BLACK AND WHITE AND COLOUR

PART II. LUMINANCE AND CHROMATIC PHENOMENOLOGY

V. The Phenomenology of Human Luminance Processing

In the first half of this paper, the central question was whether, in informational terms, trichromatic night vision is similar to looking at a monochrome photograph. There are many ways to parse the question ”what is it like to see an X?” and one of those ways asks about *what is seen*. The conclusion of the first half was that we do not see, via luminance vision, what we would see if we were given an image representing light intensity. This is why chromatic information is so useful, indeed why all diurnal mammalian species have both luminance and chromatic systems. Buried within that story was an assumption that whatever representations the luminance system derives—representations for motion, shape and so on—these properties are not encoded in luminance space. However luminance information itself is represented or registered, the implicit dimensions of relative luminance encoding do not survive its use. A biological motion system may recognize the gait of twelve close acquaintances, for example, but it does not represent them as in series of retinal images that assigns each ‘pixel’ in each image a value in intensity space—*or* in luminance space for that matter. That is not information that exists *anywhere* in the system.

It is this final conclusion that is usually denied by any believer in mental paint, be the ‘paint’ black and white or coloured. If the achromat sees cedar trunks in a dimly lit forest as dark upright columns, then the achromat sees them as *dark*. Whether that is intensity paint or luminance pigment matters not. If the tree trunk appears ‘darkly’, a dark upright expanse, then the achromat has a phenomenal image of the world that contains dark and light areas—dark and light mental paint.

There is an obvious question that springs to mind, especially if you are dubious about this phenomenal ‘substance’: where does a human mind ‘find’ mental paint? (And no, Home Depot is not an option.) (Sorry, more than almost anything, I like Home Depot jokes.) There is no direct mapping between the form of luminance representations and what we seem to see—at least not for any visual property apart for representation of albedo itself. This transformation of luminance representations into a phenomenal ‘image’ is what I will dub *phenomenal rendering*. I suspect that phenomenal rendering is not a good idea if we want to understand why the world appears to us as it does. Here I rehearse some of the familiar objections to rendering and add a few of my own.

i. *Rendering and General Luminance Processing*.

The phrase ‘phenomenal rendering’ is obviously a metaphor but it is an informative one. To render by drawing is to create a 2D image of a 3D world, one that, when observed, will look much as if you were viewing the original scene itself. “Rendering” in computer graphics’ terminology means much the same thing. Instead of starting with the world,

one starts with a model—usually a 'scene file' that defines objects in representational primitives according to their geometry, as well as specifying viewpoint, surface textures, the lighting of the scene. A rendering program then uses the contents of the scene file to create or render the 2D image, the final product of the graphics pipeline.[1] In many ways, rendering is the *reverse* process of what happens according to one class of models of albedo perception. Recall that according to one description of the problem, albedo processing begins with a *decomposition* of the Intensity Image into the Reflectance and Illuminance Images. Given the many complexities of natural lighting, the Illuminance Image is itself decomposable into further layers, varieties of non-uniform illumination, from shadowing and shading, but also inter-reflections, specular highlights, spotlights, and multiple lighting sources (cf. footnote 8 in Part I of this paper). In albedo perception, we start with an encoding of the 2D retinal image and 'decompose' the image into representations of the properties of the world; in graphics rendering, we start with formal descriptions of the scene, the scene file, and then render the 2D image according to those descriptions. (Indeed, one way to combine various elements of the 2D image is called 'compositing" in graphics terminology.) On this view of phenomenological rendering, then, visual processing involves two distinct steps: through luminance processing, the decomposition of the 2D image into encodings in a variety of representational spaces followed by the rendering of the scene files into the phenomenal 2D image composed of light and dark sensations.

The problems of rendering as a metaphorical explanation of visual experience are well known but bear repeating. For one, it is difficult to see what use a 2D 'image', thus rendered, would be to the visual brain—who would look at it or, more neutrally, what could be done with it? The purpose of luminance processing, from contrast encoding onwards, is to wring from the 2D image information about the visual scene, information that, so encoded or embodied, is *useful*. Once discerned, this information is presumably *used*. Graphics rendering is useful because we, the observers, want to see what certain scenes *would look like*. So, for example, once you have in hand a blueprint of your newly designed home, a plan that uses the standard architectural icons for blueprints, you can export the blueprints into a 3D architectural drawing system. The blueprint markings and measures are then re-interpreted in terms of the formal structures of the 3D program. Now rendering begins. You can see what the building would look like if it embodied the blueprint plans, plus the specifications for finishing detail such as the cladding, windows, lighting, and so on. You can look at it from the east or from the west vantage points, examine the lighting from various points in the living room, walk through the kitchen and notice that, in fact, you cannot open the refrigerator door because the kitchen island is positioned 30 inches in front of the refrigerator. In the case of human vision, however, there is no internal viewer who needs a 2D image in order to observe the world, no inner eye to view the re-created intensity images. Using the visual image to discern/use properties of the distal world is what luminance vision *does*. And once *that* trick is done—no mean feat—there is no reason why the visual image must be rendered, pixel-by-pixel, each point with a determinate intensity (or luminance) value. Few vision

[1] For anyone who is not familiar with the concrete realities of 3D rendering programs, there are various demonstrations to be found on the web. A good one, of a 3D architectural rendering program can be found at:LiveCad Video

researchers believe, then, that luminance vision either begins or ends with a rendered, neural encoding of the luminance image. On the input end, the retinal image is encoded in luminance contrast; at the output end, there are representations encoded (or abilities instantiated) in the dimensions of the appropriate space, of the perceived properties. This understanding of vision leaves no room or function for a process of rendering in human vision—a fortunate fact given its extraordinary complexity and high resource consumption.[2]

One other class of objections to visual rendering is worth mentioning. It has often been noted that the way in which we visually experience the world is strongly imagistic. For example, visual perception obeys the rules of perspective: perceptions are viewpoint sensitive (we see the world from where we stand, where our heads and eyes are pointing); objects such as trees in the distance subtend less of the visual field than do trees in the middle distance or foreground; parallel railway lines converge towards the vanishing point. This is certainly consistent with the "rendering" of intensity information into a final 2D image.[3] Then again, some of the properties of conscious perception do not lend themselves to the notion of imagery, a point that is brought home by psychophysical experiments. For example, by adjusting the parameters of a visual display for motion, psychophysicists can induce in observers the perception of motion direction without the perception of velocity—one sees the stimulus as moving but not as going any particular speed. That is not a state of affairs that can be rendered. In one famous experiment, Dianne Rogers-Ramachandran and Vijay Ramachandran (1998) used a stimulus display, dots on a gray background that changed in from white to black at regular intervals. The display was divided into two fields by the initial colour of the circles; one half of the display contained black circles, the other half began with white circles. Thus the fields where differentiated by whether they contained all black or all white circles. At slow speeds of reversal, it is easy to see the dots in each field reverse from black to white. As the speed of the reversal is increased, temporal summation of the signals eliminates the perception of the reversals. Both fields look like a flickering gray field, yet there is still a clear perceptual boundary between them. In other words, the stimulus produces a perception that there *is* a difference between two fields, without a perception of any particular difference between them. More generally, when we render a property, of necessity we thereby render its logical entailments. A rendered boundary—difference between two spaces—is the result of some determinate difference. But the neural *representation* of properties, discerned from luminance contrast, need not follow the same rules. The fault lines of luminance vision need not align with the joints of nature.[4]

---

[2] Daniel Dennett has put forward these objections more forcefully in many different places; the term "final product" of perception is Dennett's ("Seeing is Believing or is it?" 199x) and fits very nicely with the metaphor of the graphics pipeline which I am guessing he would find as funny as I do.

[3] In fact, Christopher Peacocke uses just this sort of example to show that there is a distinction between sensation and perception. On Peacocke's view, the pictorial aspects of experiences are the sensations of experience, because they cannot be parsed in terms of information about the external environment.

[4] The psychophysics of vision has made excellent progress on delineating the individual tasks of human vision, through the testing of the informational parameters of these processes—what we can see under exactly which conditions. These results have been illuminating but often counterintuitive. How the visual system divides up its representations of the world is not always how we imagine it should be divided up, given our usually seamless perceptions of the world's events and our understanding of the ontology of the

A very similar point can be made about determinacy and representation. Every property that is instantiated is a determinate property. If a hen has speckles, she has some determinate number of speckles, each in a determinate place. But when you see a speckled hen, you do not see each speckle, much less each speckle in a determinate place. This is true even though you represent the speckles on the hen *as* objective, as independent of our perceptions of them. A part of your perception of the speckled hen is that she has a determinate *number* of speckles; you also assume that, should you spend an evening mapping your hen's topography, you would find that each speckle occupies its own place. In other words, we represent the world *as* determinate but often we do so without representing the facts that make them determinate—e.g. each speckle and its place.

Summing up this section, we can say that while our visual phenomenology has certain similarities to 2D images, the metaphor of rendering is not a helpful one. It is hard to imagine what function rendering the 'scene files' of luminance processing could serve, the purpose of such a Herculean task. Moreover, certain central properties of visual experience do not seem to be render-able, represented by an image each pixel of which has a determinate intensity value. So if the achromat sees "in black and white", he does not do so because his visual representations have been rendered in luminance paint into a 2D image. More likely, luminance contrast information is used to discern or embody various properties, each in its appropriate representational space—and our experiences of the visual world supervene upon one subset or another of these neural events/states. Beyond that, I suspect, we have not much to say at this point—nor should say if we do not have any better ideas.

ii) *Rendering and Albedo Perception*. We come at last to the phenomenology of albedo perception, to seeing objects as light or dark. It is here, if anywhere, that philosophers and vision researchers have been tempted to say that phenomenological rendering *must* occur. Like the trichromat who sees colour *qua* a constant property of object surfaces, the achromat sees object surfaces as dark or light. Even if the achromat cannot perceive albedo, strictly speaking, the achromat sees the chalk as light, the coal as very dark. If pseudo-albedo perception is the achromat's version of colour perception, the achromat must have sensations that span the continuum from light to dark. Perhaps it is here, then, that the visual system of the rod achromat gains the 'what it is like' of seeing the world, each object rendered in the correct shade of lightness or darkness.

In the course of ordinary visual events, the rod achromat solves the problem of albedo—or more correctly of 'pseudo-albedo' given the achromat's lack of chromatic information. Still, from this position of disadvantage, the rod achromat manages to separate the

---

world. So, for example, it often comes as a surprise to people to learn how motion processing is 'sub-contracted', divided up into dissociable tasks. When you see a free throw in basketball, from the release of the ball to the satisfying swish of the net as the ball goes, the event appears continuous (and, of course, *is* continuous qua a physical event). But the processes that discern even the motion of the ball alone are discrete—onset of motion, acceleration, current velocity, direction of current motion, overall trajectory of motion qua missile, end point of trajectory, non-linear (circular?) motion as the ball circles the rim, to gravitational motion (as the ball falls to the floor), to name the processes we now know or strongly suspect.

contribution of the illuminant—the Illuminance Image—from the 'surface luminance' of objects in the scene. Looking at the scene depicted by Figure 10, the achromat distinguishes the shadowing or shading of the buildings from thier surface properties. He does not check to see whether the white 'glints' caused by specular reflectance are in fact bits of white paper floating on the canal or in the piazza. Instead, the achromat's mistakes will involve his failure to tell judge whether an object appears dark because of its surface reflectance or as a result of SPD of the illuminants (the street lights). If the 'colour' of the light changes, he will be surprised by the sudden change in appearance of the vase of flowers, of the blue hydrangeas mixed with deep red roses. And he will probably ask whether someone has turned on a new light. In other words, he will know that object surfaces are inherently light or dark but his ability to see albedo accurately will be compromised. In the same way that we trichromats, looking at the Carnovsky wallpaper under red lighting, will see all of blue figures as very dark, the rod achromat will see all red figures as very dark given the spectral sensitivity of the human rods and the SPD of natural starlight.

Now let us imagine that, in virtue of some focal brain lesion, the rod achromat looses the capacity to discern pseudo-albedo, much as a person can loose the ability to recognize faces or letters.[5] At a behavioral level, the rod achromat can no longer choose which of two papers looks darker, black construction paper or white copier paper, or say whether dark chocolate looks lighter than white chocolate. How would this neural deficit change the rod achromat's phenomenology? If you think that albedo perception is the basis of all luminance phenomenology, that representations of albedo are rendered into a phenomenal image of the world thus bringing about the appropriate sensations of lightness and darkness, then you will be hard pressed to answer this question sensibly. What is it like to see without the sensations of albedo perception? *Presumably not like anything at all*. If albedo processing is necessary for an achromat's visual phenomenology, and the achromat cannot process albedo, then the achromat lacks a phenomenology of vision. Yet an achromat without albedo vision would show almost no changes in behavior. Apart from her inability to say whether papers are light or dark, to correctly order paper samples according to their relative albedo, and perhaps to ponder her new deficit (if she notices it), the achromat's visually-guided behavior remains the same. She does not wake up, post-trauma, and cry out 'Help me, I'm blind!!' *Ex hypothesi*, none of the achromat's other visual abilities are compromised. Hence she will act, visually, just like any other rod achromat bar this one deficit. All evidence to the

[5] In fact there have been no reported cases of lost albedo perception in rod achromats, a unsurprising fact given how few rod achromats there are. There has been one case of a trichromat reported by Nijboer et al (2009) who lost albedo perception, as well as suffering a deficit in brightness perception for lights in a normal trichromat. The subject was a 66 year old female trichromat who had a subarachnoid hemorrhage. LZ had slight spatial neglect of the left side, but otherwise had normal language and memory; she was able to recognize objects as well as colours. However, when given two pieces of paper of different shades of gray, she could not say which one was darker. More strangely, she was unable to judge when a light was turned on or off, e.g. at night when walking into the bathroom. Nijboer et al. describe this a deficit in brightness perception or "brightness agnosia". I have not used this example because it is the only known case to date, which is suspicious. On theoretical grounds, I doubt that albedo perception would be a localized function or even a solely cortical function. Yet another reason to withhold judgment.

contrary, do we want to say that our achromat lacks *all* visual awareness? This is not a conclusion we want to embrace.

Let me put forward a more sophisticated argument against the rendering of albedo. We have said above that people can distinguish between the properties of the visual scene that arise from the illumination (shadows, shading) and those that arise from inherent surface properties of objects. To do so is to solve the problem of albedo, in formal terms to distinguish the Illuminance Image from the Surface Reflectance Image. However there is now good evidence that human trichromatic observers can also distinguish the intensity of light reflected from an image. [6] When looking at an image, human subjects can perceive the total amount of light reflected from an individual point in the scene, the Intensity Image. If subjects are given a complex Luminance Image and then told to concentrate on one of these three properties—illuminance, reflectance, or net intensity—they make correct judgments about each property. Changes to the illumination alone, to the lighting, for example, does not lead subjects to see a change in surface reflectance, in the inherent lightness or darkness of the objects. We do not make mistakes, *mutatis mutandis,* for the other properties either. So it would appear that, for the trichromatic viewer, brightness (of light), lightness (of objects) and the properties of illumination (shadowing, shading, spotlights) are *represented and perceived as distinct properties*. And there is no reason to think that rod achromats are any different even if they perceive only 'pseudo-albedo', an ability is less robust than trichromatic albedo perception.

The fact that we represent three separate properties is easy to gloss over when we look at the standard illustrations of the problem of albedo. Looking at Fig. 8 again, it takes a moment to realize that each of the cubes is rendered using the same grayscale or pigment, but, the grayscale represents *a different property* in each illustration. When we look at the three boxes, we immediately interpret the scenes as representing shading or surface colour or the reflected intensity as whole of the scene. Yet, when we look at the world, all three properties are present (or recoverable) in the same scene. This presents something of a mystery about the rendering of albedo information. If albedo is rendered in 'darkness' pigment, how do we distinguish phenomenal albedo from the other two properties? We see the results of illumination—shading and shadowing, highlights and spotlights—as areas of lightness and darkness as well. As good photographers, we can also see areas of an scene as being light or dark areas as a function of the total reflected light. These are distinct properties of the world, perceived by the trichromat *as distinct* from albedo. So how does one render *distinct perceptual properties* with only a single can of mental paint? There are a thousand shades of gray, as any interior designer will tell you. But are there different types of 'darkness' paint, some suitable for surfaces, others made for light, and still others for the photographic eye that sees total reflected light? If so, what makes them 'different' if not their phenomenal appearance?

It seems to me that something has gone wrong here. Instead of attempting to answer these questions, we should instead step away from all talk of rendering, as many

[6] Although this was assumed to be possible, it was not proven until 1993 when Arend and Spehar conducted a series of experiments.

philosophers have. Albedo perception is merely one computational process among the many, the results of which contribute to our phenomenology. We represent shape in however many shape-spaces the evolution of human has devised. Similarly, we represent the lightness of chalk in the representational primitives of albedo, whatever this might be. But we do not see the chalk as white because the representation of albedo has been rendered into a 2D image of the chalk in the correct shade of light gray or by having a correct *sensation* of any kind.[7] We experience shape, size *and* albedo, not by means of one sensation or another, but through a visual awareness of the property *as* a perceived property of the object. Of course there is a phenomenology of luminance vision. There is something it is like be a rod achromat who looks out the window and sees the dark trunks of trees amongst the boughs. But this is not explained by the metaphor of rendering, or by application of mental paint.

VI. Denouement

The goal of this paper was to make room for a different way to think about chromatic processing and colour perception. In a nutshell the view is as follows. You cannot explain what any particular luminance system does—how it makes vision possible for a species—without adverting to the spectral facts of its world. Nor can we explain how luminance processing works in the human trichromat without understanding the shared dance of chromatic and luminance processing. These are *paired* systems, each of which shores up the other, compensating for the other's deficiencies. In human vision, together they make possible a system for object vision, one that has high spatial, temporal and contrast resolution. We see the properties of objects that we do—virtually all of them—on the basis of both chromatic and luminance information. Twenty years ago, few neurophysiologists of vision would have seriously entertained this scientific hypothesis. (Peter Gouras is an exception I think). Now evidence in its favor is steadily accumulating. Certainly a healthy subset of colour vision researchers now assume something like this as part of their experimental paradigm. In fact, neuroethologists have been thinking about of chromatic information in this way for quite some time, as the medium, not the message, of vision as Robert Woodward once put the matter. So today, theories of this kind have the status of *empirical contenders*.

In what follows, I merely note the consequences of subsuming the principles of chromatic processing under the general principles of luminance processing as they have been discussed above. Many of these consequences are unintuitive; I suspect that some philosophers would hold that some are *a priori* false. I leave each reader to ponder what

---

[7] Of course, one could try to save rendering by claiming that all three properties *are* rendered in the same paint but they are "tagged" with the appropriate content, "Intensity", "Illumination Property" or "Surface Reflectance" in order to disambiguate them. But wouldn't the same reasoning apply to every other property depicted in phenomenal luminance image? We would need to "tag" all those properties, already discerned from the retinal image, wouldn't we have to tag the speckled hen with a label reading "Many speckles but not necessarily the same number or in the same position as represented here" or attach a pointer to the table with the message "square table top despite the oblique angle of the observer's sight line". We would appear to be right back to where we started. If the visual system has already discerned the various properties of the world, why would each property need to be rendered and labeled?

he or she believes about each one. Still, for any philosopher who thinks that one or more of the consequences *could not be true*, ask yourself this question: how could an *a priori* falsehood arise out of an seemingly 'normal' empirical theory?

Consider first the ontological and epistemic consequences of this view. If chromatic processing is like luminance processing, it follows that chromatic information is used, just as is luminance information, for discerning properties of the world. One of the properties perceive is albedo or surface reflectance; another of these properties is colour, surface spectral reflectance. In all likelihood, the trichromatic ability to see albedo and surface colour qua objective properties of the world is a product of both the chromatic and luminance systems, and whatever 'messy' informational mixture these two systems, working together, create. At present, there are many computational models of albedo perception, each of which makes different kinds of knowledge assumptions. None of these models is yet proven so there are not many things that we can say we *know* about this ability. Still, vision scientists agree on the definition, on the nature of the problem. They also agree that albedo is an objective property, albeit a dispositional one like fragility. They are now well aware of the difficulty of the computational problem. What you do so effortlessly—e.g. you just *see* your black cat curled up on your guest's white linen jacket—is a problem so complex that no artificial visual system can solve it without recourse to unrealistic simplifying assumptions. This is a problem has no local computational solution, based upon the intensity values of a discrete area of the image. And finally, the current view is that albedo perception requires information about many other, perhaps high level, properties of the visual scene. According to every class of models bar one, albedo perception requires that we first discern the properties of objects, both near and far, as well as the nature of the light within the scene (we must see shadows as shadows, shading as shading, and so on).[8] At least so far, albedo perception looks like a good candidate for *a cognitive and intentional perceptual process*. These hard-won conclusions of the vision community about albedo ought to hold for colour processing as well, which on this view, is the chromatic version of a 'luminance' property—surface spectral reflection as opposed to surface reflection.

For most readers, though, for those interested primarily in consciousness, the real 'kicker' of the view concerns trichromatic colour phenomenology. The section above, on the phenomenology of general luminance processing, ended with a rather benign conclusion: an achromat's conscious visual perceptions supervene on some proper subset of her luminance representations. That is how the achromat comes to distinguish twelve different friends and family members on the basis of each individual's gait. When we apply this same conclusion to chromatic processing, the conclusion looks less benign. We will still perceive bananas as yellow just as, at night, we perceive copier paper as light. But chromatic processing will also underlie the conscious perception of many other properties of the visual scene, properties we do not think of as being 'chromatic'. At least, this follows if chromatic and luminance processes are governed by the same principles. This would mean that there is also more to chromatic phenomenology than what we have supposed. The visual system of the human trichromat discerns much more

[8] The exception is the group of Spatial Filtering models in which information about such features affects the outcome of the perception but are not encoded per se. See Kingdom in press.

about the structure of the spectral world than the surface spectral reflectance of its objects, that an orange is orange. Indeed, assuming a very roughly hierarchical visual system, with all the feedback loops one could desire, the visual brain will understand the spectral lay of the land *prior to* seeing an orange as *orange*. For example, if an edge of the orange is made visible through chromatic edge processing alone (*pace* Hansen and Gegenfurtner), either because there is no luminance information that defines the edge or the luminance signal is too noisy, then we see that edge, chromatically defined. Now suppose that we see a shadow, one that falls across the orange, one lies across this chromatic edge. As the edge of a shadow, it is defined by means of luminance contrast alone (no chromatic cell responds). So the orange will be defined by a chromatic edge and the shadow that lies across it will be defined by a luminance edge. If we are conscious of *both* of those edges, chromatic and luminance—if we can see them both—there is no *a priori* reason why we should see those edges as being *of the same kind*, or, as our intuitions seem to shout, as 'black and white' edges. They are physically different, and these differences matter to how we perceive them. Similarly, a red apple pops out perceptually from among its background of green leaves (Mollon 1989). Yet this fact may have nothing to do with the perceived *colour* of the apple, our perception of it as *red*. Visual pop-out might result from the strong chromatic boundary between the apple and the foliage, a strong M-L chromatic contrast signal relative to a weak M+L (luminance) one. Of course, where red-green pop-out occurs, it seems likely that the perception of red and green soon follow. (By definition, pop-out grabs your attention!) Thus you see the red apples amongst the green leaves. Perhaps "colour" pop-out can occur *prior to the perception of colour*—here, the redness of the apple. Indeed, this seems likely if our beliefs about the difficulty of albedo perception apply equally to colour perception. In fact, by all accounts, the problem of colour perception is *more* difficult than albedo perception. So perception of chromatic contrast may well occur prior to colour perception.

Thinking in terms of a *chromatic phenomenology*, as opposed to a colour phenomenology, is rather strange. Whenever I look at some object in the world, I see the object as having a colour: my laptop case is red, the wall in front of me is grass green, the television screen is black. I don't see chromatic contrast or chromatic edges; nor do I see chromatically-derived biological motion *as* chromatically-derived biological motion. Or at least I am reasonably sure that I don't. This seems a good place to start a paper that explores chromatic phenomenology. For now, I leave the reader with two brief thoughts. First, it is becoming clear from other areas of neuroscience that our experiences of the world that seem 'homogenous' or 'seamless' or 'unified' are often the result of dissociable mechanisms. Just as, unbeknownst to us, the perception of "simple" trajectory motion is broken down into multiple parts by the visual system—into perceived onset, direction, acceleration, expected trajectory and endpoint—multiple chromatic mechanisms, including colour recognition, might underlie our perception of 'homogenous" coloured media. Second, it is not clear to me that our understanding of luminance phenomenology is on any firmer ground. You can train yourself to be more like Ansel Adams, to instantly understand which colour filter will heighten the rugged appearance of a cliff wall. But for the most part, *you* don't see luminance contrast or luminance edges, nor do you see luminance-derived shape *as* luminance-derived shape

and so on. As a trichromat, you can not delineate which visual experiences are derived solely from the luminance system at all. What we see is *the world*. And our concepts largely concern the events and properties of the world before us. So if your every intuition pushes you to claim that, *of course*, the results of general chromatic processing *must* be unconscious or they *must* be seen in 'black and white', you should ask yourself: Why do I think that? What do I think I know about visual processing and phenomenology that leads, inevitably, to that conclusion?

Finally, we come to the human perception of colour proper and the role—or lack thereof—of colour sensations. A predominant view in philosophy, since time immemorial, is of colour perception as essentially phenomenal. This is the view that unlike, say recognizing faces or seeing the slithering of a snake via biological motion processing, colour sensations are a necessary step in the causal chain of colour perception. To have a perception of an orange-ish orange requires that the piece of fruit cause in us an orange sensation, one that, after the mechanisms of colour constancy have been applied, allows us to see the orange as having a particular surface colour. In virtue of having the very sensation that you do, you see the fruit in the bowl as having a particular surface colour. Other philosophical assumptions follow. For example, suppose that all sensations are by definition consciousness events. If colour sensations are essential to colour perception, then colour perception is essentially a conscious process. For many philosophers, the phrase "unconscious colour perception" is an oxymoron. Similarly, if colour sensations are particulars, and particulars have determinate properties, seeing an orange as being coloured involves, as it were, a 'transfer' of determinacy to what is seen. You see the orange as having a determinate shade of colour—not yellow-orange, not red-orange, but an *orange-orange*, the determinate shade of the colour sensation. And again: assume sensations are transparent to the mind, such that we know all there is to know about them. By 'transfer', colours are transparent to the mind too. There is nothing we can learn about colour that is not there to be discerned in the sensation. These are all quite common views about the phenomenology and epistemology of colour perception proper.

If colour perception mirrors albedo perception, colour perception does not require colour sensations. If there is no "gray area' of the brain that houses the neural correlates of our 'black and white' phenomenology, there is no colour area, localized or distributed, that contains the neural correlates of colour sensations. None of the other conclusions about colour perception will follow either: colour perception can be unconscious, it does not 'present' us with determinate colours, and potentially there are many things about colour to find out by empirical investigation. Again, it is very hard to imagine what this means, the nature of the alternatives. But this does not mean the view cannot be true. In closing, I want to point out one fact (or what I take to be a fact): we are rarely 'fussed' by these same conclusions applied to luminance processing. At the retina, luminance signals are summed and compared to register chromatic contrast. This luminance 'currency' is then used to determine properties of the visual scene. One of those object properties, albedo (or 'pseudo-albedo' in the achromat), is represented in albedo space and the object surface in question is represented as having that albedo value. Moreover, there is no further step that renders this information into a 2D image of lightness. At no point are

'black and white' sensations a part of the process. Yet the achromat does have a conscious visual perception of objects as light or dark. This is what I find interesting. There are many papers published in neuroscience journals either claiming to have found 'the' colour area or doubting that the evidence to date proves that one location or another is the colour area. Other papers question whether a localized colour area is likely at all. But no one talks about the proverbial gray area. And that seems to me a good thing. After all, what makes us think there *is* a difference between black and white and colour?

Kathleen Akins
Simon Fraser University
Canada

BIBLIOGRAPHY

Arend, L. E. & Spehar, B (1993). Lightness, brightness, and brightness contrast: 1. Illuminance variation. *Percept. Psychophys. vol. 54 (4)* pp. 446-56.

Coa, D., Lee, B. B, & Sun, H. (2010). Combination of rod and cone inputs in parasol ganglion cells of the magnocellular pathway. *Journal of Vision vol. 10 (11)* pp. 4.

Dennett, D. C. (1996) Seeing is believing or is it? In Kathleen Akins, ed. *Perception: Vancouver Studies in Cognitive Science*, vol. 5: Oxford Univ. Press, pp. 158-172.

Gegenfurtner, K. (2003). Cortical mechanisms of colour vision. *Nat Rev Neurosci vol. 4 (7)* pp. 563-72.

Greenlee, M W, Magnussen, S, Nordby, K. (1988) Spatial vision of the achromat: spatial frequency and orientation-specific adaptation. *Journal of Physiology, vol. 395* pp. 661-78.

Hansen, T. & Gegenfurtner, K R (2009). Independence of color and luminance edges in natural scenes. *Vis. Neurosci. vol. 26 (1)* pp. 35-49.

Hansen, T., Olkkonen, M., Walter, S. & Gegenfurtner, K.R. (2006). Memory modulates colour appearance. *Nature Neuroscience*, *vol 9,* pp. 1367 – 1368.

Hess, R F, Nordby, K. & Pointer, J S (1987). Regional variation of contrast sensitivity across the retina of the achromat: sensitivity of human rod vision. *Journal of Physiology vol. 388* pp. 101-19.

Hubel, D H, & Livingstone, M S. (1988). Segregation of form, color, and stereopsis in primate area 18. *Journal of Neuroscience. Vol 7(11)* pp. 3378-415.

Kingdom, F A A (2003) Color brings relief to human vision. *Nat Neurosci. vol. 6 (6) pp.* 641-4.

Kingdom, F.A. (in press). Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research*.

Kingdom, FA A, Beauce, C. & Hunter, L. (2004). Colour vision brings clarity to shadows. *Perception vol. 33 (8)* pp. 907-14.

Kingdom, F A A, Kasrai, R (2006). Colour unmasks dark targets in complex displays. *Vision Research vol. 46 (6-7)* pp. 814-22.

Lind, O. & Kelber, A. (2009) The intensity threshold of colour vision in two species of parrot. *J. of Exp. Bio. vol. 212 (Pt 22)* pp. 3693-9.

Livingstone, M S & Hubel, D H (1984). Anatomy and physiology of a color system in the primate visual cortex. *J Neurosci vol. 4 (1)* pp. 309-56.

Livingstone, M S & Hubel, D H (1987). Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *J Neurosci. vol. 7 (11*) pp. 3416-68.

Livingstone, M S & Hubel, D H (1988). Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science vol. 240 (4853*) pp. 740-9.

Michna, M L, Yoshizawa, T, Mullen, K T (2006) S-cone contributions to linear and non-linear motion processing. *Vision Research vol. 47 (8)* pp. 1042-54.

Mollon, J.D. (1989). Tho' she kneel'd in that place where they grew..." The uses and origins of primate colour vision. *J. Exp. Bio. vol. 146* pp. 21-38.

Moskowitz, A., Hansen, R M, Akula, J D, Eklund, S E, & Fulton, A B. (2009). Rod and rod-driven function in achromatopsia and blue cone monochromatism. *Investigative Ophthalmology & Visual Science vol. 50 (2)* pp. 950-8.

Nijboer, T. C. W., Nys, G.M S, van der Smagt, M. J., & de Haan E.F. (2009) A selective deficit in the appreciation and recognition of brightness: brightness agnosia? *Cortex vol. 45 (7)* pp. 816-24.

den Ouden, H E M, van Ee, R. & de Haan, E H F (2005). Colour helps to solve the binocular matching problem. *The Journal of Physiology vol. 567 (Pt 2*) pp. 665-71.

Peacocke, C. (1983) Sensation and the content of experience. Chapter 1 of *Sense and Content: Experience, Thought and Their Relations.* Oxford University Press.

Rogers-Ramachandran D. C. & Ramachandran, V. S. (1998). Psychophysical evidence for boundary and surface systems in human vision. *Vision Research, vol. 38 (1)* pp. 71-7.

Sharpe, L T, Whittle, P, Nordby, K. (1993) Spatial integration and sensitivity changes in the human rod visual system. *Journal of Physiology ,vol. 461* pp. 235-46.

Shimono, K, Shioiri, S, Yaguchi, Psychophysical evidence for a purely binocular color systemH (2009). *Vision Research vol. 49 (2)* pp. 202-10.

Zaidi, Q & Li, A. (2006). Three-dimensional shape perception from chromatic orientation flows. *Visual Neuroscience vol. 23 (3-4)* pp. 323-30.