are equivalence relations. Because they are all non-transitive, they are not. "Thus", says Matthen, "it is somewhat misleading to say *tout court* that the senses assign stimuli to classes." (Matthen 2005, 20)

> In the bulk of the book, sense-features are presented in this manner as classes constituted by equivalence or indiscriminability. It is important to realize, however, that such grouping is ultimately based on similarity to a certain degree. Similarity to a degree does not, strictly speaking, give us equivalence relations, and hence defining features in this way is a simplification. Thus, as Carnap (1928) realized, the proper accommodation of sensory similarity demands a complicated logical system. Here, for the most part, we operate with the simple logic of discrete classes. (Matthen 2005, 96)

Matthen scrupulously posts this warning at different places throughout the text (20, 96, 107, 230), and part II of the book is devoted to what he takes to be the literal, more generalizable truth underlying sensory classification, namely sensory similarity. Stimuli in (what is idealized as) one class are in fact "united by some degree of similarity" (107) or by "*resemblance,* which is a matter of degree." (20) To form a class, then, we must know what degree or resemblance is required, and in what respects. Matthen suggests that the classes are in fact fuzzy sets, and the logic of classification is fuzzy logic (107); but the implications of this are not explored. Note for example that fuzzy logic alone cannot answer the question of how much dissimilarity, and in what respects, is allowed among chromatic stimuli if they are all to be assigned to one colour class.

One consequence of removing the over-simplifying assumptions is that in part II the notion of "assigning an individual to a class" becomes more complicated. Instead of the classic binary extensional copula "is a member of", we must assess degrees of resemblance, in various respects, and have some criterion available, ranging over all of them combined, that provides us with a cut-off.

This yields another difference between part I and part II. In part I Matthen is robustly confident that sensory systems assign stimuli to classes. In part II, the expressions of how the sensory system itself forms these classes become markedly more guarded. For example, he says

> Sensory Ordering permits the constructions of more and less inclusive classes: the more inclusive classes comprise stimuli that may be less similar to one another than those in the less inclusive classes. It also allows the formation of many different cross-cutting sensory classes as the context dictates: for it may well be that what is relevant in a given situation is not some category delineated in advance... (20-21)

The important difference here: "permitting the construction" of a class, or "allowing the formation of a class", are very different from "assigning" stimuli to classes. Permitting and allowing differ from actually doing the deed. In part one the statements were less guarded: sensory systems do not merely *permit* or *allow* the formation of these classes: they actually do the dirty deed themselves.

What does the "act of classification" come to if one thinks in terms of sensory ordering? In early visual processing one finds feature maps that register more or less continuous values within distinct dimensions of discriminability, but that's not enough to constitute what Matthen means by "sorting", "categorizing" and "assigning stimuli to classes". So where in sensory processing do we have enough? I don't mean "where" physiologically, but where functionally: where in the functional organization does this act takes place? Call this problem number two.

III. Specialization

Part III, specialization, opens with the question: Does anything real, in the external world, correspond to these classifications? The question becomes acute in color vision, and particularly acute in the comparative context. If we humans erect human color vision to be the touchstone for what color vision is, then the pigeon, with its greater discriminatory capacity is systematically misperceiving the colors. It finds all sorts of illusory differences, and fails to see the obvious resemblances. Contrariwise, if pigeons could erect pigeon color vision to be the touchstone, then normal trichromatic humans misclassify all the chromatic stimuli, since they fail to see the obvious differences among the colors in the vast gamut they call "one" shade.

Matthen's answer improves on his earlier "Disunity of Color". It combines two theses that seem incompatible: Pluralism (both species could be right) and Realism (the color classes picked out by both species correspond to something real). The result is "Pluralistic Realism". Matthen says:

> we must accord to pigeon colors exactly the same authority as we have just ceded to the colours that humans experience. This need not violate logic. Properties--even physically specifiable properties--are infinitely numerous... Why then can both not be correct?

(Matthen 2005, 202)

One might think the answer to that question is obvious: pigeons and humans disagree about the similarities and differences of the colors. This answer presumes that both species represent the same properties, and disagree about some of their attributes. It is precisely this assumption that Matthen denies. Even though we call both sensory capacities "color vision", perhaps the only similarity between them is that they both make discriminations of wavelength independent of luminance. Different species can meet this condition even though they "might be converging on different properties of distal stimuli" when they use those capacities (Matthen 2005, 202). They both have "color vision" (in this broad sense); but the sensory systems that instantiate color vision in the two species can have distinct functions and distinct correctness conditions. This I think is the deeper meaning of the

> ***Disunity of Colour Thesis***. I. There is no ecologically characterized class of properties such that colour vision must consist (in whatsoever kind of organism it may occur) in the capture of some or all of the members of this class. II. There is no subjectively characterized class of properties such that colour vision *must* consist (in whatsoever kind of organism it may occur) in the capture of some or all of the members of this class. (Matthen 2005, 154).

Many versions of color realism address themselves to the question: What is the real nature of colour? Matthen's response, in effect, is: Why would you think that there is just one? The real nature of pigeon colours can be a totally distinct thing from the real nature of Old World monkey colours.

For this line to be viable, it must meet a formidable challenge. It must provide a schema for determining correctness conditions--determing what a veridical chromatic perception says about its subject matter--which can be applied successfully across species. We need this in order to show that something real corresponds to the sensory classifications each species employs. Otherwise the thesis does not successfully combine Pluralism with anything worthy of the name "Realism".

IV. Content

Matthen sees the challenge, clearly and distinctly. He proposes an answer that rests on two propositions.

The first is that the sensory class must be "physically specifiable": the extension of that class can be accurately described in the language of physics (199). This does not require that the class be some fundamental physical magnitude, named in some law of physics. Instead all we need is some way of picking out the class using that language. No limit is placed on how complicated that description might be. The "power to cause 10^8 excitations in my L wavelength cones" is a physically specifiable property. Likewise, something like the "power to inhibit my (L+S)-M opponent process system to degree delta" is also, according to Matthen, physically specifiable (see 142). He says:

> it is always possible to describe what a sensory system detects in physical terms; i.e. it is always possible physically to specify the response conditions of its various states.. (Matthen 2005, 199)

Can we *always* do this? Hold that thought; I will get back to it in a few moments.

The second part of Matthen's proposal provides correctness conditions for the act of assigning stimuli to that class. The search for these drives us into the doing part of sensing, doing, and knowing. How might we find them?

> An adequate answer to this puzzle depends on a proper understanding of specialized functions. Most perceptual categories are put to some use with regard to our functioning in the external world. Now, the utility of a perceptual category may be disrupted if the organism fails to respond to its customary delineation... (205)

In fact for Matthen these uses are "associated by evolution" with the category in question (233). He assumes that "most species are in evolutionary equilibrium most of the time" (206), and argues that *therefore*, "for most species now there is no mismatch between sensory capacities and innate sensorily guided activities" (206). We would then have a set of actions "associated by evolution" with the development of that capacity. These constitute what he calls "innate sensorily guided activities", "species activities", or "specialized functions". They are "innate" in the sense that they are either "genetically or developmentally specified" (205). So: the function of a given sensory category is to guide the agent's performance of some such class of innately specified actions. If we can pick out those actions, then we can derive correctness conditions for the category.

How? The answer is pretty simple: " An act of classification is wrong if it disrupts a specialized function that this act of classification is supposed to aid" (206). More fully we have the thesis that Matthen calls

> ***Action-Relative Realism****:* Sense experience *E* (or neural state *N*) represents distal feature *F* in the action-relative sense for a member of species *S* if (a) *F* is the physically specifiable response condition of *E* (or N) and (b) misclassifying things as F disrupts some innately or developmentally specified use to which E (or N) is put by members of species S. (Matthen 2005, 206)

Not all sensory states have a content (if they fail to serve a specialized function, then they do not); but if they do, their "Primary Sensory Content" is, always,

> that the situation is right for a certain action or actions, these actions having been associated with this state by evolution. These actions may include epistemic actions. (Matthen 2005, 233)

Let us sum up. The sensory system classifies stimulus sigma in class kappa. As a result, sigma looks a certain way, and that look is a label for kappa. Being aware of the look of sigma is being aware of a message that the sensory system posts: that sigma is in kappa. There are various innate activities enabled by the ability to register class kappa. The message that sigma is in kappa helps to guide the organism in the successful performance of those actions. If, proceeding on that guidance, the subsequent actions are disrupted, or fail, then the classification is incorrect. Otherwise, the actions succeed, the situation is therefore right for those actions, and so the classification is correct. We have our correctness conditions. Or at least, we have a schema for how to find correctness conditions.

V. Example

I say "a schema to find them" because, as Matthen notes, providing a real, live, actual set of correctness conditions requires us to know all sorts of things we currently do not know. We would have to know the biological functions of color vision, and at this point almost all the statements under this heading are almost entirely speculative. Nevertheless it is useful to see how the schema might work. So let us for a moment treat some speculations as if they are true, and see what happens.

Color vision has evolved independently in many different phylogenetic groups. Among mammals trichromatism is found only among primates, and in fact only among the old world monkeys (including us) among the primates. John Mollon and coworkers have suggested that among them trichromatism could be an adaptation to their arboreal and fructivorous life style. In particular, a red-green opponent system added to the older yellow-blue one would allow its recipient more easily to spot ripe fruits and berries in the foliage. I will call this the "ripe fruit" hypothesis. Matthen uses it in a couple of examples (208, 264). The first is a fruit that is ripe when it is yellow (208). The second example switches to mangoes, which are ripe when they are reddish orange. It is worth quoting at length.

> In the supermarket, a particular mango looks reddish-orange. I infer that it is ripe. I take it home, and it looks more yellowy-green than it had before. It turns out not to be ripe. I conclude that in the supermarket it looked different from the way it really was. My assumption here is that though the mango presented different colour-looks in different conditions, it retained the same colour property. And I might well conclude that in the supermarket, my senses told me something false: they assigned the mango to a category--the category I use for ripeness-inferences in mangoes--to which it did not in fact belong. In the supermarket, I attributed to the mango the colour property designated by the look it presented there. At home, I came to realize that I was wrong to do so. At home, I realized that it *really* had--remember the Fundamental Principle of Colour Attribution--the property designated by the yellowish greenish look it presents here, not the property designated by the reddish-look it presented in the supermarket. (Matthen 2005, 264)

Notice, first, that the mango presents one appearance in the supermarket, and it (the very same mango) presents a different appearance at home. Colour appearances are attributed to a persistent physical object. Second, by appeal to the correctness conditions, one can sort out which of those appearances are veridical, and which "say something false".

How do the correctness conditions work? The key idea is that the class of mangoes that look reddish orange is significant because that class helped guide our hungry fructivorus forebears to mangoes that are ripe. An appearance that "x is reddish-orange" is nonverdicial if, when x is a mango, it leads you to one that is unripe. This inference is reiterated three times in the quote just given:

> It turns out not to be ripe. I conclude that in the supermarket it looked different from the way it really was.
>
> I might well conclude that in supermarket, my senses told me something false: they assigned the mango to a category--the category I use for ripeness-inferences in mangoes--to which it did not in fact belong.
>
> In the supermarket, I attributed to the mango the colour property designated by the look it presented there. At home, I came to realize that I was wrong to do so.

It is not really reddish-orange, because it is not ripe. It merely looked reddish-orange.

The somewhat odd consequence is that correctness conditions for colours need not hinge exclusively on visible qualities. They can refer to any property necessary for the success of an action that can be guided by chromatic discrimination. Even within one species, there can be many different candidates. We get the "Disunity" of color, or Pluralism. When veridical, the classifications correspond to something real. The combination is Pluralistic Realism.

VI. Sensory semantic atomism

This is an immensely attractive view, and (one wants to say) something along these lines *has* to be right. Other species have colour vision. They live in different niches and are interested in different features found in those niches. They use their wavelength discrimination capacities to register information about those features. So the properties detected by colour vision might differ in different species, yet they could all have veridical perceptions of the features of interest to them. How could this fail to be on the right track?

Well, the devil is in the details, and in the remainder of the talk I will focus on three. They all yield problems in satisfying the conditions for "action relative realism".

First devilish detail (which gives us problem number three). Matthen commits himself to a particular strategy for filling out correctness conditions: that they will proceed class by class. The correctness conditions themselves are given the form

> Sense experience *E* (or neural state *N*) represents distal feature *F* if ...

and this gets filled out experience by experience, and feature by feature. Rember for example the correctness conditions for the specific reddish-orange sense feature of ripe mangoes. The result is a kind of "semantic atomism" for sensory classes: we can go class by class, and the meaning of each label, if it has a meaning, can be specified without mentioning any other classes. It resides in a "punctate" relation between that class and the world (specifically, the biological world).

I don't see any good reason to argue that semantic atomism is impossible to apply in this domain. But it does seem an unnecessary restriction on the potential strengths of Matthen's account. After all, the capacity to pick out these classes is not, and cannot be, granted in a punctate fashion. In order to discriminate a single pair--reddish orange from a green surround, say--one must, thereby, be granted systematic capacities to discriminate a vast range of different pairs. One either gets a whole swath of systematic discriminabilities (the red-green dimension, for example), or none at all. So why not formulate correctness conditions systematically, for an entire dimension of discriminability, rather than step by step, for the particular classes thereby discriminated?

I propose a friendly amendment. The account could be strengthened if correctness conditions attached not to particular perceptions as of membership in a particular class, but rather to systematic abilities to discriminate.

VII. Physical specifiability

Second devilish detail; problem number four. This will sound a little less friendly. I fear the notion of "physical specifiability" shifts in the course of the argument, in two ways. The members of the class to be specified go from proximal stimuli to distal stimuli. The copula goes from *looks P* to *is P*. Both have a big effect on the plausibility of supposing that the classes involved in sensory classification can be "physically specified".

Remember where we started. In the beginning, the classes are classes of equivalent appearance. A sense feature is a property a stimulus appears to have (14). Stimuli that are co-classified *appear* the same in some sensory respect (23). Two stimuli present the same appearance in some respect if and only if they have been assigned to the same sensory class (31). Those are all direct quotes!

Likewise, "physical specification" starts firmly grounded in the psychophysics of receptors. Recall that if two stimuli have equal effects on all three classes of cones in observer O, then those two stimuli will look the same to O. Notice that we are talking about proximal stimuli--two distinct bundles of photons entering the eye on separate occasions. This indeed is a well-tested, well-confirmed, and workable scheme; I mentioned that it is the basis of CIE color specification. So when Matthen says

> it is always possible to describe what a sensory system detects in physical terms... (Matthen 2005, 199)

a friendly reader will read this as a true assertion about psychophysics. Within well understood limits, colorimetric predictions really work.

But notice how far we have moved by the time we get to ripe yellow fruit and reddish-orange mangoes. At that point we are trying to specify correctness conditions for colors of persisting material objects. The mango that looked reddish orange in the store now looks yellowy green at home. We need a persisting physical object, one that can present different appearances at different times. Throughout the book Matthen applies the word "stimulus" indifferently both to proximal stimuli and to physical objects, and sometimes (as here) the distinction matters.

Think of "a" proximal stimulus as a bundle of photons entering the eye. Many of those photons are absorbed by the retina, never to be seen again. In fact, the *only* ones of interest are the ones that are absorbed. One might have a second stimulus with all the same properties as the first one--the same numbers of photons, at the same wavelengths--but those are different photons, and it is a different occasion. It is a numerically distinct stimulus. Proximal stimuli are not persistent entities; one cannot present the same one twice.

If a proximal stimulus is an occasion, then when we name a proximal stimulus, we thereby implicitly fix values for other viewing-condition variables that dramatically affect color appearance. The observer is the observer on whose retina those photons impinged, whose cone populations and densities are like *so*. The state of adaptation is the state of adaptation of the visual system in question. The illumination is the illumination current at that time. Same for the colours of the surround; the size of the visual angle; and its position relative to the fixation point.

If all those are given, then one can indeed frame a "physical specification" of the different bundles of photons that will be treated more or less equivalently by the visual system involved. But it is an entirely different matter if we think of a "stimulus" as an enduring physical object. We must now generalize over multiple occasions: the same object can reflect different bundles of photons, under different viewing conditions, into many different retinas. There is simply no way to determine whether this bundle of photons, reflected off of that object, presented under these conditions, to this retina, will look the same as that bundle of photons, produced under different conditions, and presented to a different retina. Color science cannot do it. Even if we had a "physical specification" of *all* the physical properties of two objects we still could not predict whether they will *look* the same. They might match when viewed in light from the northern sky, but not when in direct sunlight. Hold that variable constant. Still, even among normal observers, they will match for some observers, but not others. So, hold that variable constant too. Still, they might match if viewed in a two degree field of view, but not if they fill ten degrees. And so on.

This seems fatal, but there is an even bigger problem. (Like they say in bad horror movies, death is just the beginning!) It is clear that Matthen is assuming that we can specify, not just the physical objects that look yellow or that look reddish-orange, but rather the physical objects that *are* yellow, or that *are* reddish-orange. The copula changes, from *looks P* to *is P*. If this is right, it is very bad news, for the resources to do that job also cannot be found anywhere in color science. If *P* is a color term, color science stops with verdicts of the form "*looks P*". No sum of them can get us to verdicts of the form "*is P*". More is needed, and I don't see any place in color science that provides it.

Larry Hardin dubbed this the problem of "chromatic democracy": the same object can present appearances of many different colours on different occasions, and there is no scientifically respectable way of selecting some subset of those appearances as the veridical ones. At least, there is no way if the science in question is confined to color science. Perhaps the idea is that those correctness conditions require appeal to parts of science other than color science. To get to the ripeness of the fruit that makes the appearance of reddish-orange veridical, we appeal to biochemical and physiological specifications of the sugars and starches in a fruit, their digestibility in a particular fructivore, etc. Perhaps that would give a "physical specification". It seems somehow a desperate course; at the very least we cannot simply *assume* that it will succeed.

So, problem number four: Color science can provide physical specifications of *proximal* stimuli that, in controlled conditions, *appear* to be a particular colour. Matthen's account requires physical specifications of classes of *persisting* physical objects that *are* a particular colour. I don't know anyone who can provide those.

VIII. Task oriented taxonomies

Last problem. I doubt any critic will dispute the assessment that Matthen has corralled an enormous number of theses and arguments, and he gets them all moving down the trail together in a very compelling way. At the beginning, you find your attention riveted on the sensing part of sensing, doing, and knowing. At the end, you are focused intensely on the doing part of sensing, doing, and knowing. The doing part provides correctness conditions for the sensing part. Ah, how satisfying; it all fits together in the end. What a climax. A great read.

But then, us well-dressed philosophers, doing what we do, in our wonderful profession of professional academical philosophy, stand back and start to wonder. Are we guaranteed a good climax? Is the satisfying conclusion provided as a matter of necessity?

Part IV ends with a discussion of what Matthen calls the "task-oriented conception of colour" (265) and a "task-oriented taxonomy" (263). He says:

> The proposal is that we should construe the meaning of colour-looks by reference to the equivalence classes defined by instinctive use. (264-265)

Now, here's the problem: is there any reason at all think that "equivalence classes defined by instinctive use" are going to line up perfectly, or at all, with sensory classes as originally defined, by equivalence in appearance? These two ways of defining the classes are logically independent of one another. What guarantee can we give that they will pick out the same classes?

In the end I fear the notion of sensory classification gets pulled in two opposite directions; it has to meet two independent sets of demands. It serves two masters. If sensory classifications are those that psychophysics might specify, by specifying a few layers of interaction in early sensory processing, then it is plausible to think of them as physically specifiable. But if the categories have to guide action appropriately, then it seems unlikely that "looks yellow" will do; we need "is yellow" at least; and ultimately, as seen, such task-oriented terms as "is ripe for picking". Even if "looks yellow" is physically specifiable, there is no guarantee that it will match up perfectly with the latter, "task oriented" terms.

IX. Conclusions

To sum up, here are my worries about the classes used in sensory classification, listed in the order discussed.

1. I am not convinced that all color vision is categorical in the way Matthen implies. There seem to be lots of chromatic perceptual phenomena whose content cannot be cast in the form. (I)
2. It is worrisome that the classification thesis is an idealization based on a premise known to be false. Where in early sensory processing does this "assignment to classes" actually occur? (II)
3. The account would be strengthened if we dropped the assumption that there must be a one to one mapping from particular sensory categories to specialized functions. (VI)
4. It is one thing to provide a physical specification for all the proximal stimuli that appear to be the same colour. It is a very different matter to provide a physical specification for all the persistent physical objects that are the same color. No one knows how to do the latter. (VII)

Finally,

5. Classes formed with respect to equivalent effects on early processing need not match up exactly with equivalence classes formed with respect to instinctive use. Sensory classification is left to serve two masters, who might start pulling in opposite directions. (VIII)